ETH zürich

ETH Library

A note on the computational complexity of the moment-SOS hierarchy for polynomial optimization

Conference Paper

Author(s):
Gribling, Sander; Polak, Sven; Slot, Lucas

Publication date:
2023

Permanent link:
https://doi.org/10.3929/ethz-b-000627799

Rights / license:
Creative Commons Attribution 4.0 International

Originally published in:
https://doi.org/10.1145/3597066.3597075

This page was generated automatically upon download from the ETH Zurich Research Collection.
For more information, please consult the Terms of use.

A note on the computational complexity of the moment-SOS hierarchy for polynomial optimization

Sander Gribling
gribling@irif.fr
IRIF, Université Paris Cité
Paris, France

Sven Polak
s.c.polak@tilburguniversity.edu
Tilburg University
Tilburg, The Netherlands

Lucas Slot
lucas.slot@inf.ethz.ch
ETH Zürich
Zürich, Switzerland

ABSTRACT

The moment-sum-of-squares (moment-SOS) hierarchy is one of the most celebrated and widely applied methods for approximating the minimum of an n-variate polynomial over a feasible region defined by polynomial (in)equalities. A key feature of the hierarchy is that, at a fixed level, it can be formulated as a semidefinite program of size polynomial in the number of variables n. Although this suggests that it may therefore be computed in polynomial time, this is not necessarily the case. Indeed, as O'Donnell [16] and later Raghavendra & Weitz [20] show, there exist examples where the sos-representations used in the hierarchy have exponential bit-complexity. We study the computational complexity of the moment-SOS hierarchy, complementing and expanding upon earlier work of Raghavendra & Weitz [20]. In particular, we establish algebraic and geometric conditions under which polynomial-time computation is guaranteed to be possible.

KEYWORDS

moment-SOS hierarchy, sums of squares, moments, computational complexity, polynomial optimization, semidefinite programming

ACM Reference Format:
Sander Gribling, Sven Polak, and Lucas Slot. 2023. A note on the computational complexity of the moment-SOS hierarchy for polynomial optimization. In *International Symposium on Symbolic and Algebraic Computation 2023 (ISSAC 2023), July 24–27, 2023, Tromsø, Norway.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3597066.3597075

1 INTRODUCTION

Consider the polynomial optimization problem:

$$\begin{aligned} f_{\min} := \min\ & f(\mathbf{x}) \\ \text{s.t. } & g_i(\mathbf{x}) \geq 0 \quad (1 \leq i \leq m), \\ & h_j(\mathbf{x}) = 0 \quad (1 \leq j \leq \ell), \\ & \mathbf{x} \in \mathbb{R}^n, \end{aligned} \tag{POP}$$

where $f, g_i, h_j \in \mathbb{R}[\mathbf{x}]$ are given n-variate polynomials. The feasible region of (POP) is a basic *semialgebraic* set, which we denote by:

$$S(\mathbf{g}, \mathbf{h}) := \{\mathbf{x} \in \mathbb{R}^n : g_i(\mathbf{x}) \geq 0,\ h_j(\mathbf{x}) = 0\}.$$

This work is licensed under a Creative Commons Attribution International 4.0 License.

ISSAC 2023, July 24–27, 2023, Tromsø, Norway
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0039-2/23/07.
https://doi.org/10.1145/3597066.3597075

Problems of the form (POP) are generally hard and non-convex. They naturally capture several classical combinatorial problems, and have applications in finance, energy optimization, machine learning, optimal control and quantum computing. As they are often intractable, several techniques have been proposed to approximate them. Perhaps the most well-known and studied among these techniques is the so-called *moment-SOS hierarchy*, due to Lasserre [12] and Parillo [17]. The main idea behind the hierarchy is that one can certify the nonnegativity of a polynomial $p \in \mathbb{R}[\mathbf{x}]$ on $S(\mathbf{g}, \mathbf{h})$ by representing it as a *weighted sum of squares*:

$$p(\mathbf{x}) = \sum_{i=0}^{m} g_i(\mathbf{x})\sigma_i(\mathbf{x}) + \sum_{j=1}^{\ell} h_j(\mathbf{x})p_j(\mathbf{x}), \tag{1}$$

where $\sigma_i \in \Sigma[\mathbf{x}]$ are sums of squares, $p_j \in \mathbb{R}[\mathbf{x}]$ and we set $g_0(\mathbf{x}) = 1$ for convenience. We say that a representation (1) is of degree t if $\deg(g_i\sigma_i) \leq t$ and $\deg(h_jp_j) \leq t$ for all i, j. For $t \in \mathbb{N}$, one then obtains a lower bound $\mathrm{sos}(f)_t \leq f_{\min}$ on the minimum of f by:

$$\mathrm{sos}(f)_t := \sup_{\lambda \in \mathbb{R}} \big\{\lambda : f - \lambda \text{ has a representation (1) of degree } 2t\big\}. \tag{SOS}$$

For fixed level t, this lower bound may be computed by solving a semidefinite program (SDP) involving matrices of size polynomial in n. It is often claimed that one may therefore (approximately) compute $\mathrm{sos}(f)_t$ in polynomial time, for instance by applying the ellipsoid algorithm. As was noted by O'Donnell [16] and later by Raghavendra & Weitz [20], this is not necessarily the case. Indeed, polynomial runtime of the ellipsoid algorithm is only guaranteed when the feasible region of the SDP contains an *inner ball* which is not too small, and is contained in an *outer ball* which is not too large. Informally, these two balls ensure that it is possible to choose the coefficients of the multipliers σ_i, p_j in the representation (1) so that their *bit-complexity* is polynomial in n. We call such a representation *compact*. Roughly speaking, the following mild algebraic boundedness assumption on the feasible region $S(\mathbf{g}, \mathbf{h})$ guarantees the existence of the inner ball for the SDP-formulation of $\mathrm{sos}(f)_t$.

Definition 1. We say that the problem (POP) is *explicitly bounded* if $g_1(\mathbf{x}) = R^2 - \|\mathbf{x}\|_2^2$ for some $R \geq 0$.

Remark 2. If (POP) is explicitly bounded, then (SOS) is feasible for all $t \geq \lceil \deg(f)/2 \rceil$.

Remark 3. Explicit boundedness is a slightly stronger assumption than the usual Archimedean condition, which merely requires $R^2 - \|\mathbf{x}\|_2^2$ to have a representation of the form (1). For example,

the semialgebraic set defined by $x^2(1-x^2) \geq 0$ is Archimedean[1] since we have the identity $1 - x^2 = (1-x^2)^2 + x^2(1-x^2)$, but it clearly does not have a constraint of the form $R^2 - x^2$ in its description. Likewise, we point out that a compact semialgebraic set is not necessarily Archimedean (see, e.g. [19, Ex. 6.3.1], [13, Ex. 3.19]).

The remaining question, then, is whether an outer ball always exists. O'Donnell [16] shows that in fact, it does not; he constructs an example where every representation (1) of $f(\mathbf{x}) - \mathrm{sos}(f)_2$ necessarily involves multipliers σ_i, h_j whose coefficients are doubly-exponentially large in n. Raghavendra & Weitz [20] subsequently show that it is possible to construct such an example even when the equalities $\mathbf{h}$ include the boolean constraints $\mathbf{x}_i - \mathbf{x}_i^2 = 0$, negatively answering a question posed by O'Donnell [16]. On the positive side, they show conditions under which existence of a compact representation (1) is guaranteed. These conditions are met for the reformulation of several well-known combinatorial problems as a (POP), as well as for optimization over the unit hypersphere. To state our and their results, we make the natural assumption that the coefficients of the objective function f and the polynomials g_i, h_j defining the feasible region $S(\mathbf{g}, \mathbf{h})$ of (POP) have polynomial bit-complexity.

Assumption 4. Throughout, we assume that the coefficients of the polynomials f, g_i, h_j in (POP) have polynomial bit-complexity in n and that their degree is independent of n. We also assume that the number of constraints $(m + \ell)$ is polynomial in n.

Theorem 5 (Main positive result of [20], paraphrased). *Let $S(\mathbf{g}, \mathbf{h})$ be a semialgebraic set and let $t \geq \lceil \deg(f)/2 \rceil$ be fixed. Assume that $S(\mathbf{g}, \mathbf{h})$ is explicitly bounded: $g_1(\mathbf{x}) = R^2 - \sum_{i=1}^n \mathbf{x}_i^2$ for some $1 \leq R \leq 2^{\mathrm{poly}(n)}$. Suppose furthermore that the following conditions are satisfied:*

(1) *For any $p \in \mathbb{R}[\mathbf{x}]_t$ with $p(\mathbf{x}) = 0$ for all $\mathbf{x} \in S(\mathbf{g}, \mathbf{h})$, there are $p_1, p_2, \ldots, p_\ell \in \mathbb{R}[\mathbf{x}]$ such that:*

$$p(\mathbf{x}) = \sum_{j=1}^{\ell} p_j(\mathbf{x}) h_j(\mathbf{x}),$$

and $\deg(p_j h_j) = O(t)$.

(2) *Let μ be the uniform probability measure on $S(\mathbf{g}, \mathbf{h})$. The* moment matrix *$M_t(\mu)$ defined by:*

$$M_t(\mu)_{\alpha,\beta} := \int_{S(\mathbf{g},\mathbf{h})} \mathbf{x}^{\alpha+\beta} d\mu(\mathbf{x}) \quad (\alpha, \beta \in \mathbb{N}_t^n)$$

has smallest non-zero eigenvalue $\geq 2^{-\mathrm{poly}(n)}$.

(3) *There exists an $\eta \geq 2^{-\mathrm{poly}(n)}$ such that $g_i(\mathbf{x}) \geq \eta$ for all $\mathbf{x} \in S(\mathbf{g}, \mathbf{h})$ and $1 \leq i \leq m$.*

Then the program (SOS) has an (approximately) optimal solution involving only multipliers σ_i, p_j whose coefficients are at most $2^{\mathrm{poly}(n)}$.

The conditions of Theorem 5 have a natural interpretation in the dual formulation of (SOS), the *moment formulation*, which reads (see, e.g., [5]):

$$\mathrm{mom}(f)_t := \inf_{L \in \mathbb{R}[\mathbf{x}]_{2t}^*} \big\{ L(f) : L(1) = 1, L(p) \geq 0 \text{ for } p \text{ with a representation (1) of degree } 2t \big\}. \quad \text{(MOM)}$$

We refer to Section 3.3 for a detailed discussion of (MOM) and the duality between it and (SOS). For now let us introduce two types of matrices associated to linear functionals $L \in \mathbb{R}[\mathbf{x}]_{2t}^*$: the *moment matrix* $M_t(L)$ and, for a polynomial g, the localizing matrix $M_t(gL)$. Their entries are given by

$$M_t(L)_{\alpha,\beta} = L(\mathbf{x}^{\alpha+\beta}) \quad (|\alpha|, |\beta| \leq t), \tag{2}$$

$$M_t(gL)_{\alpha,\beta} = L(g(\mathbf{x})\mathbf{x}^{\alpha+\beta}) \quad (|\alpha|, |\beta| \leq t - \lceil \deg(g)/2 \rceil). \tag{3}$$

For ease of notation, we use a subscript t for the localizing matrix $M_t(gL)$ even though it is indexed by monomials of degree at most $t - \lceil \deg(g)/2 \rceil$. Here we use terminology associated to measures since the linear functionals L in (MOM) can be viewed as relaxations of probability measures: any probability measure μ whose support is contained in $S(\mathbf{g}, \mathbf{h})$ gives rise to a feasible L_μ defined by

$$L_\mu(p) = \int_{S(\mathbf{g},\mathbf{h})} p(\mathbf{x}) d\mu(\mathbf{x}) \text{ for } p \in \mathbb{R}[\mathbf{x}]_{2t}. \tag{4}$$

We can then interpret the conditions in Theorem 5 as follows: (2) says that $M_t(L_\mu)$ has smallest non-zero eigenvalue $\geq 2^{-\mathrm{poly}(n)}$ and (3) implies that the same holds for all localizing matrices $M_t(gL_\mu)$ with $g \in \mathbf{g}$.

1.1 Our contributions

The main goal of this paper is to carefully map under what circumstances computation of the bounds $\mathrm{sos}(f)_t$ (and/or $\mathrm{mom}(f)_t$) and the corresponding representation (SOS) is guaranteed to be possible in polynomial time.

Our starting point will be the following proposition which concerns the bounds $\mathrm{mom}(f)_t$. The proof of this proposition consists of a straightforward reformulation of the conditions of strict feasibility for the SDP formulation of $\mathrm{mom}(f)_t$ for explicitly bounded polynomial optimization problems, see Section 3.4. As we will see there, the conditions (1) and (2) in Proposition 6 below will guarantee the existence of a ball in the feasible region of (MOM), whose radius depends on the smallest non-zero eigenvalue of the (localizing) moment matrices corresponding to a feasible solution L.

Proposition 6. *Let $S(\mathbf{g}, \mathbf{h})$ be a semialgebraic set and let $t \geq \lceil \deg(f)/2 \rceil$ be fixed. Assume that $S(\mathbf{g}, \mathbf{h})$ is explicitly bounded: $g_1(\mathbf{x}) = R^2 - \sum_{i=1}^n \mathbf{x}_i^2$ for some $1 \leq R \leq 2^{\mathrm{poly}(n)}$. Suppose furthermore that there exists an $L \in \mathbb{R}[\mathbf{x}]_{2t}^*$ with $L(1) = 1$ and the following properties:*

(1) *For any $g \in \mathbf{g}$ and any $p \in \mathbb{R}[\mathbf{x}]_{t - \lceil \deg(g)/2 \rceil}$, if $L(gp^2) = 0$, then there are $p_1, p_2, \ldots, p_\ell \in \mathbb{R}[\mathbf{x}]$ such that:*

$$gp^2(\mathbf{x}) = \sum_{j=1}^{\ell} p_j(\mathbf{x}) h_j(\mathbf{x}),$$

and $\deg(p_j) \leq 2t - \deg(h_j)$ for each $j \in [\ell]$. We recall that $1 \in \mathbf{g}$ by convention.

(2) *The matrices $M_t(L)$ and $M_t(gL)$ ($g \in \mathbf{g}$) have smallest non-zero eigenvalue at least $2^{-\mathrm{poly}(n)}$.*

[1] Technically, both the Archimedean and explicit boundedness condition are properties of the description of a semialgebraic set and not of the set itself. However since we always fix the description, we often write 'the set $S(\boldsymbol{g}, \boldsymbol{h})$ is explicitly bounded'.

Then for $\varepsilon \geq 2^{-\mathrm{poly}(n)}$, the bound $\mathrm{mom}(f)_t$ *(which equals* $\mathrm{sos}(f)_t$*) may be computed in polynomial time in n up to an additive error of at most ε.*

The statement of Proposition 6 is stronger than the result of Raghavendra & Weitz in the sense that it guarantees polynomial-time computation of the bound $\mathrm{sos}(f)_t$, whereas Theorem 5 only guarantees existence of a compact representation (1). Furthermore, its conditions do not require *strict positivity* of the inequality constraints $\mathbf{g}$ on $S(\mathbf{g}, \mathbf{h})$. As we see below, it therefore applies to several natural settings where Theorem 5 may not be applied. On the other hand, the first condition of Proposition 6 is more restrictive than the first condition of Theorem 5. We note that it is satisfied for example when L is the linear operator associated to a positive Borel measure supported on $S(\mathbf{g}, \mathbf{h})$, and the constraints $\mathbf{h}$ form a Gröbner basis of a real radical ideal (cf. [13, Sec. 2]).

Our first contribution is a sufficient condition for the second requirement of Proposition 6.

Theorem 7. *Let $S(\mathbf{g}, \mathbf{h})$ be an explicitly bounded semialgebraic set with $R \leq 2^{\mathrm{poly}(n)}$ and let $L \in \mathbb{R}[\mathbf{x}]_{2t}^*$ be a feasible solution to* $\mathrm{mom}(f)_t$*. Assume that $L(\mathbf{x}^\alpha) \in \mathbb{Q}$ has polynomial bit-complexity for all $\alpha \in \mathbb{N}_{2t}^n$. Then the smallest non-zero eigenvalue of $M_t(L)$ is at least $2^{-\mathrm{poly}(n)}$ and the same holds for the localizing matrices $M_t(gL)$ for $g \in \mathbf{g}$.*

Our second contribution is an alternative, *geometric* condition on the feasible region $S(\mathbf{g}, \mathbf{h})$ of (POP) which guarantees polynomial-time computation of $\mathrm{sos}(f)_t$ in the special case where the formulation does not contain any equality constraints. We write $S(\mathbf{g})$ for $S(\mathbf{g}, \emptyset)$.

Theorem 8. *Let $S(\mathbf{g}) \subseteq \mathbb{R}^n$ be a semialgebraic set defined only by inequalities. Assume that the following two conditions are satisfied:*

(1) *$S(\mathbf{g})$ is explicitly bounded: $g_1(\mathbf{x}) = R^2 - \sum_{i=1}^n \mathbf{x}_i^2$ with constant $1 \leq R \leq 2^{\mathrm{poly}(n)}$.*
(2) *$S(\mathbf{g})$ contains a ball of radius $r \geq 2^{-\mathrm{poly}(n)}$, i.e., $B(z, r) \subseteq S(\mathbf{g})$ for some $z \in \mathbb{R}^n$.*

Then, for fixed $t \geq \lceil \deg(f)/2 \rceil$ and $\varepsilon \geq 2^{-\mathrm{poly}(n)}$, the bound $\mathrm{sos}(f)_t$ *may be computed in polynomial time in n up to an additive error of at most ε.*

The inclusions $B(z, r) \subseteq S(\mathbf{g}) \subseteq B(z, R)$ for $\frac{1}{r}, R \leq 2^{\mathrm{poly}(n)}$ are a natural way to ensure that f has an approximate minimizer over $S(\mathbf{g})$ whose bit-complexity is $\mathrm{poly}(n)$. Furthermore, they are very reminiscent of the sufficient conditions for solving semidefinite programs in polynomial time, see Theorem 15 below.

As we will see in Proposition 14, it is possible to choose constraints $\mathbf{g} = (g_i)$, each with *constant* bit-complexity, such that the second condition of Theorem 8 is not satisfied. Notably, the resulting semialgebraic set $S(\mathbf{g})$ does not satisfy the conditions of Theorem 5 or Theorem 6 either.

Finally, as a third contribution, we make explicit the connection between computational aspects of the primal formulation (SOS) of the sos-hierarchy, and its dual formulation (MOM) in terms of moments. This connection is implicitly present in the proof of Theorem 5 in [20].

Theorem 9. *Let $S(\mathbf{g}, \mathbf{h})$ be a semialgebraic set and suppose that the conditions of Theorem 6 or Theorem 8 are satisfied. Then, for a fixed $t \geq \lceil \deg(f)/2 \rceil$ and $\varepsilon > 0$, there exists a sum-of-squares representation (1) proving nonnegativity of $f - \mathrm{sos}(f)_t + \varepsilon$ on $S(\mathbf{g}, \mathbf{h})$ with bit-complexity* $\mathrm{poly}(n, \log(1/\varepsilon))$.

2 ON THE GEOMETRIC CONDITION

Before we move on to the proofs of our results, let us give some examples of natural settings where they may be applied. In general, Theorem 8 is better equipped to deal with non-discrete semialgebraic sets $S(\mathbf{g}, \mathbf{h})$ than Theorem 5. The third condition of Theorem 5, which demands in particular that $g_i(\mathbf{x}) > 0$ for each $\mathbf{x} \in S(\mathbf{g}, \mathbf{h})$, is rather hard to satisfy:

Example 1. The unit hypercube $[-1, 1]^n$, the unit ball $B^n \subseteq \mathbb{R}^n$ and the standard simplex $\Delta^n \subseteq \mathbb{R}^n$ are semialgebraic sets, defined by:

$$\begin{aligned} [-1,1]^n &= \{\mathbf{x} \in \mathbb{R}^n : 1 - \mathbf{x}_i^2 \geq 0, \quad i = 1, 2, \ldots, n\}, \\ B^n &= \{\mathbf{x} \in \mathbb{R}^n : 1 - \|\mathbf{x}\|_2^2 \geq 0\}, \\ \Delta^n &= \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}_i \geq 0,\ 1 - \textstyle\sum_{i=1}^n \mathbf{x}_i \geq 0\}. \end{aligned}$$

It is straightforward to see that they each satisfy the conditions of Theorem 8 (after adding a ball constraint $g_1(\mathbf{x}) = R^2 - \|\mathbf{x}\|_2^2$ if needed). They do not, however, satisfy the third condition of Theorem 5.

The following proposition gives an alternative sufficient condition related to strict feasibility of (POP), which implies that the conditions for Theorem 8 are satisfied and thus guarantees polynomial-time computability of the bound $\mathrm{mom}(f)_t$.

Proposition 10. *Let $S(\mathbf{g})$ be a full-dimensional semialgebraic set contained in a ball of radius $R \leq 2^{\mathrm{poly}(n)}$. Assume that there exists a rational point $\mathbf{x} \in S(\mathbf{g})$ with $g_i(\mathbf{x}) > 0$ for all i, whose bit-complexity is polynomial in n. Then $S(\mathbf{g})$ contains a ball of radius $r \geq 2^{-\mathrm{poly}(n)}$.*

Proof. As the polynomials g_i are of fixed degree and have bounded coefficients, their Lipschitz constants on $B(0, R)$ can be bounded by $2^{\mathrm{poly}(n)}$. Furthermore, as $\mathbf{x}$ has polynomial bit-complexity in n, we know that $g_i(\mathbf{x}) \in \mathbb{Q}$ has polynomial bit-complexity as well, and therefore $g_i(\mathbf{x}) > 0$ implies that $g_i(\mathbf{x}) \geq 2^{-\mathrm{poly}(n)}$. Together, this implies that there exists an $r \geq 2^{-\mathrm{poly}(n)}$ such that $g_i(\mathbf{y}) \geq 0$ for all $\mathbf{y}$ with $\|\mathbf{x} - \mathbf{y}\|_2 \leq r$, i.e., such that $S(\mathbf{g})$ contains the ball of radius $r \geq 2^{-\mathrm{poly}(n)}$ centered at $\mathbf{x}$. □

2.1 Semialgebraic sets with large volume contain a large ball

Another class of semialgebraic sets satisfying the conditions of Theorem 8 are those of sufficiently large volume. In the case that $S(\mathbf{g})$ is convex, we have the following consequence of John's theorem.

Lemma 11. *Let $S(\mathbf{g}) \subseteq \mathbb{R}^n$ be a* convex *semialgebraic set. Assume that $S(\mathbf{g})$ is contained in a ball of radius $R \leq 2^{\mathrm{poly}(n)}$. Then there exists a constant*

$$r \geq \sqrt{\mathrm{vol}(S(\mathbf{g}))} \cdot 2^{-\mathrm{poly}(n)}$$

such that $S(\mathbf{g})$ contains a (translated) ball of radius r.

Proof. By John's theorem [10] there exists an ellipsoid E with center $c \in \mathbb{R}^n$ such that

$$E \subseteq S(\mathbf{g}) \subseteq c + n(E - c).$$

These inclusions show that $n^n \operatorname{vol}(E) = \operatorname{vol}(c+n(E-c)) \geq \operatorname{vol}(S(\mathbf{g}))$ and thus $\operatorname{vol}(E) \geq \operatorname{vol}(S(\mathbf{g}))n^{-n}$. Let $E = \{x \in \mathbb{R}^n : (x-c)^T A^{-1}(x-c) \leq 1\}$ for a positive definite matrix A. Let v be an eigenvector of A corresponding to eigenvalue λ. Then $c \pm \sqrt{\lambda}v \in E \subseteq B(0,R)$ and hence $\lambda \leq R^2$. Moreover, $\operatorname{vol}(E) = \operatorname{vol}(B(0,1)) \prod_{i=1}^n \sqrt{\lambda_i}$ where $\lambda_1, \ldots, \lambda_n$ are the eigenvalues of A. Combined with the lower bound on $\operatorname{vol}(E)$, this shows for each $i \in [n]$ that

$$\lambda_i \geq \operatorname{vol}(S(\mathbf{g}))n^{-n}(R^2)^{-(n-1)} \operatorname{vol}(B(0,1))^{-1}.$$

We thus have $B(z,r) \subseteq E \subseteq S(\mathbf{g})$ for $z = c$ and

$$r = \left(\frac{\operatorname{vol}(S(\mathbf{g}))}{n^n R^{2(n-1)} \operatorname{vol}(B(0,1))}\right)^{\frac{1}{2}} \geq \sqrt{\operatorname{vol}(S(\mathbf{g}))} \cdot 2^{-\operatorname{poly}(n)}.$$

□

Lemma 11 can be used to show that full-dimensional, rational polytopes contain a large enough ball (which is well known, see, e.g., [8, Thm. 3.6], [7, Lem. 1-2]).

Corollary 12. *Let $P \subseteq \mathbb{R}^n$ be a full-dimensional polytope defined by rational linear inequalities with polynomial bit-complexity. Then P is contained in a ball of radius $R \leq 2^{\operatorname{poly}(n)}$ and $\operatorname{vol} P \geq 2^{-\operatorname{poly}(n)}$. Using Lemma 11, P thus contains a ball of radius $r \geq 2^{-\operatorname{poly}(n)}$.*

Proof. Let $\{p_1, \ldots, p_N\}$ be the vertices of P, each of which has polynomial bit-complexity in n, as a consequence of Cramer's rule. As P is full-dimensional, we may assume w.l.o.g. that $0 \in P$ and that $p_1, p_2, \ldots p_n$ are linearly independent. Therefore,

$$0 < \operatorname{vol}\left(\operatorname{span}\{p_1, \ldots p_n\}\right) = \frac{1}{n!}|\det\left(p_1\ p_2\ \ldots\ p_n\right)|.$$

But since the p_i's have polynomial bit-complexity in n, we may conclude that

$$\operatorname{vol} P \geq \frac{1}{n!}|\det\left(p_1\ p_2\ \ldots\ p_n\right)| \geq 2^{-\operatorname{poly}(n)}.$$

□

In fact, a result similar to Lemma 11 can be shown *without* the assumption that $S(\mathbf{g})$ is convex. That is, *any* bounded semialgebraic set $S(\mathbf{g}) \in B(0,R)$ with $R \leq 2^{\operatorname{poly}(n)}$ and $\operatorname{vol}(S(\mathbf{g})) \geq 2^{-\operatorname{poly}(n)}$ satisfies the conditions of Theorem 8. To show this, we use an upper bound on the volume of neighborhoods of algebraic varieties from [1]. Obtaining such volume bounds is a problem with a long history, see, e.g., [1, 14, 24] for a discussion.

Proposition 13. *Let $S(\mathbf{g}) \subseteq \mathbb{R}^n$. Assume that $S(\mathbf{g})$ is contained in a ball of radius $R \leq 2^{\operatorname{poly}(n)}$. Then there exists a constant*

$$r \geq \operatorname{vol}(S(\mathbf{g})) \cdot 2^{-\operatorname{poly}(n)}$$

such that $S(\mathbf{g})$ contains a (translated) ball of radius r.

Proof. Let $G = \prod_{g \in \mathbf{g}} g$. First observe that the boundary of $S(\mathbf{g})$ is contained in the variety $V(G) = \{\mathbf{x} \in \mathbb{R}^n : G(\mathbf{x}) = 0\}$. Now consider the δ-neighborhood of $V(G)$, i.e., the set of $\mathbf{x} \in \mathbb{R}^n$ such that $\operatorname{dist}(\mathbf{x}, V(G)) \leq \delta$. In [1, Theorem 3.2], building on [14], it is shown that for all positive δ we have

$$\frac{\operatorname{vol}(V(G) + B(0,\delta))}{\operatorname{vol}(B(0,R))} \leq 4 \sum_{i=1}^n \binom{n}{i} \left(\frac{4D\delta}{R}\right)^i \left(1 + \frac{\delta}{R}\right)^{n-i}$$

where $D = \prod_{g \in \mathbf{g}} \deg(g)$.[2] In particular, this shows that

$$\operatorname{vol}\left(V(G) + B(0,\delta)\right) \leq \sum_{j=1}^n c_j \delta^j$$

where the coefficients $c_j \in \mathbb{R}$ satisfy $0 < c_j \leq 2^{\operatorname{poly}(n)}$. For $\delta \leq 1$, the upper bound is of the form $\delta \cdot 2^{p(n)}$ for some polynomial p, and hence for $\delta < \operatorname{vol}(S(\mathbf{g}))/2^{p(n)}$ we have

$$\operatorname{vol}(\partial S(\mathbf{g}) + B(0,\delta)) \leq \operatorname{vol}(V(G) + B(0,\delta)) < \operatorname{vol}(S(\mathbf{g})),$$

which implies that there exists an $\mathbf{x} \in S(\mathbf{g})$ that has distance greater than δ to the boundary of $S(\mathbf{g})$. The ball $B(\mathbf{x}, r)$ with center $\mathbf{x}$ and radius $r = \delta \geq \operatorname{vol}(S(\mathbf{g}))/2^{\operatorname{poly}(n)}$ is thus contained in $S(\mathbf{g})$. □

2.2 A counterexample

To end this section, we show that the second condition of Theorem 8 is not superfluous by exhibiting a full-dimensional semialgebraic set which does not contain a hypercube $[-r,r]^n$ of size $r \geq 2^{-\operatorname{poly}(n)}$. As a direct result, it cannot contain a ball of radius $r \geq 2^{-\operatorname{poly}(n)}$, either. We use a simple repeated-squaring argument.

Proposition 14. *There exists a full-dimensional semialgebraic set $S(\mathbf{g})$, defined only by polynomial inequalities $\mathbf{g} = (g_1, g_2, \ldots, g_m)$ of fixed degree, whose coefficients have constant bit-complexity, which does not contain a (translated) cube $[-r,r]^n$ for $r \geq 2^{-\operatorname{poly}(n)}$.*

Proof. Let $S(\mathbf{g})$ be the set defined by the system of inequalities:

$$\begin{aligned} \mathbf{x}_i &\geq 0 && (1 \leq i \leq n), \\ \mathbf{x}_i - \mathbf{x}_{i+1}^2 &\leq 0 && (1 \leq i \leq n-1), \\ \mathbf{x}_n &\leq 1/2. \end{aligned}$$

Set $r := 2^{-2^{n-1}}$. Then $[0,r]^n \subseteq S(\mathbf{g})$ and so $S(\mathbf{g})$ is full-dimensional. But from the inequalities it follows that $0 \leq \mathbf{x}_1 \leq r$ for any $\mathbf{x} \in S(\mathbf{g})$, meaning $S(\mathbf{g})$ cannot contain a cube of size at least $2^{-\operatorname{poly}(n)}$. □

3 PRELIMINARIES

3.1 Notation

We denote by $\mathbb{R}[\mathbf{x}]$ the space of n-variate polynomials. For $d \in \mathbb{N}$, we write $\mathbb{R}[\mathbf{x}]_d$ for the subspace of polynomials of degree at most d, whose dimension is equal to $h(n,d) := \binom{n+d}{d}$. It has a basis of monomials $\mathbf{x}^\alpha = \mathbf{x}_1^{\alpha_1} \ldots \mathbf{x}_n^{\alpha_n}$, with $\alpha \in \mathbb{N}_d^n := \{\alpha \in \mathbb{N}^n : \|\alpha\|_1 \leq d\}$. We denote by $\mathcal{S}^n$ the space of n-by-n real symmetric matrices. We have the trace inner product $\langle X, Y \rangle = \operatorname{Tr}(XY)$ on $\mathcal{S}^n$, which induces the Frobenius norm $\|X\|_F = \sqrt{\langle X, X \rangle}$.

3.2 Complexity of semidefinite programming

We will make use of the following result on the complexity of semidefinite programming. Let $C \in \mathbb{Q}^{n \times n}$, $A_i \in \mathbb{Q}^{n \times n}$, $b_i \in \mathbb{Q}$, $1 \leq i \leq m$ be rational input data. Consider the following SDP in standard form:

$$\begin{aligned} \operatorname{val} = \inf\ & \langle C, X \rangle \\ & \langle A_i, X \rangle = b_i \quad (1 \leq i \leq m), \\ & X \in \mathcal{S}^n \text{ is positive semidefinite.} \end{aligned} \tag{SDP}$$

[2] Theorem 3.2 of [1] applies to more general algebraic varieties V, here we use the special case where V is the zero set of a single polynomial and $\dim_{\mathbb{R}}(V) = n - 1$.

We denote the feasible region of (SDP) by:

$$\mathcal{F} := \{X \in \mathcal{S}^n : X \succeq 0, \langle A_i, X\rangle = b_i \quad (1 \le i \le m)\}.$$

One can show, for example using the ellipsoid method [9], that under certain assumptions (SDP) can be solved in polynomial time. We use the following explicit formulation from [6] where a similar result was shown using an interior point method.

Theorem 15 (Thm. 1.1 in [6]). *Let $r, R > 0$ be given and suppose that there exists an $X_0 \in \mathcal{F}$ so that:*

$$B(X_0, r) \subseteq \mathcal{F} \subseteq B(X_0, R),$$

where $B(X_0, \lambda)$ is the ball of radius $\lambda \in \mathbb{R}$ (in the norm $\|\cdot\|_F$) centered at X_0 in the subspace:

$$V = V(\mathcal{F}) := \{X \in \mathcal{S}^n : \langle A_i, X\rangle = b_i \quad (1 \le i \le m)\}.$$

Then for any rational $\varepsilon > 0$ one can find a rational matrix $X^* \in \mathcal{F}$ such that:

$$\langle C, X^*\rangle - \mathrm{val} \le \varepsilon,$$

in time polynomial in $n, m, \log(R/r), \log(1/\varepsilon)$ and the bit-complexity of the input data C, A_i, b_i and the feasible point X_0.

3.3 Dual formulation and moments of measures

It will be convenient to work with the dual formulation of (SOS), which we recall is (see, e.g., [5]):

$$\begin{aligned} \mathrm{mom}(f)_t := \inf\ & L(f) \\ \text{s.t. } & L(1) = 1, \\ & L(g_i p^2) \ge 0, \quad (g_i p^2 \in \mathbb{R}[\mathbf{x}]_{2t}) \\ & L(h_j \mathbf{x}^\alpha) = 0, \quad (h_j \mathbf{x}^\alpha \in \mathbb{R}[\mathbf{x}]_{2t}) \\ & L \in \mathbb{R}[\mathbf{x}]_{2t}^*. \end{aligned} \tag{MOM}$$

Assuming that (POP) is explicitly bounded, these formulations are equivalent.

Theorem 16 ([3, Sec. 2.6.3, Ex. 2.12], see also [11]). *If (POP) is explicitly bounded[3], we have strong duality between the primal and dual formulations (SOS) and (MOM) of the moment-SOS hierarchy. That is, we then have:*

$$\mathrm{sos}(f)_t = \mathrm{mom}(f)_t \quad \forall t \ge \lceil \deg(f)/2 \rceil.$$

As is well known (see e.g., [13, Prop. 6.2]), explicit boundedness of $S(\mathbf{g}, \mathbf{h})$ also gives a bound on the feasible region of (MOM).

Lemma 17. *Assume that $S(\mathbf{g}, \mathbf{h})$ is explicitly bounded for some $R > 0$. Let $L \in \mathbb{R}[x]_{2t}^*$ be a feasible solution to (MOM). Then $|L(\mathbf{x}^\alpha)| \le R^{|\alpha|}$ for all $\alpha \in \mathbb{N}_{2t}^n$.*

There is a natural relation between the dual formulation (MOM) and *moments* of measures supported on $S(\mathbf{g}, \mathbf{h})$, which clarifies the assumptions made in Theorem 5 and Theorem 6. For a measure μ supported on $S(\mathbf{g}, \mathbf{h})$, the *moment* of degree $\alpha \in \mathbb{N}^n$ is defined by:

$$m(\mu)_\alpha := \int_{S(\mathbf{g},\mathbf{h})} \mathbf{x}^\alpha d\mu(\mathbf{x}).$$

[3] If we assume instead that $S(\mathbf{g}, \mathbf{h})$ is Archimedean with non-empty interior, then it is known that $\mathrm{sos}(f)_t = \mathrm{mom}(f)_t$ for t large enough [12, Theorem 4.2].

For $t \in \mathbb{N}$, the (truncated) moment matrix $M_t(\mu)$ of order t for μ is the matrix of size $h(n,t) = \binom{n+t}{t}$ given by:

$$M_t(\mu)_{\alpha,\beta} = m(\mu)_{\alpha+\beta} \quad (|\alpha|, |\beta| \le t). \tag{5}$$

Consider the linear functional $L_\mu \in \mathbb{R}[\mathbf{x}]_{2t}^*$ defined by:

$$L_\mu(p) := \int_{S(\mathbf{g},\mathbf{h})} p(\mathbf{x}) d\mu(x) \quad (p \in \mathbb{R}[\mathbf{x}]_{2t}). \tag{6}$$

For any constraint g_i and $p \in \mathbb{R}[\mathbf{x}]$ with $\deg(g_i p^2) \le 2t$, we have:

$$L_\mu(g_i p^2) = \mathbf{p}^\top M_t(g_i \mu)\mathbf{p} = \int_{S(\mathbf{g},\mathbf{h})} p^2(\mathbf{x}) g_i(\mathbf{x}) d\mu(\mathbf{x}) \ge 0,$$

where $\mathbf{p}$ denotes the vector of coefficients of $p \in \mathbb{R}[\mathbf{x}]_t$ in the monomial basis. Here the (α, β)-entry of the *localizing matrix* $M_t(g_i\mu)$ is defined as $\int_{S(\mathbf{g},\mathbf{h})} g_i(\mathbf{x}) \mathbf{x}^{\alpha+\beta} d\mu(\mathbf{x})$. In particular, for each i the matrix $M_t(g_i\mu)$ is positive semidefinite. Note that $M_t(\mu) = M_t(L_\mu)$ and similarly for the localizing matrices. Furthermore, for any constraint h_j and $\alpha \in \mathbb{N}^n$ with $\deg(\mathbf{x}^\alpha h_j) \le 2t$, we have:

$$L_\mu(h_j \mathbf{x}^\alpha) = \int_{S(\mathbf{g},\mathbf{h})} h_j(\mathbf{x}) \mathbf{x}^\alpha d\mu(\mathbf{x}) = 0.$$

If μ is a probability measure, we get $L_\mu(1) = 1$, and it follows that L_μ is a feasible solution to (MOM).

3.4 Standard form of the moment formulation: Proof of Proposition 6

In order to apply Theorem 15 in our proof of Proposition 6 below, we need to express the formulation (MOM) in the standard form (SDP). Set $N = h(n,t) + \sum_{i\in[m]} h(n, t - \lceil \deg(g_i)/2 \rceil)$. One can construct $A_1, \ldots, A_K \in \mathbb{Q}^{N\times N}$ and $b_1, \ldots, b_K \in \mathbb{Q}$ with entries either -1, 0, 1, or a coefficient of the polynomials g_i, h_j, such that the conditions $\langle A_j, X\rangle = b_j$ for $j \in [K]$ are equivalent to the statement that X is a block-diagonal matrix $X = M_t(L) \oplus \left(\bigoplus_{i\in[m]} M_t(g_i L)\right)$ where

$$M_t(L)_{\alpha,\beta} = L(\mathbf{x}^{\alpha+\beta}) \qquad (|\alpha|, |\beta| \le t), \tag{7}$$

$$M_t(g_i L)_{\alpha,\beta} = L(g_i(\mathbf{x}) \mathbf{x}^{\alpha+\beta}) \quad (|\alpha|, |\beta| \le t - \lceil \deg(g_i)/2 \rceil), \tag{8}$$

for some linear functional $L \in \mathbb{R}[\mathbf{x}]_{2t}^*$ satisfying $L(1) = 1$ and $L(h_j \mathbf{x}^\alpha) = 0$ for each $j \in [\ell]$ and $|\alpha| \le 2t - \deg(h_j)$. Let $C_f \in \mathcal{S}^N$ be such that $\langle C_f, X\rangle = L(f)$, then we have

$$\begin{aligned} \mathrm{mom}(f)_t = \inf\ & \langle C_f, X\rangle \\ & \langle A_i, X\rangle = b_i \quad (1 \le i \le K), \\ & X \in \mathcal{S}^N \text{ is positive semidefinite.} \end{aligned} \tag{MOM-SDP}$$

To see the equivalence it suffices to observe that the conditions $L(g_i p^2) \ge 0$ for all g_i and p of appropriate degree of (MOM) are equivalent to $X \succeq 0$.

The upshot is that we may apply Theorem 15 to the formulation (MOM) in the following way.

Proposition 18. *Consider an instance of (MOM) with feasible solution $L_0 \in \mathbb{R}[\mathbf{x}]_{2t}^*$ that satisfies the property: $L_0(g_i p^2) = 0$ implies $g_i p^2 = \sum_{j=1}^{\ell} p_j h_j$ for some polynomials p_j with $\deg(p_j) \le 2t - \deg(h_j)$. Let $X_0 = M_t(L_0) \oplus \left(\bigoplus_{i\in[m]} M_t(g_i L_0)\right)$ be the matrix associated to L_0 via (7) and (Eq. (8)). If all non-zero eigenvalues of X_0 are at least $r > 0$, then $B(X_0, r/2)$ is contained in the feasible region of (MOM-SDP).*

Proof. Let $X \in B(X_0, r/2)$ and write $X = X_0 + \tilde{X}$. Here $B(X_0, r/2)$ is the ball of radius $r/2$ (in the Frobenius norm) in the affine space $V(\mathcal{F})$ defined by the linear equalities of (MOM-SDP). We first show that a zero-eigenvector of X_0 is a zero-eigenvector of X. Since X_0 is block-diagonal, a zero-eigenvector of X_0 corresponds to a zero-eigenvector v of a single block $M_t(g_i L_0)$. For such a vector v we have

$$0 = v^T M_t(g_i L_0) v = L_0(g_i(\mathbf{x}) p_v(\mathbf{x})^2)$$

where $p_v(\mathbf{x}) = \sum_{|\alpha| \le t - \lceil \deg(g_i)/2 \rceil} v_\alpha \mathbf{x}^\alpha$. By assumption on L_0, we therefore have $g_i p_v^2 = \sum_{j=1}^{\ell} p_j h_j$ for polynomials p_j with $\deg(p_j) \le 2t - \deg(h_j)$.

To show that v corresponds to a zero-eigenvector of X it remains to observe that, since $X \in V(\mathcal{F})$, it corresponds to a linear functional $L \in \mathbb{R}[\mathbf{x}]_{2t}^*$ via (7) and (8) with the property that $L(h_j p) = 0$ for all $j \in [m]$ and polynomials p of degree at most $2t - \deg(h_j)$. In particular,

$$L(g_i(\mathbf{x}) p_v(\mathbf{x})^2) = \sum_{j \in [\ell]} L(h_j p_j) = 0,$$

and therefore v corresponds to a zero-eigenvector of the g_i-th block of X.

Finally, to show that X is positive semidefinite, note that by assumption $X = X_0 + \tilde{X}$ where $\|\tilde{X}\|_F \le r/2$. In particular, the eigenvalues of $\tilde{X}$ are all at most $r/2$. Since the kernel of X_0 is contained in the kernel of X, this means that all non-zero eigenvalues of X are at least $r/2$ and hence X is positive semidefinite. □

As a corollary of Proposition 18 we obtain Proposition 6.

Proof of Proposition 6. By assumption L is such that for all $g \in \mathbf{g}$ and $p \in \mathbb{R}[\mathbf{x}]_{t - \lceil \deg(g)/2 \rceil}$ we have $L(gp^2) = 0$ implies $gp^2 = \sum_{j=1}^{\ell} p_j h_j$ for some polynomials p_j with $\deg(p_j) \le 2t - \deg(h_j)$. Moreover, all non-zero eigenvalues of $M_t(L) \oplus \left(\bigoplus_{i \in [m]} M_t(g_i L) \right)$ are at least $2^{-\operatorname{poly}(n)}$. By Proposition 18, the feasible region of (MOM-SDP) contains a ball of radius $2^{-\operatorname{poly}(n)}$ (in $V(\mathcal{F})$), and by Lemma 17 it is contained in a ball of radius $2^{\operatorname{poly}(n)}$. Theorem 15 thus shows that we can compute an ε-additive approximation to the value $\operatorname{mom}(f)_t$ in time polynomial in n and $\log(1/\varepsilon)$. □

4 AN ALGEBRAIC CONDITION FOR POLYNOMIAL-TIME COMPUTABILITY: PROOF OF THEOREM 7

We split the proof of Theorem 7 into two parts: we first consider the moment matrix $M_t(L)$ and then the localizing matrices $M_t(gL)$ for $g \in \mathbf{g}$. A key tool is the following auxiliary lemma.

Lemma 19 (Lemma 3.1 in [20].). *Let $M \in \mathbb{Z}^{N \times N}$ be a symmetric integer matrix with $|M_{ij}| \le B$ for all $i, j \in [N]$. Then each non-zero eigenvalue of M has absolute value at least $(BN)^{-N}$.*

Using this lemma, we are able to show the following.

Proposition 20. *Assume that $S(\mathbf{g}, \mathbf{h}) \subseteq B(0, R)$ is explicitely bounded with $R \le 2^{\operatorname{poly}(n)}$, and let $L \in \mathbb{R}[\mathbf{x}]_{2t}^*$ be a feasible solution to (MOM). Assume further that $L(\mathbf{x}^\alpha) \in \mathbb{Q}$ has polynomial bit-complexity for all $\alpha \in \mathbb{N}_{2t}^n$. Then $\lambda_{\min}(M_t(L))$, the smallest non-zero eigenvalue of the moment matrix $M_t(L)$ of (7), satisfies:*

$$\lambda_{\min}(M_t(L)) \ge 2^{-\operatorname{poly}(n)}.$$

Proof. By assumption, there exists an integer $C \le 2^{\operatorname{poly}(n)}$ such that $C \cdot M_t(L)$ is an integer matrix. Furthermore, using Lemma 17, the largest entry of $C \cdot M_t(L)$ in absolute value is bounded from above by $C \cdot R^{2t} \le 2^{\operatorname{poly}(n)}$. Writing $N := h(n, t) = \binom{n+t}{t} \le \operatorname{poly}(n)$ for the number of monomials of degree at most t, we may thus use Lemma 19 to conclude that the smallest non-zero eigenvalue of $M_t(L)$ is at least $\lambda_{\min}(M_t(L)) \ge (2^{\operatorname{poly}(n)} \cdot N)^{-N} \ge 2^{-\operatorname{poly}(n)}$. □

Proposition 21. *Under the same assumptions as in Proposition 20, let $g \in \mathbb{R}[\mathbf{x}]$ be one of the constraints defining $S(\mathbf{g}, \mathbf{h})$. Then the smallest non-zero eigenvalue of the localizing matrix $M_t(gL)$ of (8) satisfies $\lambda_{\min}(M_t(gL)) \ge 2^{-\operatorname{poly}(n)}$.*

Proof. We may express g in the monomial basis as:

$$g(\mathbf{x}) = \sum_{|\alpha| \le d} g_\alpha \mathbf{x}^\alpha \quad (g_\alpha \in \mathbb{Q}).$$

The coefficients g_α have polynomial bit-complexity by assumption, and so $|g_\alpha| \le 2^{\operatorname{poly}(n)}$ for each $|\alpha| \le d$. For the same reason, there exists an integer $C_g \le 2^{\operatorname{poly}(n)}$ such that $C_g g_\alpha \in \mathbb{Z}$ for all $|\alpha| \le d$. Recall that the entries of $M_t(gL)$ are linear combinations of the entries of $M_t(L)$, namely for $|\alpha|, |\beta| \le t - \lceil \deg(g_i)/2 \rceil$ the (α, β)-th entry is of the form:

$$M_t(gL)_{\alpha,\beta} = \sum_{|\gamma| \le \deg(g_i)} g_\gamma L(\mathbf{x}^{\alpha+\beta+\gamma}).$$

Now, as in Proposition 20, we see that $C_g \cdot C \cdot M_t(gL)$ is an integer matrix for some integer $C \le 2^{\operatorname{poly}(n)}$. Furthermore, since the entries of $M_t(L)$ are at most $2^{\operatorname{poly}(n)}$, the entries of $M_t(gL)$ are bounded from above by

$$2^{\operatorname{poly}(n)} \cdot \sum_{|\alpha| \le d} |g_\alpha| \le 2^{\operatorname{poly}(n)} \cdot \binom{n+d}{d} \cdot C_g \le 2^{\operatorname{poly}(n)}.$$

As before, we may thus invoke Lemma 19 to conclude the proof. □

5 A GEOMETRIC CONDITION FOR POLYNOMIAL-TIME COMPUTABILITY: PROOF OF THEOREM 8

Recall that we consider in Theorem 8 an explicitly bounded semialgebraic set $S(\mathbf{g}) \subseteq \mathbb{R}^n$ with the additional geometric assumption:

$$B(z, r) \subseteq S(\mathbf{g}) \tag{9}$$

for some $r \ge 2^{-\operatorname{poly}(n)}$ and $z \in \mathbb{R}^n$. To prove Theorem 8, we will exploit this assumption to exhibit a feasible solution $L \in \mathbb{R}[\mathbf{x}]_{2t}^*$ to (MOM) that satisfies the conditions of Theorem 7.

We begin by noting that the inclusion (9) implies that $S(\mathbf{g})$ contains a translated hypercube $B := [-r, r]^n + z$, for a (slightly smaller) $r \ge 2^{-\operatorname{poly}(n)}$. We consider the probability measure μ_z obtained by restricting the Lebesgue measure to B and renormalizing. We show that the operator $L_{\mu_z} \in \mathbb{R}[\mathbf{x}]_{2t}^*$ associated to μ_z via (4) satisfies the conditions of Proposition 6. Let us first note that condition (1) is satisfied automatically as B is full-dimensional. Indeed, this means that $\int_{S(\mathbf{g})} gp^2 d\mu = 0$ if and only if $gp^2 = 0$. It remains to show that condition (2) also holds, for which we use Theorem 7.

For simplicity, we assume first that $z = 0$, so that $B = [-r, r]^n$. In this case, we may use an explicit formula for the moments of μ_0

(see, e.g., [5]) to find:

$$L_{\mu_0}(\mathbf{x}^\alpha) = \int_{[-r,r]^n} \mathbf{x}^\alpha d\mu_0(\mathbf{x}) = \begin{cases} \prod_{i=1}^n \frac{r^{\alpha_i}}{\alpha_i+1} & \text{if } \alpha \in (2\mathbb{N})^n, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

As $r \geq 2^{-\operatorname{poly}(n)}$, it follows that $L_{\mu_0}(\mathbf{x}^\alpha) \in \mathbb{Q}$ has polynomial bit-complexity for all $|\alpha| \leq 2t$, i.e., that L_{μ_0} satisfies the condition of Theorem 7.

It remains to consider the case $z \neq 0$. Since $S(\mathbf{g})$ is explicitly bounded, we must have $\|z\|_2 \leq R \leq 2^{\operatorname{poly}(n)}$. After possibly choosing a slightly smaller r, we may assume that $z \in \mathbb{Q}^n$ and that z has polynomial bit-complexity in n. First, note for all $|\alpha| \leq 2t$ that

$$\begin{aligned} L_{\mu_z}(\mathbf{x}^\alpha) &= \int_{[-r,r]^n+z} \mathbf{x}^\alpha d\mu_z(\mathbf{x}) = \int_{[-r,r]^n} (\mathbf{x}-z)^\alpha d\mu_0(\mathbf{x}) \\ &= L_{\mu_0}((\mathbf{x}-z)^\alpha). \end{aligned}$$

Second, note that the coefficients c_γ in the expansion of $(\mathbf{x}-z)^\alpha$ in the monomial basis all have polynomial bit-complexity. The nonzero coefficients are those for which γ is entrywise less than or equal to α, which we denote $\gamma \leq \alpha$. Note that there are at most $2^{\operatorname{poly}(n)}$ such coefficients. We may thus use (10) to find that

$$L_{\mu_z}(\mathbf{x}^\alpha) = L_{\mu_0}((\mathbf{x}-z)^\alpha) = \sum_{\gamma \leq \alpha} c_\gamma \cdot L_{\mu_0}(\mathbf{x}^\gamma)$$

has polynomial bit-complexity for all $\alpha \in \mathbb{N}^n$ with $|\alpha| \leq 2t$. It follows that L_{μ_z} satisfies the condition of Theorem 7.

6 FROM MOMENTS TO SUMS OF SQUARES: PROOF OF THEOREM 9

Proposition 6 and Theorem 8 show that, under their respective conditions, we can find an additive ε-approximation to $\operatorname{mom}(f)_t = \operatorname{sos}(f)_t$ in time $\operatorname{poly}(n, \log(1/\varepsilon))$ by solving (MOM-SDP). We now show that, under these conditions, we have a compact sum-of-squares proof for this bound as well.

Theorem 22 (Detailed version of Theorem 9). *Let $S(\mathbf{g},\mathbf{h})$ be a semialgebraic set and suppose that the conditions of Theorem 6 or Theorem 8 are satisfied for some fixed $t \geq \lceil \deg(f)/2 \rceil$. Suppose that $f - \lambda$ has a sum-of-squares decomposition:*

$$f(\mathbf{x}) - \lambda = \sigma_0(\mathbf{x}) + \sum_{i=1}^m g_i(\mathbf{x})\sigma_i(\mathbf{x}) + \sum_{j=1}^\ell h_j(\mathbf{x})p_j(\mathbf{x}), \quad (11)$$

where $\deg(h_j p_j) \leq 2t$ and the sum-of-squares are of the form $\sigma_i = \sum_{k=1}^K s_{i,k}^2$ for $s_{i,k} \in \mathbb{R}[\mathbf{x}]$ with $\deg(g_i s_{i,k}^2) \leq 2t$. Then, for a fixed $\varepsilon > 0$, there exists a polynomial $\mathcal{E}(\mathbf{x})$ of degree $2t$ such that

$$f(\mathbf{x}) - \lambda + \mathcal{E}(\mathbf{x}) = \tilde{\sigma}_0(\mathbf{x}) + \sum_{i=1}^m g_i(\mathbf{x})\tilde{\sigma}_i(\mathbf{x}) + \sum_{j=1}^\ell h_j(\mathbf{x})\tilde{p}_j(\mathbf{x}) \quad (12)$$

where, for $i = 0, \ldots, m$, we have $\tilde{\sigma}_i(\mathbf{x}) = \sum_{k=1}^K \tilde{s}_{i,k}^2$ for polynomials $\tilde{s}_{i,k}$ with bit-complexity $\operatorname{poly}(n, \log(1/\varepsilon))$ and $\|\mathcal{E}\|_1 \leq \varepsilon R^{2t} 2^{\operatorname{poly}(n)}$. In particular, this proves nonnegativity of $f(\mathbf{x}) - \lambda + R^{2t}\|\mathcal{E}\|_1$ on $S(\mathbf{g},\mathbf{h})$.

From the SDP-formulation of (SOS) it follows that one can bound K by the rank of the matrices involved, i.e., we have $K \leq \binom{n+t}{t}$ and thus $K \leq \operatorname{poly}(n)$.

Proof of Theorem 22. Let $\sigma_i = \sum_{k=1}^K s_{i,k}^2$ and p_j as in (11) be given. We first show that the coefficients of the $s_{i,k}$ are upper bounded. To do so, let L be a linear functional that satisfies the conditions of Proposition 6 (in the proof of Theorem 8 we also construct such an L). We then have $L(f-\lambda) = L(\sigma_0) + \sum_{i=1}^m L(g_i\sigma_i)$, and since all terms on the right-hand side are nonnegative, this implies in particular that (using Lemma 17):

$$L(g_i s_{i,k}^2) \leq L(f-\lambda) \leq R^{2t} \cdot \operatorname{poly}(\operatorname{bitcomplexity}(f)),$$

for all $0 \leq i \leq m$ and $1 \leq k \leq K$, where we set $g_0 = 1$ for convenience. We can now distinguish two cases: (i) $L(g_i s_{i,k}^2) \neq 0$, or (ii) $L(g_i s_{i,k}^2) = 0$. In the first case, if $L(g_i s_{i,k}^2) \neq 0$, then we also have

$$L(g_i s_{i,k}^2) \geq \lambda_{\min}(M_t(g_i L)) \cdot \|s_{i,k}\|_2^2,$$

where $\lambda_{\min}(M_t(g_i L))$ is the smallest non-zero eigenvalue of $M_t(g_i L)$, and so

$$\|s_{i,k}\|_2^2 \leq R^{2t} \cdot \lambda_{\min}(M_t(g_i L))^{-1} \cdot \operatorname{poly}(\operatorname{bitcomplexity}(f)).$$

For the second case, if on the other hand $L(g_i s_{i,k}^2) = 0$, then by condition (1) of Proposition 6 we have $g_i s_{i,k}^2 = \sum_{j=1}^\ell h_j q_j$ for some polynomials q_j. We may thus remove such $s_{i,k}$ from the sum-of-squares part of (11) and add them to the ideal part of the certificate.

Using the above bound on the coefficients of the $s_{i,k}$, we now show a bound on the size of the coefficients of the p_j. The polynomial identity (11) allows us to view the coefficients of the p_j as the solution of a linear system $A\mathbf{p} = \mathbf{b}$ where $\mathbf{p}$ is a vector that contains the coefficients of the p_j ($\mathbf{p}$ is at most $\ell\binom{n+2t}{2t}$-dimensional), A contains coefficients of the h_j, and $\mathbf{b}$ is the $\binom{n+2t}{2t}$-dimensional vector that contains the coefficients of $f(\mathbf{x}) - \lambda - \sum_{i=0}^m g_i(\mathbf{x})\sigma_i(\mathbf{x})$. The system $A\mathbf{p} = \mathbf{b}$ is feasible, $\mathbf{b} \neq 0$, and therefore $r = \operatorname{rank}(A)$ is strictly positive. Let $\overline{A}$ be an invertible r-by-r submatrix of A and write $\overline{\mathbf{p}}$ and $\overline{\mathbf{b}}$ for the restrictions of $\mathbf{p}$ and $\mathbf{b}$ to the corresponding rows/columns. Cramer's rule then shows that the ith coordinate of $\overline{\mathbf{p}}$ can be written as

$$\overline{\mathbf{p}}_i = \frac{\det(\overline{A}_i)}{\det(\overline{A})}$$

where $\overline{A}_i$ is the matrix formed by replacing the ith column of $\overline{A}$ with the vector b. To upper bound $|\overline{\mathbf{p}}_i|$ we must give a lower bound on $|\det(\overline{A})|$ and an upper bound on $|\det(\overline{A}_i)|$. Let us first observe that $|\det(\overline{A})| \geq 2^{-\operatorname{poly}(n)}$. Indeed, $\overline{A}$ is an invertible r-by-r matrix with $r \in \operatorname{poly}(n)$, its entries have bit-complexity $\operatorname{poly}(n)$ since they correspond to coefficients of the h_j, and therefore applying Lemma 19 to a suitable integer multiple of $\overline{A}$ shows that $|\det(\overline{A})| \geq 2^{-\operatorname{poly}(n)}$. To upper bound $|\det(\overline{A}_i)|$ it suffices to observe that all entries of $\overline{A}_i$ are upper bounded in absolute value by $2^{\operatorname{poly}(n)}$: for the ith column this follows from the above-derived bound on the σ_i, for the other columns, as before, we observe that they contain coefficients of the h_j. By setting all remaining coordinates to zero, we can extend $\overline{\mathbf{p}}$ to a feasible solution $\mathbf{p}$ of $A\mathbf{p} = \mathbf{b}$. To summarize, this shows that there exists a sum-of-squares decomposition of $f - \lambda$ as in (11) where $\|s_{i,k}\|_2^2 \leq R^{2t}\lambda_{\min}(M_t(g_i L))^{-1} \cdot \operatorname{poly}(n)$ for all i, k and $\|p_j\|_\infty \leq 2^{\operatorname{poly}(n)}$ for all j.

We finally show that rounding each coefficient of this certificate to few bits introduces a small error. For each $i \in [m]$, $j \in [K]$, let $\tilde{s}_{i,k}$ be the polynomial $s_{i,k}$ with each of the coefficients rounded

to the nearest integer multiple of ε. Therefore, for $|\alpha| \leq t$ we have $|(s_{i,k})_\alpha - (\tilde{s}_{i,k})_\alpha| \leq \varepsilon$. Using the identity $\tilde{s}_{i,k}^2 - s_{i,k}^2 = (\tilde{s}_{i,k} + s_{i,k})(\tilde{s}_{i,k} - s_{i,k})$, this shows $|(\tilde{s}_{i,k}^2 - s_{i,k}^2)_\alpha| \leq \varepsilon(2(s_{i,k})_\alpha + \varepsilon)$ and thus $\|\tilde{s}_{i,k}^2 - s_{i,k}^2\|_1 \leq \varepsilon(2\|s_{i,k}\|_1 + \varepsilon\binom{n+2t}{2t})$. Similarly, for each $j \in [\ell]$, let $\tilde{p}_j$ be the polynomial p_j with each of the coefficients rounded to the nearest integer multiple of ε. Then we have (12) for the polynomial $\mathcal{E}(\mathbf{x})$ defined as

$$\mathcal{E}(\mathbf{x}) := \sum_{i=0}^{m} g_i(\tilde{\sigma}_i - \sigma_i) + \sum_{j=1}^{\ell} h_j(\tilde{p}_j - p_j)$$

and hence

$$\|\mathcal{E}\|_1 \leq \sum_{i=0} \|g_i\|_1 \left(\sum_{k=1}^{K} \|\tilde{s}_{i,k}^2 - s_{i,k}^2\|_1\right) + \sum_{j=0}^{\ell} \|h_j\|_1 \|\tilde{p}_j - p_j\|_\infty.$$

As shown above, we have $K \in \mathrm{poly}(n)$, $\|\tilde{p}_j - p_j\|_\infty \leq \varepsilon$, and

$$\|\tilde{s}_{i,k}^2 - s_{i,k}^2\|_1 \leq \varepsilon\left(R^{2t}\lambda_{\min}(M_t(g_iL))^{-1}\,\mathrm{poly}(n) + \varepsilon\binom{n+2t}{2t}\right).$$

Using Proposition 21 we moreover have $\lambda_{\min}(M_t(g_iL))^{-1} \leq 2^{\mathrm{poly}(n)}$. Combining these estimates shows $\|\mathcal{E}\|_1 \leq \varepsilon R^{2t} 2^{\mathrm{poly}(n)}$.

The statement $f(\mathbf{x}) - \lambda + R^{2t}\|\mathcal{E}\|_1 \geq 0$ on $S(\mathbf{g}, \mathbf{h})$ follows from (12) by adding $R^{2t}\|\mathcal{E}\|_1 - \mathcal{E}(\mathbf{x})$ to both sides of the equation, using the fact that $R^{2t}\|\mathcal{E}\|_1 - \mathcal{E}(\mathbf{x}) \in \mathcal{M}(\mathbf{g})_{2t}$, which follows from Lemma 23 below. □

Lemma 23. *Let $g_1, \ldots, g_m \in \mathbb{R}[\mathbf{x}]$ be such that for $R > 0$ we have $R^2 - x_i^2 \in \mathcal{M}(\mathbf{g})_2$ for each $i \in [n]$. For $t \in \mathbb{N}$ and $\alpha \in \mathbb{N}^n$ with $|\alpha| \leq 2t$, we have $R^{|\alpha|} - x^\alpha \in \mathcal{M}(\mathbf{g})_{2t}$.*

Proof. We first consider the case of coordinatewise even exponents, i.e., $R^{2|\alpha|} - x^{2\alpha}$. For this we use induction on $|\alpha|$. The case $|\alpha| = 1$ holds by assumption and for $|\alpha| > 1$ we use the identity

$$R^{2|\alpha|} - x^{2\alpha} = R^{2|\alpha|-2}(R^2 - x_j^2) + x_j^2(R^{2|\alpha|-2} - x^{2(\alpha - e_j)}) \in \mathcal{M}(\mathbf{g})_{2|\alpha|}$$

where $j \in [n]$ is an index for which $\alpha_j > 0$ and e_j is the j-th unit vector. Here we use that the first term on the right hand side belongs to $\mathcal{M}(\mathbf{g})_2$ and the second term belongs to $\mathcal{M}(\mathbf{g})_{2|\alpha|}$ as $R^{2|\alpha|-2} - x^{2(\alpha-e_j)} \in \mathcal{M}(\mathbf{g})_{2|\alpha|-2}$ by the induction hypothesis.

Now let $\gamma \in \mathbb{N}^n$ with $|\gamma| \leq 2t$. We distinguish two cases: $|\gamma|$ is odd or even. When $|\gamma|$ is odd we write $\gamma = \alpha + \beta$ for $\alpha, \beta \in \mathbb{N}^n$ with $|\alpha|, |\beta| \leq t$ and $|\alpha| + 1 = |\beta| = \frac{|\gamma|+1}{2}$. We then observe that we have the following identity

$$\frac{(Rx^\alpha - x^\beta)^2 + R^2(R^{2|\alpha|} - x^{2\alpha}) + (R^{2|\beta|} - x^{2\beta})}{2R} = \frac{R^{2+2|\alpha|} + R^{2|\beta|} - 2Rx^{\alpha+\beta}}{2R} = R^{|\gamma|} - x^\gamma$$

where in the last equality we use the identities $R^{|\gamma|} = R^{|\alpha|+|\beta|} = R^{1+2|\alpha|} = R^{2|\beta|-1}$. In the first part of the proof we have shown that $R^{2|\alpha|} - x^{2\alpha}, R^{2|\beta|} - x^{2\beta} \in \mathcal{M}(\mathbf{g})_{2t}$ and since moreover $(Rx^\alpha - x^\beta)^2 \in \mathcal{M}(\mathbf{g})_{2t}$, the above identity thus shows that $R^{|\gamma|} - x^\gamma \in \mathcal{M}(\mathbf{g})_{2t}$. Finally, for the case where $|\gamma|$ is even we use a similar argument. We write $\gamma = \alpha + \beta$ for $\alpha, \beta \in \mathbb{N}^n$ with $|\alpha| = |\beta|$ and use the identity

$$R^{|\gamma|} - x^\gamma = \frac{(x^\alpha - x^\beta)^2 + (R^{2\alpha} - x^{2\alpha}) + (R^{2\beta} - x^{2\beta})}{2}. \qquad \square$$

7 DISCUSSION

We have given algebraic and geometric conditions that guarantee polynomial-time computability of the moment-SOS hierarchy for polynomial optimization problems (POP). In the general, explicitly bounded setting, our conditions are similar to the ones considered by Raghavendra & Weitz [20] to show existence of compact sum-of-squares certificates. For full-dimensional feasible regions $S(\mathbf{g})$, we give explicit, geometric conditions, which include for instance that $S(\mathbf{g})$ either contains a small ball, a strictly feasible point of low bit-complexity, or has sufficient volume. Furthermore, we make explicit the connection between polynomial-time computability of the bound $\mathrm{mom}(f)_t$ and the existence of compact feasible solutions to the sum-of-squares formulation $\mathrm{sos}(f)_t$, which is only implicitly present in [20].

A general geometric condition

Theorem 8 applies only when the feasible region $S(\mathbf{g})$ of (POP) is a full-dimensional semialgebraic set. It would be very interesting to formulate a similar, geometric condition that guarantees polynomial-time computability of the moment-SOS hierarchy in the general case. This requires finding an appropriate analog of the second condition of Theorem 8 in the setting where $S(\mathbf{g}, \mathbf{h})$ might not be full-dimensional.

Relation to the complexity of SDP

Our present discussion relates closely to the more general study of the computational complexity of semidefinite programming. It is an open question whether SDPs can be solved to (near-)optimality in polynomial-time. Even the exact complexity of testing feasibility of SDPs is not known. We do know that in the bit-model, membership of the feasibility problem in NP and Co-NP is simultaneous [21]. (In the real number model of Blum-Shub-Smale it lies in NP ∩ Co-NP [21].)

On the positive side, polynomial-time solvability of SDPs is guaranteed when the feasible region contains an 'inner ball' and is contained in an 'outer ball' of appropriate size [6, 9]. On the negative side, there are several classes of relatively simple SDPs whose feasible solutions nonetheless have exponential bit-complexity, see, e.g., [18] and the discussion therein.

In principle, these positive and negative results on SDPs provide conditions on (SOS) and (MOM) that (partly) show when polynomial-time computation can and cannot be guaranteed. The key difference with our results is that we only impose conditions on the original polynomial optimization problem (POP), rather than on the semidefinite programs resulting from the moment/sum-of-squares relaxations.

Finding exact SOS-decompositions

In the setting of polynomial optimization, it usually suffices to find *approximate* SOS-decompositions, for which one can use (standard) SDP-solvers. The problem of finding exact SOS-decompositions is more complicated. In the general case one could in principle use, for example, quantifier-elimination algorithms [2, 22]. In the univariate case specialized algorithms have been developed [4, 23], see also [15]. We note however that none of these methods come with polynomial runtime guarantees.

ACKNOWLEDGMENTS

We thank Harold Nieuwboer for bringing to our attention the literature relating to the volume of neighborhoods of algebraic sets. We further thank the anonymous referees for their helpful comments and suggestions.

REFERENCES

[1] Saugata Basu and Antonio Lerario. 2022. Hausdorff approximations and volume of tubes of singular algebraic sets. *Math. Ann.* (2022). https://doi.org/10.1007/s00208-022-02458-w

[2] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. 2006. *Algorithms in Real Algebraic Geometry (Algorithms and Computation in Mathematics).* Springer-Verlag, Berlin, Heidelberg.

[3] Aharon Ben-Tal and Arkadi Nemirovski. 2001. *Lectures on Modern Convex Optimization.* Society for Industrial and Applied Mathematics.

[4] Sylvain Chevillard, John Harrison, Mioara Maria Joldes, and Christoph Lauter. 2011. Efficient and accurate computation of upper bounds of approximation errors. *Theoretical Computer Science* 416 (2011), 1523–1543.

[5] Etienne de Klerk and Monique Laurent. 2019. A Survey of Semidefinite Programming Approaches to the Generalized Problem of Moments and Their Error Analysis. In *World Women in Mathematics 2018: Proceedings of the First World Meeting for Women in Mathematics (WM)²*. Springer International Publishing, Cham, 17–56.

[6] Etienne de Klerk and Frank Vallentin. 2016. On the Turing Model Complexity of Interior Point Methods for Semidefinite Programming. *SIAM Journal on Optimization* 26, 3 (2016), 1944–1961. https://doi.org/10.1137/15M103114X arXiv:https://doi.org/10.1137/15M103114X

[7] Peter Gács and Laszlo Lovász. 1981. Khachiyan's algorithm for linear programming. In *Mathematical Programming at Oberwolfach*, H. König, B. Korte, and K. Ritter (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 61–68. https://doi.org/10.1007/BFb0120921

[8] Martin Grötschel, László Lovász, and Alexander Schrijver. 1981. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* 1, 2 (1981), 169–197.

[9] Martin Grötschel, László Lovász, and Alexander Schrijver. 1993. *Geometric Algorithms and Combinatorial Optimization.* Springer Berlin, Heidelberg.

[10] Fritz John. 1948. Extremum problems with inequalities as subsidiary conditions. In *Studies and Essays Presented to R. Courant on his 60th Birthday, January 8, 1948.* Interscience Publishers, Inc., New York, N.Y., 187–204.

[11] Cédric Josz and Didier Henrion. 2016. Strong duality in Lasserre's hierarchy for polynomial optimization. *Opt. Lett.* 10 (2016), 3–10.

[12] Jean-Bernard Lasserre. 2001. Global Optimization with Polynomials and the Problem of Moments. *SIAM Journal on Optimization* 11, 3 (2001), 796–817. https://doi.org/10.1137/S1052623400366802 arXiv:https://doi.org/10.1137/S1052623400366802

[13] Monique Laurent. 2009. Sums of Squares, Moment Matrices and Optimization Over Polynomials. In *Emerging Applications of Algebraic Geometry*, Mihai Putinar and Seth Sullivant (Eds.). Springer New York, New York, NY, 157–270. https://doi.org/10.1007/978-0-387-09686-5_7

[14] Martin Lotz. 2015. On the volume of tubular neighborhoods of real algebraic varieties. *Proc. Amer. Math. Soc.* 143, 5 (2015), 1875–1889.

[15] Victor Magron, Mohab Safey El Din, and Markus Schweighofer. 2019. Algorithms for weighted sum of squares decomposition of non-negative univariate polynomials. *Journal of Symbolic Computation* 93 (2019), 200–220.

[16] Ryan O'Donnell. 2017. SOS Is Not Obviously Automatizable, Even Approximately. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 67)*, Christos H. Papadimitriou (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 59:1–59:10. https://doi.org/10.4230/LIPIcs.ITCS.2017.59

[17] Pablo Parrilo. 2003. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming Series B* 96 (2003), 293–320.

[18] Gábor Pataki and Aleksandr Touzov. 2021. How do exponential size solutions arise in semidefinite programming? *www.arxiv.org/abs/2103.00041* (2021). https://doi.org/10.48550/arXiv.2103.00041

[19] Alexander Prestel and Charles N. Delzell. 2001. *Positive Polynomials – From Hilbert's 17th Problem to Real Algebra.* Springer, Berlin.

[20] Prasad Raghavendra and Benjamin Weitz. 2017. On the Bit Complexity of Sum-of-Squares Proofs. *ICALP* 80 (2017), 1–13.

[21] Motakuri V. Ramana. 1997. An exact duality theory for semidefinite programming and its complexity implications. *Mathematical Programming* 77, 1 (1997), 129–162.

[22] James Renegar. 1992. On the computational complexity and geometry of the first-order theory of the reals. Part I: Introduction. Preliminaries. The geometry of semi-algebraic sets. The decision problem for the existential theory of the reals. *Journal of Symbolic Computation* 13, 3 (1992), 255–299. https://doi.org/10.1016/S0747-7171(10)80003-3

[23] Markus Schweighofer. 1999. *Algorithmische Beweise für Nichtnegativ- und Positivstellensätze.* Master's thesis. Universität Passau.

[24] Richard Wongkew. 1993. Volumes of tubular neighbourhoods of real algebraic varieties. *Pacific J. Math.* 159, 1 (1993), 177 – 184. https://doi.org/pjm/1102634385