

LOW-RANK UNIVARIATE SUM OF SQUARES HAS NO SPURIOUS LOCAL MINIMA*

BENOÎT LEGAT[†], CHENYANG YUAN[†], AND PABLO PARRILO[†]

Abstract. We study the problem of decomposing a polynomial p into a sum of r squares by minimizing a quadratically penalized objective $f_p(\mathbf{u}) = \|\sum_{i=1}^r u_i^2 - p\|^2$. This objective is nonconvex and is equivalent to the rank- r Burer–Monteiro factorization of a semidefinite program (SDP) encoding the sum of squares decomposition. We show that for all univariate polynomials p , if $r \geq 2$, then $f_p(\mathbf{u})$ has no spurious second-order critical points, showing that all local optima are also global optima. This is in contrast to previous work showing that for general SDPs, in addition to genericity conditions, r has to be roughly the square root of the number of constraints (the degree of p) for there to be no spurious second-order critical points. Our proof uses tools from computational algebraic geometry and can be interpreted as constructing a certificate using the first- and second-order necessary conditions. We also show that by choosing a norm based on sampling equally spaced points on the circle, the gradient ∇f_p can be computed in nearly linear time using fast Fourier transforms. Experimentally we demonstrate that this method has very fast convergence using first-order optimization algorithms such as L-BFGS, with near-linear scaling to million-degree polynomials.

Key words. nonconvex optimization, sum of squares, trigonometric polynomials, Burer–Monteiro method, global landscape, semidefinite programming

MSC codes. 90C23, 90C26, 90C22

DOI. 10.1137/22M1516208

1. Introduction. Burer–Monteiro factorization [9] is a methodology to solve large-scale semidefinite programs (SDPs) by replacing positive semidefinite (PSD) variables $X \succeq 0$ with a factorization $X = UU^\top$. This automatically enforces the PSD constraint and lets us find low-rank solutions by choosing the rank of the new variable U . In addition, this factorization results in a nonlinear optimization problem that can be solved with first-order methods with fast per-iteration times, especially when $\text{rank}(U)$ is small. However, the resulting problem is nonconvex, so these methods may get stuck in local optima. We show that this will not happen to the SDP finding the sum of squares decomposition of univariate polynomials; in this setting all local optima are also global.

In this work we study SDPs arising from sum of squares optimization [32]. The ability to represent the cone of sum of squares polynomials as a SDP enables many applications in polynomial optimization, control, and relaxations of combinatorial problems [27, 5]. To determine if a polynomial $p(x) \in \mathbb{R}[x]_{2d}$ is a sum of squares, it suffices to find a feasible solution to the following SDP:

$$(1.1) \quad \begin{aligned} p(x) &= b(x)^\top X b(x), \\ X &\succeq 0, \end{aligned}$$

*Received by the editors August 16, 2022; accepted for publication (in revised form) April 18, 2023; published electronically August 8, 2023.

<https://doi.org/10.1137/22M1516208>

Funding: The first author was supported by a BAEF Postdoctoral Fellowship and the NSF grant OAC-1835443. The second author was supported by the NSF grant CCF-1565235 and AFOSR grant FA8750-19-2-1000.

[†]Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (blegat@mit.edu, ycy@mit.edu, parrilo@mit.edu).

where $b(x)$ is a suitable polynomial basis of $\mathbb{R}[x]_d$. The constraint $p(x) = b(x)^\top X b(x)$ defines an affine subspace of the space of symmetric matrices, which can be expressed by matching the coefficients of $p(x)$ with the corresponding coefficients of the polynomial $b(x)^\top X b(x)$.

Given the factorization $X = UU^\top$, where $\bar{u}_1, \dots, \bar{u}_r$ are the column vectors of U , we have the explicit sum of squares decomposition

$$(1.2) \quad p(x) = \sum_{i=1}^r u_i(x)^2,$$

where $u_i(x) = b(x)^\top \bar{u}_i$. In other words, (1.2) is a formulation for the Burer–Monteiro factorization of (1.1) independent of any particular basis. Instead of solving a SDP to find $u_i(x)$, we apply a quadratic penalty to the equality constraint (1.2) to arrive at the following nonconvex objective:

$$(1.3) \quad f_p(\mathbf{u}) = \left\| \sum_{i=1}^r u_i(x)^2 - p(x) \right\|^2,$$

where $\mathbf{u} = [u_1 \ \cdots \ u_r]$ is a vector of r degree- d polynomials and the norm is induced by any inner product on polynomials of degree- $2d$. Then a degree- $2d$ polynomial $p(x)$ is a sum of r squares if and only if

$$\min_{\mathbf{u} \in \mathbb{R}[x]_d^r} f_p(\mathbf{u}) = 0.$$

We say that \mathbf{u} is a *first-order critical point* (FOCP) of $f_p(\mathbf{u})$ when $\nabla f_p(\mathbf{u}) = 0$, where the derivatives are taken with respect to the variables in \mathbf{u} . Since for every p we can find spurious FOCPs (for example, taking $\mathbf{u} = 0$), it is essential to consider second-order necessary conditions. If \mathbf{u} also satisfies $\nabla^2 f_p(\mathbf{u}) \succeq 0$, then it is a *second-order critical point* (SOCP) of $f_p(\mathbf{u})$.

Every local minimum of a function is a SOCP, but the converse is not true when the function is nonconvex.¹ If \mathbf{u} is not a SOCP, then we can produce a descent direction using the gradient or Hessian. This leads to efficient first-order methods [20, 24] that converge to SOCPs. Thus if we can show that at every SOCP $f_p(\mathbf{u}) = 0$, these algorithms will always converge to a global minimum. A recent line of work [8, 3] has shown that for general SDPs under smoothed analysis or genericity conditions, when the rank of the factorization is above the Barvinok–Pataki bound (roughly the square root of the number of constraints), there are no spurious SOCPs. Moreover, the smoothed analysis or genericity conditions are necessary, as [3] constructed a SDP where only a full-rank factorization can guarantee no spurious SOCPs.

In this paper we consider the setting of univariate polynomial optimization, which is a class of problems with applications in signal processing, control [35, 17], and computing equilibria of polynomial games [33]. These optimization problems involving nonnegative univariate polynomials can be transformed (i.e., by a bisection on the objective value) into feasibility problems for finding the sum of squares decomposition of a univariate polynomial. The main result of our paper shows that without any additional assumptions, the rank-2 quadratic-penalized Burer–Monteiro factorization of the SDP describing the sum of squares decomposition (1.3) of a univariate polynomial has no spurious SOCPs.

¹Consider, for example, $f(x) = x^3$ at $x = 0$.

THEOREM 1.1. *For all nonnegative univariate polynomials $p(x) \in \mathbb{R}[x]_{2d}$ and any $r \geq 2$, if $\mathbf{u} \in \mathbb{R}[x]_d^r$ satisfies $\nabla f_p(\mathbf{u}) = 0$ and $\nabla^2 f_p(\mathbf{u}) \succeq 0$, then $f_p(\mathbf{u}) = 0$.*

In particular, the rank bound in our result matches the Pythagoras number for univariate polynomials [12, Example 2.13].² In comparison, applying rank bounds for general SDPs to this setting requires $r \gtrsim \sqrt{d}$ (see section 2 for more details).

Theorem 1.1 is proved in section 4, by deriving a series of increasingly stronger sufficient conditions ((C1), (C2), and (C3)) implying $f_p(\mathbf{u}) = 0$ for increasingly larger classes of $\mathbf{u} = (u_1, u_2)$, eventually proving the result for all $\mathbf{u} \in \mathbb{R}[x]^2$. To illustrate this, we first show that when $r = 2$ and u_1, u_2 are coprime, $\nabla f_p(\mathbf{u}) = 0$ implies that $f_p(\mathbf{u}) = 0$ (the precise statement and proof of this case are in subsection 4.1). Note that in this simplified setting only the first-order gradient condition is needed. By computing the gradient, $\nabla f_p(\mathbf{u}) = 0$ is equivalent to

$$\nabla f_p(\mathbf{u})(\mathbf{v}) = \langle u_1 v_1 + u_2 v_2, u_1^2 + u_2^2 - p \rangle = 0$$

for all $\mathbf{v} = (v_1, v_2) \in \mathbb{R}[x]^2$. Since u_1, u_2 are coprime, Bézout’s identity (Lemma 3.3) implies that we can find $\mathbf{v}' = (v'_1, v'_2)$ so that

$$(1.4) \quad u_1 v'_1 + u_2 v'_2 = u_1^2 + u_2^2 - p,$$

thus showing that $f_p(\mathbf{u}) = \|u_1^2 + u_2^2 - p\|^2 = 0$. However, we cannot assume a priori that u_1, u_2 are coprime as we are only given p as the input. The main technical contribution of our proof is how to handle the more involved case when u_1, u_2 share a common factor.

We can also interpret our proof as a certificate. When we choose \mathbf{v}' satisfying (1.4), we obtain the identity

$$\nabla f_p(\mathbf{u})(\mathbf{v}') = f_p(\mathbf{u}).$$

This implies that $f_p(\mathbf{u}) = 0$ when $\nabla f_p(\mathbf{u}) = 0$. In section 5 we generalize this example to our full proof of Theorem 1.1, showing for all \mathbf{u} and p how to find \mathbf{v}' and $Q \succeq 0$ satisfying the following identity:

$$\nabla f_p(\mathbf{u})(\mathbf{v}') + \langle Q, \nabla^2 f_p(\mathbf{u}) \rangle = -f_p(\mathbf{u}).$$

From this identity it is clear that if \mathbf{u} is a SOCP, then $f_p(\mathbf{u}) = 0$. This compact form of our proof allows us to easily extend our result to other problems in section 6.

Since Theorem 1.1 holds for any inner product, we can choose one that enables efficient computation of $\nabla f_p(\mathbf{u})$. When p is a degree- $2d$ univariate polynomial, an equivalent way of ensuring the constraint $p(x) = b(x)^\top X b(x)$ in (1.1) is to write $2d + 1$ constraints $p(\hat{x}_i) = b(\hat{x}_i)^\top X b(\hat{x}_i)$, where $\hat{x}_0, \dots, \hat{x}_d$ are distinct sample points. This formulation can be cast into the following least-squares objective:

$$(1.5) \quad \min_{U \in \mathbb{R}^{(2d+1) \times r}} f_p(U) = \frac{1}{2d+1} \sum_{i=1}^{2d+1} \left(\|U^\top b(\hat{x}_i)\|_2^2 - p(\hat{x}_i) \right)^2.$$

This is equivalent to choosing an inner product in (1.3) that evaluates the polynomial on $2d + 1$ points. If we choose $2d + 1$ points on the complex unit circle, we can compute $\nabla f_p(U)$ in $O(d \log d)$ time using the fast Fourier transform (FFT). In section 7, we show that this method exhibits linear convergence experimentally using unconstrained optimization algorithms such as L-BFGS and has near-linear scaling to million-degree polynomials.

²The Pythagoras number for $\Sigma[x]_{2d}$ is the smallest r such that all polynomials in $\Sigma[x]_{2d}$ can be written as a sum of r squares of polynomials in $\mathbb{R}[x]_d$.

TABLE 1

Time and iterations to convergence for sum of squares decomposition of random nonnegative trigonometric polynomials. All values are median of 50 runs (with range based on 25th and 75th percentile). Table 1(a) fixes the rank r and varies the polynomial degree, whereas (b) fixes the polynomial degree and varies the rank r .

Degree	Time (s)	Iterations	r	Time	Iters	FFT Calls
2,000	2 (1 – 2)	340 (306 – 384)	2	50	4124	20618
10,000	6 (5 – 6)	530 (497 – 592)	3	9	896	6272
20,000	9 (8 – 10)	632 (587 – 695)	4	6	530	4774
100,000	53 (46 – 59)	1126 (980 – 1248)	5	5	446	4900
200,000	160 (139 – 174)	1375 (1212 – 1532)	6	5	396	5142
1,000,000	1461 (1212 – 1532)	2303 (1934 – 2437)	7	5	374	5618

(a) Varying degree, $r = 4$ (b) Varying r , degree $2d = 10,000$

1.1. Contributions. In summary, our main contributions in this paper are:

1. Proving that the quadratic penalty form of Burer–Monteiro factorization for univariate polynomial sum of squares decomposition has no spurious SOCPs (Theorem 1.1). Our result holds where the rank of the factorization is at least 2, matching the Pythagoras number for univariate polynomials. This is in contrast to previous work requiring rank $r \gtrsim \sqrt{d}$ in addition to genericity conditions or smoothed analysis for general SDPs, or showing that no spurious local minima exist in statistical problems.
2. Developing a new framework for proving that there are no spurious SOCPs for a quadratic-penalized factorized SDP, by constructing a certificate (5.3) using the first- and second-order necessary conditions. This certificate representation helps us extend our results to projection onto the sum of squares cone (Corollary 6.1), certifying nonnegativity on intervals (Corollary 6.2) and sum of squares optimization.
3. Showing that by choosing a special norm (based on the evaluation of the polynomial on points on the unit circle), the full gradient of the objective can be computed in near-linear time using FFTs. It enables us to efficiently scale first-order methods to instances with millions of variables (Table 1a). This is possible because our result (1.1) is independent of the penalty function.

2. Background and related work. Let \mathbb{S}_n be the space of $n \times n$ symmetric matrices. Given $A_i \in \mathbb{S}_n$, we consider the standard-form semidefinite feasibility problem with variable $X \in \mathbb{S}_n$:

$$\text{(SDP)} \quad \begin{aligned} \langle A_i, X \rangle &= b_i, \\ X &\succeq 0. \end{aligned}$$

Nonconvex formulation. Burer and Monteiro [9] introduced the nonconvex reformulation $X = UU^\top$ to enforce the semidefinite constraint, where $U \in \mathbb{R}^{n \times r}$:

$$\text{(NSDP}_r) \quad \langle A_i, UU^\top \rangle = b_i.$$

This motivates the following least-squares formulation:

$$\text{(SDPLS}_r) \quad \min_{U \in \mathbb{R}^{n \times r}} \sum_i \|\langle A_i, UU^\top \rangle - b_i\|^2.$$

If in addition to the equality constraints in (SDP) one wishes to minimize the objective $\langle C, X \rangle$, the works [10, 3, 14] formulate an augmented Lagrangian problem, where the term $\lambda \langle C, UU^\top \rangle$ is added to the objective of (SDPLS_r).

General SDP rank bounds. Burer and Monteiro subsequently showed in [10] that when $r \geq n$, there are no spurious SOCPs to (SDPLS_r). This result is in fact tight and [3] constructed an explicit instance where if $r = n - 1$, one can find a SOCP that is not a global minimum. Thus for general SDP feasibility, additional conditions on the objective or analysis must be imposed. Then the rank bound can be improved to the maximum rank of extreme points of the section of the PSD cone with m affine constraints. This maximum rank is $O(\sqrt{m})$, also known as the Barvinok–Pataki bound [2, 34]. In the same work, Burer and Monteiro [10] showed that when a linear objective is added to (SDP) and (NSDP_r), if $r \gtrsim \sqrt{m}$, any local minimum of (NSDP_r) is also a local minimum of (SDP) with an additional rank- r constraint. Then they showed that such a local minimum is either an optimal extreme point or contained within the relative interior of a face of the feasible set of (SDP), which is constant with respect to the objective function.

Subsequent work [3] then showed that if C is generic enough, all local minima of (NSDP_r) are global minima of (SDP) (see Cifuentes and Moitra [14] for more references). In summary, this line of work requires generic constraints, and in addition either smoothness of the constraint set [7, 8] or smoothed analysis [3, 14, 13]. In addition, [40] showed that when r is smaller than \sqrt{m} , SOCPs are not generically optimal.

Structured SDPs. Problems such as matrix completion and matrix sensing can be expressed as instances of (SDPLS_r). There has been a lot of recent interest in studying the global landscape of matrix sensing problems [22, 4, 21]. A recent line of work [21, 22, 1] shows that for certain statistical problems aiming to recover a signal in the form of a low-rank matrix corrupted by noise (where the SDP has a rank-1 solution in the noiseless setting), there are no spurious SOCPs when the noise level is low enough. Similar results [21, 29, 39] can be obtained for matrix sensing, where a low-rank matrix is reconstructed from linear measurements called sensing operators. See [11] for a survey of these problems. In summary, for a wide range of statistical problems, local minima are also global minima. These results are either satisfied with high probability or require that the sensing operators A_i satisfy the restricted isometry property.

Sampling basis. The sampling or interpolation basis for sum of squares optimization is studied in [30] and [15]. Informally, they showed that if the sampled points are “generic” enough, the problem expressed in the sampled basis is equivalent to the original problem. This idea is also used for univariate polynomial optimization in [31, 26].

Univariate polynomials. Univariate/trigonometric polynomial optimization and their applications are studied in [41, 35, 17] and the references therein. The decomposition of a nonnegative trigonometric polynomial into a sum of squares is also known as its spectral factorization [17, Theorem 1.1]. Previous methods for spectral factorization either require finding all n roots of the polynomial, solving linear systems of order n , or using an approximate $O(N \log N)$ FFT-based algorithm by sampling $N \gg n$ points [41, 17]. Design problems involving constraints on nonnegative trigonometric polynomials can be formulated as SDPs. Due to their special structure, [35] used FFTs to speed up per-iteration complexity for interior point methods solving these SDPs to $O(n^3)$.

The set of all X satisfying (1.1) is known as the Gram spectrahedron of $p(x)$. As the bounds developed for general SDPs depend on the rank of extreme points of the Gram spectrahedra, one may wonder if this quantity can be tightly bounded (i.e., better than the Barvinok–Pataki bound) in the special case of univariate polynomials. A recent work by Scheiderer [36] showed that this is not possible. If $p(x)$ is a sufficiently general positive univariate polynomial of degree d , its Gram spectrahedron has extreme points of all ranks up to $O(\sqrt{d})$.

2.1. Notation. Let $\mathbb{R}[x]_d$ be the space of univariate polynomials of degree at most d . Let $\mathbf{u}(x) \in \mathbb{R}[x]_d^r$ be a vector of r polynomials where each $u_i(x)$ is a polynomial in $\mathbb{R}[x]_d$. Let $\sigma : \mathbb{R}[x]_d^r \rightarrow \mathbb{R}[x]_{2d}$ be the quadratic map defined by $\mathbf{v}(x) \mapsto \sum_{i=1}^r v_i(x)^2$, and $\Sigma[x]_{2d} := \text{cone}(\sigma(\mathbb{R}[x]_d^r)) \subseteq \mathbb{R}[x]_{2d}$ be the cone of sum of squares univariate polynomials of degree- $2d$. A *binary form* is a homogeneous polynomial in two variables. Let $\mathbb{R}[x]_{\mathbf{d}}$ be the space of binary forms of degree- d and $\Sigma[x]_{2\mathbf{d}} \subseteq \mathbb{R}[x]_{2\mathbf{d}}$ be the space of sum of squares binary forms of degree- $2d$. $\mathbb{R}[x]_{\mathbf{d}}$ and $\Sigma[x]_{2\mathbf{d}}$ are isomorphic to $\mathbb{R}[x]_d$ and $\Sigma[x]_{2d}$, respectively.

2.2. Univariate polynomials. Any monic univariate polynomial $p \in \mathbb{R}[x]_d$ can be uniquely factored as $p(x) = \prod_{i=1}^d (x - \alpha_i)$, where $\alpha_i \in \mathbb{C}$ are the roots of p . Given univariate polynomials p , g , and q , we define an equivalence relation $p \equiv q \pmod{g}$ if there exists $w \in \mathbb{R}[x]$ such that $p = q + wg$. We say that g is a *divisor* of p if $p \equiv 0 \pmod{g}$. In addition, if g is also a divisor of q , then we say that g is a *common divisor* of p and q . Let $\text{gcd}(p, q)$ be the *greatest common divisor* of p and q , which is the common divisor with the highest degree. By the unique factorization of p and q , $\text{gcd}(p, q)$ is unique up to multiplication by a scalar. We say that p and q are *coprime* if $\text{gcd}(p, q) = 1$.

Any binary form $p \in \mathbb{R}[x]_{\mathbf{d}}$ can be factored as $p(x_1, x_2) = \prod_{i=1}^d (\alpha_i x_1 - \beta_i x_2)$, where $(\alpha_i, \beta_i) \in \mathbb{C}^2$. This factorization is unique up to multiplication of (α_i, β_i) by a scalar. The equivalence relation $p \equiv q \pmod{g}$ and gcd on binary forms are defined analogously to those on univariate polynomials.

3. Preliminary results. Since the polynomials we are optimizing over have fixed degrees, to better keep track of this degree we will work with homogeneous polynomials (forms). In subsection 3.1 we derive expressions for the first- and second-order necessary conditions ((3.3) and (3.4)). Next, in sections 3.2 and 3.3 we review some results on binary forms and prove our main lemma, Lemma 3.7. Particularly important is our decomposition of \mathbf{u} in Proposition 3.4.

3.1. First- and second-order necessary conditions. First we derive expressions for the gradient and Hessian of $f_p(\mathbf{u})$. In addition to the binary forms we consider in this paper, the derivations in this section also hold for general multivariate forms. For any inner product on forms $\langle \cdot, \cdot \rangle : \mathbb{R}[x]_{2\mathbf{d}} \times \mathbb{R}[x]_{2\mathbf{d}} \rightarrow \mathbb{R}$ and its associated norm $\|\cdot\| : \mathbb{R}[x]_{2\mathbf{d}} \rightarrow \mathbb{R}$, the objective function is written as

$$f_p(\mathbf{u}(x)) = \left\| \sum_{j=1}^r u_j(x)^2 - p(x) \right\|^2.$$

To find the gradient and Hessian of the objective, we compute the first- and second-order terms of $f_p(\mathbf{u}(x) + \epsilon \mathbf{v}(x))$ to obtain

$$(3.1) \quad \frac{1}{4} \nabla_{\mathbf{u}} f_p(\mathbf{u}(x))(\mathbf{v}(x)) = \left\langle \sum_{j=1}^r u_j(x)v_j(x), \sum_{j=1}^r u_j(x)^2 - p(x) \right\rangle,$$

$$(3.2) \quad \frac{1}{4} \nabla_{\mathbf{u}}^2 f_p(\mathbf{u}(x))(\mathbf{v}(x), \mathbf{v}(x)) = \left\langle \sum_{j=1}^r v_j(x)^2, \sum_{j=1}^r u_j(x)^2 - p(x) \right\rangle + 2 \left\| \sum_{j=1}^r u_j(x)v_j(x) \right\|^2.$$

Given a vector of polynomials $\mathbf{u}(x) \in \mathbb{R}[x]_{\mathbf{d}_1}^r$, we define the linear map $\mathcal{A}_{\mathbf{u}(x)} : \mathbb{R}[x]_{\mathbf{d}_2}^r \rightarrow \mathbb{R}[x]_{\mathbf{d}_1+\mathbf{d}_2}$ as

$$\mathcal{A}_{\mathbf{u}(x)} : \mathbf{v}(x) \mapsto \sum_{i=1}^r u_i(x)v_i(x).$$

For example, we can write $\sigma(\mathbf{u}) = \mathcal{A}_{\mathbf{u}}(\mathbf{u})$. When $n = 2$, $\mathcal{A}_{\mathbf{u}}$ is up to a constant the map induced by the *Sylvester* matrix [38]. If $d_1 = d_2$, the determinant of $\mathcal{A}_{\mathbf{u}}$ is up to a constant the *resultant* of u_1 and u_2 ; it vanishes if and only if u_1 and u_2 share a common divisor. Next we concisely define SOCPs using this new notation.

DEFINITION 3.1. We say that $\mathbf{u} \in \mathbb{R}[x]_{\mathbf{d}}^r$ is a SOCP of $f_p(\mathbf{u})$ if its gradient is zero and its Hessian is positive semidefinite. In other words, for all $\mathbf{v} \in \mathbb{R}[x]_{\mathbf{d}}^r$,

$$(3.3) \quad \frac{1}{4} \nabla_{\mathbf{u}} f_p(\mathbf{u})(\mathbf{v}) = \langle \mathcal{A}_{\mathbf{u}}(\mathbf{v}), \sigma(\mathbf{u}) - p \rangle = 0,$$

$$(3.4) \quad \frac{1}{4} \nabla_{\mathbf{u}}^2 f_p(\mathbf{u})(\mathbf{v}, \mathbf{v}) = \langle \sigma(\mathbf{v}), \sigma(\mathbf{u}) - p \rangle + 2 \|\mathcal{A}_{\mathbf{u}}(\mathbf{v})\|^2 \geq 0.$$

Remark 3.2. The first-order condition (3.3) alone is insufficient to guarantee global optimality, even when p is generic. For example, $\mathbf{u} = 0$ is always a FOCP, but is spurious if $p \neq 0$. Since we can always construct spurious FOCPs, even when p is generic, we need to consider the second-order conditions.

3.2. Pairs of binary forms. In this section we present some results on pairs of binary forms that will be helpful for characterizing the sets³ $\text{Im}(\mathcal{A}_{\mathbf{u}})$ and $\text{cone}(\sigma(\ker(\mathcal{A}_{\mathbf{u}})))$ and proving Theorem 1.1 in section 4. From now on we assume that $r = 2$ and $\mathbf{u} = (u_1, u_2)$.

The following lemma is a restatement of Bézout’s lemma for univariate polynomials using our notation. It states that any form in $\mathbb{R}[x]_{2\mathbf{d}}$ is in the image of the map $\mathcal{A}_{\mathbf{u}}$, as long as u_1 and u_2 are coprime. We provide a proof below for completeness.

LEMMA 3.3. Given $\mathbf{u} = (u_1, u_2) \in \mathbb{R}[x]_{\mathbf{d}_1}^2$ and $d_2 \geq d_1 - 1$, consider the map $\mathcal{A}_{\mathbf{u}} : \mathbb{R}[x]_{\mathbf{d}_2}^2 \rightarrow \mathbb{R}[x]_{\mathbf{d}_1+\mathbf{d}_2}$. Then u_1 and u_2 are coprime if and only if $\text{Im}(\mathcal{A}_{\mathbf{u}}) = \mathbb{R}[x]_{\mathbf{d}_1+\mathbf{d}_2}$.

Proof. For the “if” direction, when u_1 and u_2 are not coprime, they share a common factor w with degree at least 1, and so does every form in $\text{Im}(\mathcal{A}_{\mathbf{u}})$. Thus $\text{Im}(\mathcal{A}_{\mathbf{u}})$ is not equal to $\mathbb{R}[x]_{\mathbf{d}_1+\mathbf{d}_2}$.

Next we prove the “only if” direction. Since $\mathcal{A}_{\mathbf{u}}(\mathbf{v}) = v_1u_1 + v_2u_2$, for every $\mathbf{v} \in \ker(\mathcal{A}_{\mathbf{u}})$ we have

$$(3.5) \quad u_1v_1 = -u_2v_2.$$

³For the more algebraically inclined reader, the sets $\text{Im}(\mathcal{A}_{\mathbf{u}})$ and $\ker(\mathcal{A}_{\mathbf{u}})$ are the graded parts of the ideal [16, Definition 1.4.1] and syzygy module [16, Definition 10.4.3] of \mathbf{u} , respectively.

If $d_2 = d_1 - 1$, since u_1 and u_2 are coprime, by evaluating (3.5) on the roots of u_1 and u_2 we can conclude that $v_1 = v_2 = 0$ and $\dim(\ker(\mathcal{A}_{\mathbf{u}})) = 0$. Otherwise if $d_2 \geq d_1$, this implies that u_1 is a divisor of v_2 and u_2 is a divisor of v_1 . So there exists $w_1, w_2 \in \mathbb{R}[x]_{d_2-d_1}$ such that $v_1 = w_1 u_2, v_2 = w_2 u_1$. Then (3.5) becomes $(w_1 + w_2)u_1 u_2 = 0$, implying that $w_2 = -w_1$. Thus $\ker(\mathcal{A}_{\mathbf{u}}) = \{(wu_2, -wu_1) \mid w \in \mathbb{R}[x]_{d_2-d_1}\}$ and $\dim(\ker(\mathcal{A}_{\mathbf{u}})) = \dim(\mathbb{R}[x]_{d_2-d_1})$. In summary, $d_2 \geq d_1 - 1$ implies that $\dim(\ker(\mathcal{A}_{\mathbf{u}})) = d_2 + 1 - d_1$. Therefore, by the rank-nullity theorem,

$$\dim(\text{Im}(\mathcal{A}_{\mathbf{u}})) = 2 \dim(\mathbb{R}[x]_{d_2}) - \dim(\ker(\mathcal{A}_{\mathbf{u}})) = \dim(\mathbb{R}[x]_{d_1+d_2}).$$

Since $\text{Im}(\mathcal{A}_{\mathbf{u}}) \subseteq \mathbb{R}[x]_{d_1+d_2}$ and these two sets have the same dimensions, they must be equal. \square

In particular, Lemma 3.3 motivates the decomposition $u_1 = u'_1 \hat{g}$ and $u_2 = u'_2 \hat{g}$, where u'_1 and u'_2 are coprime and $\hat{g} = \gcd(u_1, u_2)$. Indeed when $\hat{g} = 1$ we can apply the sufficient condition (C1) in subsection 4.1 to show that \mathbf{u} is not a SOCP. When $\gcd(\hat{g}, \sigma(\mathbf{u}')) = 1$ we can apply (C2) in subsection 4.2. Otherwise we need to partition the roots of $\hat{g} = gh$ by whether each root is also a root of $\sigma(\mathbf{u}')$, then apply (C3) in subsection 4.3.

PROPOSITION 3.4. *Given $\mathbf{u} = (u_1, u_2) \in \mathbb{R}[x]_{\mathbf{d}}^2$, we can always find $g \in \mathbb{R}[x]_{\mathbf{m}}$, $h \in \mathbb{R}[x]_{\mathbf{k}}$, and $\mathbf{u}' = (u'_1, u'_2) \in \mathbb{R}[x]_{\mathbf{d}-\mathbf{m}-\mathbf{k}}^2$ so that*

$$\begin{aligned} (u_1, u_2) &= (u'_1 gh, u'_2 gh), \\ \gcd(u_1, u_2) &= gh, \\ \gcd(u'_1, u'_2) &= 1, \\ \gcd(\sigma(\mathbf{u}'), g) &= \gcd(u_1'^2 + u_2'^2, g) = 1, \\ r \text{ is a (possibly complex) root of } h &\Rightarrow r \text{ is a root of } \sigma(\mathbf{u}'). \end{aligned}$$

Moreover, this decomposition is unique up to multiplication by constants.

Proof. Let $\hat{g} = \gcd(u_1, u_2)$ so that $u'_1 = u_1/\hat{g}$ and $u'_2 = u_2/\hat{g}$. By partitioning the roots of \hat{g} we can decompose $\hat{g} = gh$, where g has no common roots with $\sigma(\mathbf{u}')$ and every root of h is also a root of $\sigma(\mathbf{u}')$. The uniqueness of this decomposition follows from the unique factorization theorem for binary forms. \square

We demonstrate this decomposition with an example.

Example 3.5. Let $u_1(x_1, x_2) = (x_1^2 + x_2^2)^2(2x_1^2 + x_2^2)x_1^2$ and $u_2(x_1, x_2) = (x_1^2 + x_2^2)^2(2x_1^2 + x_2^2)x_1 x_2$. Then the decomposition in Proposition 3.4 gives $u'_1 = x_1, u'_2 = x_2, g = x_1(2x_1^2 + x_2^2)$, and $h = (x_1^2 + x_2^2)^2$.

The following observation about the roots of $\sigma(\mathbf{u}')$ and h is useful in our proofs.

PROPOSITION 3.6. *In the decomposition of Proposition 3.4, both $\sigma(\mathbf{u}')$ and h have no real roots and $\deg(h)$ is even.*

Proof. If $\sigma(\mathbf{u}')$ has a real root \hat{x} , then $u'_1(\hat{x})^2 + u'_2(\hat{x})^2 = 0$. Thus \hat{x} is a root of u'_1 and u'_2 , contradicting the fact that $\gcd(u'_1, u'_2) = 1$. Thus every root of h is complex, and h must have even degree. \square

3.3. Main lemma. In our proof of Theorem 1.1 we will use the following main result, which may be of independent interest.

LEMMA 3.7. *Given binary forms $g \in \mathbb{R}[x]_{\mathbf{m}}$, $q \in \Sigma[x]_{2\mathbf{d}-2\mathbf{m}}$, and $p \in \Sigma[x]_{2\mathbf{d}}$, if g and q are coprime, then there exists a sum of squares binary form $s \in \Sigma[x]_{2\mathbf{m}}$ such that*

$$p \equiv sq \pmod{g}.$$

The main ingredient of the proof of Lemma 3.7 is a result stating that any univariate polynomial $a(x)$ strictly positive on the real zeros of $g(x)$ can be written as a single square⁴ modulo $g(x)$.

PROPOSITION 3.8. *Let $g(x)$ and $a(x)$ be coprime univariate polynomials where $\deg(g) = m$. If $a(x) > 0$ for all $\{x \in \mathbb{R} \mid g(x) = 0\}$, then there exists a polynomial $t \in \mathbb{R}[x]_m$ such that*

$$a \equiv t^2 \pmod{g}.$$

This result is related to Schmüdgen’s certificate [37], which states that if a polynomial is strictly positive on a compact semialgebraic set, then it has a Positivstellensatz certificate in terms of the equations describing the set. We prove Proposition 3.8 using Hermite interpolation on the series expansion of $\sqrt{a(x)}$ around the roots of g .

Proof of Proposition 3.8. Let r_i be the roots (possibly complex) of $g(x)$, each with multiplicity n_i , so that $\sum_i n_i = m$. Consider the Taylor series expansion of $f(x) = \sqrt{a(x)}$ centered at r_i . Since a and g do not share any common roots as they are coprime, this Taylor series is well defined around any root of g . Let the polynomials $\gamma_i(x)$ be the first n_i terms of the Taylor expansion of $f(x)$ centered at r_i . The polynomials γ_i have real coefficients if r_i is real, and if r_i and r_j are a pair of conjugate roots, $\gamma_i = \bar{\gamma}_j$. We can then use the Chinese Remainder Theorem [18, section 7.6] to construct the unique polynomial $t(x)$ with real coefficients and $\deg(t) < m$ such that

$$\text{for all } i, \quad t(x) \equiv \gamma_i(x) \pmod{(x - r_i)^{n_i}}.$$

By construction, for all roots r_i of g and any $k = 0, \dots, n_i - 1$, we have

$$\frac{d^k}{dx^k} f(r_i) = \frac{d^k}{dx^k} \gamma_i(r_i) = \frac{d^k}{dx^k} t(r_i).$$

For each root r_i , we have

$$\sqrt{a(r_i)} = f(r_i) = \gamma_i(r_i) = t(r_i)$$

and

$$\frac{d}{dx} a(r_i) = 2f(r_i) \frac{d}{dx} f(r_i) = 2t(r_i) \frac{d}{dx} t(r_i) = \frac{d}{dx} t(r_i)^2.$$

By induction we get $\frac{d^k}{dx^k} a(r_i) = \frac{d^k}{dx^k} t(r_i)^2$ for $k = 0, \dots, n_i - 1$. This is a generalization of Hermite interpolation for a variable number of consecutive derivatives at each point [23, section 17.6].⁵ Since $a(x)$ and $t(x)^2$ match at all the roots of g (including derivatives up to the multiplicity of the root), we have shown that $a \equiv t^2 \pmod{g}$. \square

Then we prove the affine version of Lemma 3.7.

⁴The number of squares is not important in our proof of Lemma 3.7; we only need the property that $a(x)$ can be written as a sum of squares modulo $g(x)$.

⁵This is referred to as Birkhoff interpolation in [23], and the existence of a unique interpolating polynomial crucially depends on the use of consecutive derivatives.

LEMMA 3.9. *Let g, p , and q be univariate polynomials where $\deg(g) = m$, $\deg(p) = 2d$, and $\deg(q) = 2d - 2m$, p and q are sum of squares, and g and q are coprime. Then there exists a sum of squares polynomial $s \in \Sigma[x]_{2m}$ such that*

$$p \equiv sq \pmod{g}.$$

Proof. Since q is coprime with g , Lemma 3.3 (after reducing q modulo g) guarantees that there exists a polynomial $a \in \mathbb{R}[x]_m$ such that

$$(3.6) \quad aq \equiv 1 \pmod{g}.$$

We have $q(x) > 0$ for all real roots x of g , since q is nonnegative and coprime with g . Thus $a(x) > 0$ for all real roots x of g , by evaluation of (3.6) at these roots. Since $a(x)q(x) = 1$ for all roots x of g , a is also coprime with g . Then we can apply Proposition 3.8 to find $t \in \mathbb{R}[x]_m$ so that $a \equiv t^2 \pmod{g}$. Multiplying both sides of (3.6) by p , we get

$$p \equiv t^2 pq \pmod{g}.$$

Since $t^2 p$ is a sum of squares, we can reduce each squared polynomial modulo g to get $s \in \Sigma[x]_{2m}$. \square

Finally we prove Lemma 3.7, which is the projective version of Lemma 3.9.

Proof of Lemma 3.7. We first apply a linear change of coordinates so that $(0, 1)$ is not a root of g , p , or q . Then let $g'(x) = g(x, 1)$, $p'(x) = p(x, 1)$, and $q'(x) = q(x, 1)$. Since this dehomogenization procedure preserves the degree of g , p , and q ,⁶ we can apply Lemma 3.9 to find polynomials $s' \in \Sigma[x]_{2m}$ and $t' \in \mathbb{R}[x]_{2d-m}$ so that

$$p' = s'q' + t'g'.$$

We can then homogenize by letting $s(x_1, x_2) = x_2^{2m}s'(x_1/x_2)$ and $t(x_1, x_2) = x_2^{2d-m}t'(x_1/x_2)$. Thus s is also a sum of squares and

$$p = sq + tg. \quad \square$$

4. Main theorem and proof. In this section we prove Theorem 1.1, which states that for univariate polynomials, a rank-2 decomposition has no spurious second-order critical points. Using the decomposition in Proposition 3.4, we first prove simplified versions of Theorem 1.1 in sections 4.1 and 4.2, before proving the full version in subsection 4.3.

4.1. Coprime case: $g = 1$, $h = 1$. This is the case explained in the introduction. In the decomposition of Proposition 3.4, $g = h = 1$ implies that u_1 and u_2 are coprime. This happens generically and implies that for a fixed p , the gradient condition (3.3) is sufficient for almost all \mathbf{u} .

PROPOSITION 4.1. *Suppose $\mathbf{u} \in \mathbb{R}[x]_{\mathbf{d}}^2$ and $p \in \Sigma[x]_{2\mathbf{d}}$ satisfies $\nabla f_p(\mathbf{u}) = 0$. If*

$$(C1) \quad p \in \text{Im}(\mathcal{A}_{\mathbf{u}}),$$

then we have $f_p(\mathbf{u}) = 0$.

⁶If the degree of g is not preserved after dehomogenization, the degree of t' after applying Lemma 3.9 could be larger than $2d - m$. For example, if $d = m = 2$, $g = x_1x_2$, $p = (2x_1^2 + x_2^2)x_2^2$, and $q = 1$, we get that $s' = (x^2 + 1)^2$ and $t' = -x^3$ after dehomogenizing and applying Lemma 3.9. This issue will not occur if the dehomogenization is degree-preserving.

Proof. Since $p \in \text{Im}(\mathcal{A}_{\mathbf{u}})$, we can find $\mathbf{v} \in \mathbb{R}[x]_{\mathbf{d}}^2$ so that $\mathcal{A}_{\mathbf{u}}(\mathbf{v}) = \sigma(\mathbf{u}) - p$. Evaluating the gradient condition (3.3) at \mathbf{v} , we conclude that $f_p(\mathbf{u}) = \|\sigma(\mathbf{u}) - p\|^2 = 0$. \square

Since Lemma 3.3 implies that $\text{Im}(\mathcal{A}_{\mathbf{u}}) = \mathbb{R}[x]_{2\mathbf{d}}$ if and only if u_1 and u_2 are coprime, we have shown that when $g = h = 1$, we always have $p \in \text{Im}(\mathcal{A}_{\mathbf{u}})$ and there are no spurious FOCs and SOCPs.

4.2. Special case: $h = 1$. When u_1 and u_2 are not coprime, $\sigma(\mathbf{u}) - p$ might not be in $\text{Im}(\mathcal{A}_{\mathbf{u}})$ and we cannot use the argument in Proposition 4.1. Thus we need to use make use of the Hessian condition (3.4).

PROPOSITION 4.2. *Suppose $\mathbf{u} \in \mathbb{R}[x]_{\mathbf{d}}^2$ and $p \in \Sigma[x]_{2\mathbf{d}}$ satisfies $\nabla f_p(\mathbf{u}) = 0$ and $\nabla^2 f_p(\mathbf{u}) \succeq 0$. If*

$$(C2) \quad p \in \text{Im}(\mathcal{A}_{\mathbf{u}}) + \text{cone}(\sigma(\ker(\mathcal{A}_{\mathbf{u}}))),$$

then we have $f_p(\mathbf{u}) = 0$.

This means that if for all \mathbf{u} we can decompose $p = q + r$, where $q \in \text{Im}(\mathcal{A}_{\mathbf{u}})$ and $r \in \text{cone}(\sigma(\ker(\mathcal{A}_{\mathbf{u}})))$, then $f_p(\mathbf{u})$ has no spurious SOCPs. In particular, similar to how $\text{Im}(\mathcal{A}_{\mathbf{u}})$ is related to the gradient condition (3.3) in Proposition 4.1, $\text{cone}(\sigma(\ker(\mathcal{A}_{\mathbf{u}})))$ is related to the Hessian condition (3.4). The following result states that (C2) is satisfied if $h = 1$.

LEMMA 4.3. *Given $\mathbf{u} \in \mathbb{R}[x]_{\mathbf{d}}^2$, if in the decomposition of Proposition 3.4 we have $h = 1$, then for every $p \in \Sigma[x]_{2\mathbf{d}}$,*

$$p \in \text{Im}(\mathcal{A}_{\mathbf{u}}) + \text{cone}(\sigma(\ker(\mathcal{A}_{\mathbf{u}}))).$$

Proof. We want to show that any $p(x) \in \Sigma[x]_{2\mathbf{d}}$ can be written as the sum of polynomials in $\text{Im}(\mathcal{A}_{\mathbf{u}})$ and $\sigma(\ker(\mathcal{A}_{\mathbf{u}}))$. Therefore it is useful to have a characterization of these sets. Lemma 3.3 tells us that

$$\begin{aligned} \text{Im}(\mathcal{A}_{\mathbf{u}}) &= \{\mathcal{A}_{\mathbf{u}'}(\mathbf{v})g \mid \mathbf{v} \in \mathbb{R}[x]_{\mathbf{d}}\} \\ &= \{wg \mid w \in \mathbb{R}[x]_{2\mathbf{d}-\mathbf{m}}\}. \end{aligned}$$

Since for all $t \in \mathbb{R}[x]_{\mathbf{m}}$ we have $(-tu'_2, tu'_1) \in \ker(\mathcal{A}_{\mathbf{u}})$,

$$(4.1) \quad \begin{aligned} \{t^2\sigma(\mathbf{u}') \mid t \in \mathbb{R}[x]_{\mathbf{m}}\} &\subseteq \sigma(\ker(\mathcal{A}_{\mathbf{u}})), \\ \{s\sigma(\mathbf{u}') \mid s \in \Sigma[x]_{2\mathbf{m}}\} &\subseteq \text{cone}(\sigma(\ker(\mathcal{A}_{\mathbf{u}}))). \end{aligned}$$

Since $\sigma(\mathbf{u}')$ is coprime with g by assuming $h = 1$, we can apply Lemma 3.7 to show that there exists $w \in \mathbb{R}[x]_{2\mathbf{d}-\mathbf{m}}$ and $s \in \Sigma[x]_{2\mathbf{m}}$ such that $p = s\sigma(\mathbf{u}') + wg$. \square

Finally we prove Proposition 4.2.

Proof of Proposition 4.2. The condition (C2) implies that there exist $\mathbf{v} \in \mathbb{R}[x]_{\mathbf{d}}^2, \mathbf{w}^{(i)} \in \ker(\mathcal{A}_{\mathbf{u}})$ such that

$$p = \mathcal{A}_{\mathbf{u}}(\mathbf{v}) + \sum_i \sigma(\mathbf{w}^{(i)}).$$

Since $\nabla f_p(\mathbf{u}) = 0$, (3.3) implies that

$$(4.2) \quad \langle \mathcal{A}_{\mathbf{u}}(\mathbf{v}), \sigma(\mathbf{u}) - p \rangle = 0.$$

Since $\nabla^2 f_p(\mathbf{u}) \succeq 0$ and $\mathbf{w}^{(i)} \in \ker(\mathcal{A}_{\mathbf{u}})$, (3.4) implies that

$$(4.3) \quad \langle \sigma(\mathbf{w}^{(i)}), \sigma(\mathbf{u}) - p \rangle \geq 0.$$

Combining (4.2) and (4.3) gives

$$(4.4) \quad \langle p, \sigma(\mathbf{u}) - p \rangle \geq 0.$$

Since $\nabla f_p(\mathbf{u}) = 0$ implies that

$$\langle \mathcal{A}_{\mathbf{u}}(\mathbf{u}), \sigma(\mathbf{u}) - p \rangle = \langle \sigma(\mathbf{u}), \sigma(\mathbf{u}) - p \rangle = 0,$$

we have

$$f_p(\mathbf{u}) = \|\sigma(\mathbf{u}) - p\|^2 = \langle \sigma(\mathbf{u}) - p, \sigma(\mathbf{u}) - p \rangle = -\langle p, \sigma(\mathbf{u}) - p \rangle.$$

This together with (4.4) implies that $f_p(\mathbf{u}) \leq 0$. However $f_p(\mathbf{u})$ is always nonnegative, and thus it must be 0. \square

4.3. General case. Lemma 4.3 alone is insufficient to prove Theorem 1.1. It is possible for $\hat{g} = \gcd(u_1, u_2)$ to share complex roots with $\sigma(\mathbf{u}')$ (recall from Proposition 3.6 that all roots of $\sigma(\mathbf{u}')$ are complex), as seen in Example 3.5. Hence the argument in the proof of Lemma 4.3 fails as \hat{g} is not coprime with $\sigma(\mathbf{u}')$. To get around this issue, we will derive the sufficient condition (C3) in Proposition 4.4, a stronger version of (C2), by carefully examining the Hessian condition (3.4). Roughly speaking, Proposition 4.4 shows that we can replace every root of h (which must be complex) with any real root.⁷ Since $\sigma(\mathbf{u}')$ has no real roots, $\hat{g}x_1^k/h$ is now coprime with $\sigma(\mathbf{u}')$, and we can then complete the proof by following the argument in the previous section.

PROPOSITION 4.4. *Suppose $\mathbf{u} \in \mathbb{R}[x]_{\mathbf{d}}^2$ and $p \in \Sigma[x]_{2\mathbf{d}}$ satisfies $\nabla f_p(\mathbf{u}) = 0$ and $\nabla^2 f_p(\mathbf{u}) \succeq 0$, with the decomposition in Proposition 3.4, where $k = \deg(h)$ and $\mathbf{u}x_1^k/h = (u'_1gx_1^k, u'_2gx_1^k)$. If*

$$(C3) \quad p \in \text{Im} \left(\mathcal{A}_{\mathbf{u}x_1^k/h} \right) + \text{cone} \left(\Sigma \left(\ker(\mathcal{A}_{\mathbf{u}}) \right) \right),$$

then $f_p(\mathbf{u}) = 0$.

Proof. We first prove that if $r \in \mathbb{R}[x]_{\ell}$ is a common divisor of $\sigma(\mathbf{u}')$ and $\hat{g} = \gcd(u_1, u_2)$, then

$$(4.5) \quad \langle \mathcal{A}_{\mathbf{u}}(\mathbf{b})x_1^{\ell}/r, \sigma(\mathbf{u}) - p \rangle = 0 \text{ for all } \mathbf{b} = (b_1, b_2) \in \mathbb{R}[x]_{\mathbf{d}}^2.$$

Given any $b_1, b_2 \in \mathbb{R}[x]_{\mathbf{d}}$ and $\eta \in \mathbb{R}$, let $v_1 = \eta x_1^{\ell} u_2 / r + b_2$ and $v_2 = -\eta x_1^{\ell} u_1 / r - b_1$. We have

$$(4.6) \quad \mathcal{A}_{\mathbf{u}}(\mathbf{v}) = u_1(\eta x_1^{\ell} u_2 / r + b_2) - u_2(\eta x_1^{\ell} u_1 / r + b_1) = u_1 b_2 - u_2 b_1,$$

$$(4.7) \quad \sigma(\mathbf{v}) = \eta^2 x_1^{2\ell} (u_1^2 + u_2^2) / r^2 + 2\eta x_1^{\ell} (b_1 u_1 + b_2 u_2) / r + (b_1^2 + b_2^2).$$

Since r is a divisor of both $\sigma(\mathbf{u}')$ and \hat{g} , $x_1^{2\ell} (u_1^2 + u_2^2) / r^2 = x_1^{2\ell} \sigma(\mathbf{u}') \hat{g}^2 / r^2$ is a multiple of \hat{g} . Thus $x_1^{2\ell} (u_1^2 + u_2^2) / r^2 \in \text{Im}(\mathcal{A}_{\mathbf{u}})$ and we have

$$\langle \eta^2 x_1^{2\ell} (u_1^2 + u_2^2) / r^2, \sigma(\mathbf{u}) - p \rangle = 0.$$

⁷Without loss of generality we choose this real root to be x_1 .

Therefore, the Hessian condition (3.4) implies that for all $\eta \in \mathbb{R}$,

$$(4.8) \quad 2\eta \langle (b_1u_1 + b_2u_2)x_1^\ell/r, \sigma(\mathbf{u}) - p \rangle + \langle b_1^2 + b_2^2, \sigma(\mathbf{u}) - p \rangle + 2\|u_1b_2 - u_2b_1\|^2 \geq 0.$$

This implies the identity (4.5); otherwise there exists η such that (4.8) is negative.

Since $\sigma(\mathbf{u}')$ and h have no real roots (Proposition 3.6), we can write $h = \prod_{i=1}^{k/2} r_i$, where each $r_i \in \mathbb{R}[x]_{\mathbf{2}}$ is a quadratic form corresponding to the product of a pair of complex roots. We first apply the same argument from above to show that

$$(4.9) \quad \langle \mathcal{A}_{\mathbf{u}}(\mathbf{b})x_1^2/r_1, \sigma(\mathbf{u}) - p \rangle = 0 \text{ for all } \mathbf{b} \in \mathbb{R}[x]_{\mathbf{d}}^2.$$

Next we show that (4.9) implies that

$$(4.10) \quad \langle \mathcal{A}_{\mathbf{u}}(\mathbf{b})x_1^4/(r_1r_2), \sigma(\mathbf{u}) - p \rangle = 0 \text{ for all } \mathbf{b} \in \mathbb{R}[x]_{\mathbf{d}}^2.$$

Similar to before, let $r = r_1r_2$ so we have the identities (4.6) and (4.7) as before. Since r_2 is a divisor of both $\sigma(\mathbf{u}')$ and \hat{g}/r_1 , $x_1^8\sigma(\mathbf{u}')\hat{g}^2/(r_1r_2)^2 = x_1^8\frac{\sigma(\mathbf{u}')}{r_2}\frac{\hat{g}}{r_1r_2}\frac{\hat{g}}{r_1}$ is a multiple of \hat{g}/r_1 . Thus $x_1^8\sigma(\mathbf{u}')\hat{g}^2/(r_1r_2)^2 \in \text{Im}(\mathcal{A}_{\mathbf{u}x_1^2/r_1})$ and we then use (4.9) to show (4.10).

Thus by iteratively applying the previous arguments,⁸ we show that for every $1 \leq k' \leq k/2$ and $\mathbf{b} \in \mathbb{R}[x]_{\mathbf{d}}^2$,

$$\left\langle \mathcal{A}_{\mathbf{u}}(\mathbf{b})x_1^{2k'} / \prod_{i=1}^{k'} r_i, \sigma(\mathbf{u}) - p \right\rangle = 0.$$


This is because each r_i is a divisor of $\sigma(\mathbf{u}')$ and $\prod_{i=1}^{k'} r_i$ divides \hat{g} . From here we can finish our proof by following the same steps as in the proof of Proposition 4.2. \square

With Proposition 4.4 we can prove Theorem 1.1, by showing that every $p \in \Sigma[x]_{\mathbf{2d}}$ has the required decomposition.

Proof of Theorem 1.1. Since $\langle \mathcal{A}_{\mathbf{u}'}(\mathbf{b})gx_1^k, \sigma(\mathbf{u}) - p \rangle = 0$ for all $\mathbf{b} \in \mathbb{R}[x]_{\mathbf{d}}^2$ and u'_1, u'_2 are coprime, Lemma 3.3 implies that $\langle wgx_1^k, \sigma(\mathbf{u}) - p \rangle = 0$ for all $w \in \mathbb{R}[x]$. Since $\sigma(\mathbf{u}')$ has no real roots (Proposition 3.6), it is coprime with x_1^k . As $\sigma(\mathbf{u}')$ is coprime with g , it is also coprime with gx_1^k . Thus Lemma 3.7 tells us that there exists a sum of squares polynomial s such that $p \equiv s\sigma(\mathbf{u}') \pmod{gx_1^k}$. Since $s\sigma(\mathbf{u}') \in \text{cone}(\sigma(\ker(\mathcal{A}_{\mathbf{u}})))$ by (4.1), we are done. \square

5. Geometric interpretation and certificates. In this section, we provide a geometric interpretation of our proof of Theorem 1.1, which allows us to turn the proof into a certificate. In order to prove that there are no spurious second-order critical points when minimizing $f_p(\mathbf{u}) = \|\sigma(\mathbf{u}) - p\|^2$, we have to show that for all $\mathbf{u} \in \mathbb{R}[x]^r$ and for all $p \in \Sigma[x]$, $\nabla f_p(\mathbf{u}) = 0$ and $\nabla^2 f_p(\mathbf{u}) \succeq 0$ imply that $f_p(\mathbf{u}) = 0$ and $p = \sigma(\mathbf{u})$. One way to tackle this problem is to fix p then characterize the set of \mathbf{u} satisfying the second-order critical point conditions. This is the approach taken by [3] and related works, where they used an argument based on the dimension of the subspace generated by the constraints of the SDP. However, the SOCP conditions

⁸The argument here is subtle because although every root of h is a root of $\sigma(\mathbf{u}')$, a root may have higher multiplicity in h than in $\sigma(\mathbf{u}')$. For example, it is possible that $h = (x_1^2 + x_2^2)^2$ but $\sigma(\mathbf{u}') = x_1^2 + x_2^2$. In this case, to obtain (4.5) we need to iteratively “peel off” the factors $r_1 = r_2 = x_1^2 + x_2^2$, by first proving (4.9) and then proving (4.10).

FIG. 1. *The geometric interpretation.*

are nonconvex in \mathbf{u} . In order to do better than a dimension-counting argument, our proof takes a different approach. If we fix \mathbf{u} , the set of all p satisfying the gradient condition (3.3) is an affine subspace, whereas the set of all $p \in \Sigma[x]$ satisfying the Hessian condition (3.4) is a convex semidefinite-representable set. We need to show that these two sets intersect at only one point, $p = \sigma(\mathbf{u})$ (see Figure 1).

Our proof can be interpreted as constructing a certificate to show that these two sets only intersect at one point. This is true if and only if the following optimization problem has a zero optimal objective value:

$$(5.1) \quad \max_{p \in \Sigma[x]} \min_{Q \succeq 0, \lambda \in \mathbb{R}[x]^r} \|\sigma(\mathbf{u}) - p\|^2 + \nabla f_p(\mathbf{u})(\lambda) + \langle Q, \nabla^2 f_p(\mathbf{u}) \rangle.$$

Expanding the gradient and Hessian, we get

$$\begin{aligned} \nabla f_p(\mathbf{u})(\lambda) &= \langle \mathcal{A}_{\mathbf{u}}(\lambda), \sigma(\mathbf{u}) - p \rangle, \\ \langle Q, \nabla^2 f_p(\mathbf{u}) \rangle &= \sum_i \langle \sigma(\mathbf{v}^{(i)}), \sigma(\mathbf{u}) - p \rangle + 2 \left\| \mathcal{A}_{\mathbf{u}}(\mathbf{v}^{(i)}) \right\|^2, \end{aligned}$$

where $Q = \sum_i \mathbf{v}^{(i)} \mathbf{v}^{(i)\top}$. If for every $p \in \Sigma[x]$ we can find λ and Q such that

$$(5.2) \quad \nabla f_p(\mathbf{u})(\lambda) + \langle Q, \nabla^2 f_p(\mathbf{u}) \rangle = -\|\sigma(\mathbf{u}) - p\|^2,$$

then the objective of (5.1) is at most 0 and cannot be positive, showing that $p = \sigma(\mathbf{u})$ is the only point satisfying the gradient and Hessian conditions. Since $\mathcal{A}_{\mathbf{u}}(\mathbf{u}) = \sigma(\mathbf{u})$, this is equivalent to finding λ and Q such that $\nabla f_p(\mathbf{u})(\lambda) + \langle Q, \nabla^2 f_p(\mathbf{u}) \rangle = \langle p, \sigma(\mathbf{u}) - p \rangle$.

5.1. Warmup. As a warmup, we construct such a certificate if $h = 1$ in the decomposition of \mathbf{u} in Proposition 3.4. Recall that in this case $g = \gcd(u_1, u_2)$, $u_1 = gu'_1$, $u_2 = gu'_2$, and $\sigma(\mathbf{u}')$ is coprime with $g \in \mathbb{R}[x]_{\mathbf{m}}$. Therefore, by Lemma 3.7, there exists $s \in \Sigma[x]_{2\mathbf{m}}$ such that $p \equiv s\sigma(\mathbf{u}') \pmod{g}$. Let

$$Q = s \begin{bmatrix} u_2'^2 & -u_1' u_2' \\ -u_1' u_2' & u_1'^2 \end{bmatrix}.$$

As both $\sigma(\mathbf{u})$ and $p - \sigma(\mathbf{u}')$ are divisible by g , Lemma 3.3 implies that there exists λ such that:

$$\mathcal{A}_{\mathbf{u}}(\lambda) = -\sigma(\mathbf{u}) + p - s\sigma(\mathbf{u}').$$

These values of λ and Q give

$$\begin{aligned} \nabla f_p(\mathbf{u})(\lambda) &= -\|\sigma(\mathbf{u}) - p\|^2 - \langle s\sigma(\mathbf{u}'), \sigma(\mathbf{u}) - p \rangle, \\ \langle Q, \nabla^2 f_p(\mathbf{u}) \rangle &= \langle s\sigma(\mathbf{u}'), \sigma(\mathbf{u}) - p \rangle, \end{aligned}$$

and hence taking the sum we have the identity (5.2).

5.2. Certificate. Now we can present the proof of Theorem 1.1 in the form of a certificate. First we decompose $\mathbf{u} \in \mathbb{R}[x]_{\mathbf{d}}^2$ as in Proposition 3.4. Since g is coprime with $\sigma(\mathbf{u}')$ and $\sigma(\mathbf{u}')$ has no real roots, gx_1^k is also coprime with $\sigma(\mathbf{u}')$. Then by Lemma 3.7, there exists $s \in \Sigma[x]_{\mathbf{2}(\mathbf{m}+\mathbf{k})}$ such that $p \equiv s\sigma(\mathbf{u}') \pmod{gx_1^k}$. Next we apply Lemma 3.3 to find $\mathbf{b}^0 \in \mathbb{R}[x]_{\mathbf{d}}^2$ so that

$$2\frac{x_1^k}{h}\mathcal{A}_{\mathbf{u}}(\mathbf{b}^0) = 2gx_1^k\mathcal{A}_{\mathbf{u}'}(\mathbf{b}^0) = p - s\sigma(\mathbf{u}').$$

As in the proof of Proposition 4.4, we write $h = \prod_{i=1}^{k/2} r_i$. For every $1 \leq j \leq k/2$, since r_j divides $\sigma(\mathbf{u}')$, by Lemma 3.3 there exists $\mathbf{b}^j \in \mathbb{R}[x]_{\mathbf{d}}^2$ so that

$$2gx_1^{k-2j}\mathcal{A}_{\mathbf{u}'}(\mathbf{b}^j) \prod_{i=1}^j r_i = -g^2x_1^{2k-4(j-1)}\sigma(\mathbf{u}') \prod_{i=1}^{j-1} r_i^2.$$

Given any $\mathbf{a} = (a_1, a_2) \in \mathbb{R}[x]_{\mathbf{d}}^2$, we define $\bar{\mathbf{a}} := (a_2, -a_1)$. Given a parameter $\eta \in \mathbb{R}$, for every $0 \leq j \leq k/2$ let $\eta_j = \eta^{3^j}$ and

$$\mathbf{v}^j = \eta_j^{3/2}gx_1^{k-2j}\bar{\mathbf{u}}' \prod_{i=1}^j r_i + \eta_j^{-1/2}\bar{\mathbf{b}}^j.$$

Then define

$$\begin{aligned} Q &= s\bar{\mathbf{u}}'\bar{\mathbf{u}}'^{\top} + \frac{1}{\eta} \sum_{j=0}^{k/2} \mathbf{v}^j \mathbf{v}^{j\top}, \\ \lambda &= -(1 + \eta^{-1}\eta_{k/2+1})\mathbf{u}. \end{aligned}$$

Since $\mathcal{A}_{\mathbf{u}}(\bar{\mathbf{u}}') = 0$, $\sigma(\bar{\mathbf{u}}') = \sigma(\mathbf{u}')$, $\sigma(\bar{\mathbf{b}}^j) = \sigma(\mathbf{b}^j)$, $\mathcal{A}_{\bar{\mathbf{u}}'}(\bar{\mathbf{b}}^j) = \mathcal{A}_{\mathbf{u}'}(\mathbf{b}^j)$, and $\eta_{j+1} = \eta_j^3$, we have

$$\begin{aligned} \frac{1}{\eta} \sum_{j=0}^{k/2} \sigma(\mathbf{v}^j) &= \frac{1}{\eta} \sum_{j=0}^{k/2} \left(\eta_j^3 g^2 x_1^{2k-4j} \sigma(\mathbf{u}') \prod_{i=1}^j r_i^2 + 2\eta_j g x_1^{k-2j} \mathcal{A}_{\mathbf{u}'}(\mathbf{b}^j) \prod_{i=1}^j r_i + \eta_j^{-1} \sigma(\mathbf{b}^j) \right) \\ &= p - s\sigma(\mathbf{u}') + \eta^{-1}\eta_{k/2+1}\sigma(\mathbf{u}) + \eta^{-1} \sum_{j=0}^{k/2} \eta_j^{-1} \sigma(\mathbf{b}^j), \end{aligned}$$

$$\mathcal{A}_{\mathbf{u}}(\mathbf{v}^j) = \eta_j^{-1/2} \mathcal{A}_{\mathbf{u}}(\bar{\mathbf{b}}^j),$$

$$\begin{aligned} \langle Q, \nabla^2 f_p(\mathbf{u}) \rangle &= \left\langle p + \eta^{-1}\eta_{k/2+1}\sigma(\mathbf{u}) + \eta^{-1} \sum_{j=0}^{k/2} \eta_j^{-1} \sigma(\mathbf{b}^j), \sigma(\mathbf{u}) - p \right\rangle \\ &\quad + \sum_{j=0}^{k/2} 2\eta_j^{-1}\eta^{-1} \left\| \mathcal{A}_{\mathbf{u}}(\bar{\mathbf{b}}^j) \right\|^2, \end{aligned}$$

$$\nabla f_p(\mathbf{u})(\lambda) = -(1 + \eta^{-1}\eta_{k/2+1}) \langle \sigma(\mathbf{u}), \sigma(\mathbf{u}) - p \rangle.$$

So we have proven the identity

$$(5.3) \quad \nabla f_p(\mathbf{u})(\boldsymbol{\lambda}) + \langle Q, \nabla^2 f_p(\mathbf{u}) \rangle = -\|\sigma(\mathbf{u}) - p\|^2 + \sum_{j=0}^{k/2} \frac{1}{\eta \eta_j} \left(\langle \sigma(\mathbf{b}^j), \sigma(\mathbf{u}) - p \rangle + 2 \|\mathcal{A}_{\mathbf{u}}(\bar{\mathbf{b}}^j)\|^2 \right).$$

This implies that for every $\eta > 0$ and every \mathbf{u} that satisfies $\nabla f_p(\mathbf{u}) = 0$ and $\nabla^2 f_p(\mathbf{u}) \succeq 0$,

$$\|\sigma(\mathbf{u}) - p\|^2 \leq \sum_{j=0}^{k/2} \eta^{-(3^j+1)} \left(\langle \sigma(\mathbf{b}^j), \sigma(\mathbf{u}) - p \rangle + 2 \|\mathcal{A}_{\mathbf{u}}(\bar{\mathbf{b}}^j)\|^2 \right).$$

Since \mathbf{b}^j does not depend on η , we can make the right hand side arbitrarily small by taking the limit $\eta \rightarrow \infty$. Thus we can conclude that $\|\sigma(\mathbf{u}) - p\| = 0$.

6. Extensions and generalizations. The certificate interpretation discussed in the previous section allows us to generalize Theorem 1.1 to other settings, such as projecting onto the sum of squares cone, certifying nonnegativity on intervals, and imposing linear constraints on coefficients of univariate sum of squares polynomials.

6.1. Projection onto the sum of squares cone. A natural question to consider is what happens to the optimization landscape of $f_p(\mathbf{u})$ when p cannot be expressed as a sum of squares. In this case the objective $f_p(\mathbf{u})$ can never be zero, but we show that all SOCPs have the same objective value, which is the projection of p to the sum of squares cone.

COROLLARY 6.1. *For all $\mathbf{u} \in \mathbb{R}[x]_{\mathbf{d}}^2$ where $\nabla f_p(\mathbf{u}) = 0$ and $\nabla^2 f_p(\mathbf{u}) \succeq 0$, $\sigma(\mathbf{u})$ is the projection of p to the sum of squares cone with respect to the inner product used to define f_p . In other words, $f_p(\mathbf{u}) = \|\sigma(\mathbf{u}) - p\|^2 \leq \|q - p\|^2$ for all $q \in \Sigma[x]_{2\mathbf{d}}$.*

Proof. Corollary 6.1 can be proved by a simple modification of the certificate (5.3). Although p is no longer a sum of squares, we can use (5.3) to show that for all $q \in \Sigma[x]_{2\mathbf{d}}$,

$$\langle \sigma(u) - q, \sigma(u) - p \rangle \leq 0.$$

This is exactly the variational characterization of projection onto the convex cone $\Sigma[x]_{2\mathbf{d}}$. \square

6.2. Certifying nonnegativity on intervals. Suppose we wish to certify that a univariate polynomial $p(x)$ is nonnegative in a union of intervals $I = \bigcup_{i=1}^m I_i$ where $I_i = \{x \in \mathbb{R} \mid \alpha_i \leq x \leq \beta_i\}$. This can be accomplished by finding a decomposition

$$p(x) = \sum_i^m a_i(x) q_i(x),$$

where $a_i(x)$ are fixed polynomials depending on the intervals I_i and $q_i(x)$ are sum of squares polynomials (see, e.g., [5, Theorem 3.72]). This objective can also be written as a nonconvex optimization problem by the decomposition $q_i(x) = \sum_{j=1}^r u_{ij}(x)^2 = \sigma(\mathbf{u}_i)$. If we let

$$s(\mathbf{u})(x) = \sum_{i=1}^m a_i(x) \sigma(\mathbf{u}_i),$$

then the objective $f_p^I(\mathbf{u})$ and its gradient and Hessian can be written as

$$\begin{aligned}
 f_p^I(\mathbf{u}) &= \left\| \sum_{i=1}^m a_i \sigma(\mathbf{u}_i) - p \right\|^2 = \|s(\mathbf{u}) - p\|^2, \\
 \nabla f_p^I(\mathbf{u})(\mathbf{v}) &= \left\langle \sum_{i=1}^m a_i \mathcal{A}_{\mathbf{u}_i}(\mathbf{v}_i), s(\mathbf{u}) - p \right\rangle, \\
 \nabla^2 f_p^I(\mathbf{u})(\mathbf{v}, \mathbf{v}) &= \left\langle \sum_{i=1}^m a_i \sigma(\mathbf{v}_i), s(\mathbf{u}) - p \right\rangle + 2 \left\| \sum_{i=1}^m a_i \mathcal{A}_{\mathbf{u}_i}(\mathbf{v}_i) \right\|^2.
 \end{aligned}$$

COROLLARY 6.2. *Suppose $r \geq 2$ and we are given $p = \sum_{i=1}^m a_i q_i \in \mathbb{R}[x]_{2\mathbf{d}+\mathbf{k}}$, where $a_i \in \mathbb{R}[x]_{\mathbf{k}}$, $q_i \in \Sigma[x]_{2\mathbf{d}}$. For all $\mathbf{u} \in \mathbb{R}[x]_{\mathbf{d}}^{m \times r}$ such that $\nabla f_p^I(\mathbf{u}) = 0$ and $\nabla f_p^I(\mathbf{u}) \succeq 0$, $f_p^I(\mathbf{u}) = 0$.*

Proof. We can prove this by constructing a certificate of the form (5.3). For each i we can choose $\mathbf{v}_j = 0$ for all $j \neq i$, then follow the reasoning in subsection 5.2 to find λ_i and Q_i such that for all $\eta_i > 0$,

$$(6.1) \quad \nabla f_p^I(\mathbf{u})(\lambda_i) + \langle Q_i, \nabla^2 f_p^I(\mathbf{u}) \rangle = \langle a_i q_i, s(\mathbf{u}) - p \rangle + C_i,$$

where C_i is a value that can be made arbitrarily small by taking a limit. We then sum (6.1) for all i , along with the equality $\nabla f_p^I(\mathbf{u})(-\mathbf{u}) = -\langle s(\mathbf{u}), s(\mathbf{u}) - p \rangle$ to get $\|s(\mathbf{u}) - p\|^2 \leq \sum_i C_i$, which implies that $f_p^I(\mathbf{u}) = \|s(\mathbf{u}) - p\|^2 = 0$. \square

6.3. Sum of squares optimization. More generally, we can consider the problem of finding a feasible point in the intersection of the cone $\Sigma[x]_{2\mathbf{d}}$ with any affine subspace. This allows us to solve sum of squares optimization problems involving univariate polynomials. Let $\mathcal{B} : \mathbb{R}[x]_{2\mathbf{d}} \rightarrow \mathbb{R}^m$ be a linear map. Given $b \in \mathbb{R}^m$, we want to find $p \in \Sigma[x]_{2\mathbf{d}}$ so that $\mathcal{B}(p) = b$. This is equivalent to minimizing the quadratic-penalized problem

$$(6.2) \quad f_{\mathcal{B}}(\mathbf{u}) = \|\mathcal{B}(\sigma(\mathbf{u})) - b\|^2.$$

COROLLARY 6.3. *Suppose there exists $p \in \Sigma[x]_{2\mathbf{d}}$ such that $\mathcal{B}(p) = b$. Then $\nabla f_{\mathcal{B}}(\mathbf{u}) = 0$ and $\nabla^2 f_{\mathcal{B}}(\mathbf{u}) \succeq 0$ implies that $f_{\mathcal{B}}(\mathbf{u}) = 0$.*

Proof. The gradient and Hessian of the objective (6.2) can be written as

$$\begin{aligned}
 \frac{1}{4} \nabla f_{\mathcal{B}}(\mathbf{u})(\mathbf{v}) &= \langle \mathcal{B}(\mathcal{A}_{\mathbf{u}}(\mathbf{v})), \mathcal{B}(\sigma(\mathbf{u}) - p) \rangle, \\
 \frac{1}{4} \nabla^2 f_{\mathcal{B}}(\mathbf{u})(\mathbf{v}, \mathbf{v}) &= \langle \mathcal{B}(\sigma(\mathbf{v})), \mathcal{B}(\sigma(\mathbf{u}) - p) \rangle + 2 \|\mathcal{B}(\mathcal{A}_{\mathbf{u}}(\mathbf{v}))\|^2.
 \end{aligned}$$

Thus by the linearity of \mathcal{B} we can use the same construction as in the certificate (5.3) to show that $f_{\mathcal{B}}(\mathbf{u}) = 0$. \square

7. Implementation and experiments. In this section we describe an efficient implementation of finding a sum of squares decomposition of trigonometric polynomials. A trigonometric polynomial of degree- d is defined by $2d + 1$ coefficients and has the form

$$p(t) = a_0 + \sum_{k=1}^d (a_k \cos(kt) + a_{-k} \sin(kt)).$$

By the substitution $\cos(t) = \frac{1-x^2}{1+x^2}$ and $\sin(t) = \frac{2x}{1+x^2}$, $p(x)$ becomes a rational function with the denominator a power of $1+x^2$ and the numerator a degree- $2d$ polynomial in x . Thus certifying the nonnegativity of the numerator is equivalent to certifying the nonnegativity of $p(t)$. By this correspondence the result of Theorem 1.1 also applies to trigonometric polynomials.

Since the proof of Theorem 1.1 does not depend on the norm used for $f_p(\mathbf{u}) = \|p - \sum_i u_i^2\|^2$, we can choose one most suitable for the gradient computation. For the rest of this section we assume that d is even for simplicity of notation; a similar decomposition exists for odd d by choosing “half-angles” (see [30] for more details). We then choose the inner product defined by evaluation at $2d+1$ points on the circle,

$$\langle p, q \rangle = \frac{1}{2d+1} \sum_{k=1}^{2d+1} p(x_k)q(x_k), \quad x_k = \frac{2k\pi}{2d+1}.$$

Since a trigonometric polynomial of degree- d is uniquely defined by evaluation on $2d+1$ unique points, $\|p(x)\|^2 = 0$ if and only if p is identically zero.

Let $U \in \mathbb{R}^{(d+1) \times r}$ be a matrix with column U_i representing the coefficients of $u_i(x)$, and $B \in \mathbb{R}^{(d+1) \times (2d+1)}$ be the evaluation map on $2d+1$ points with columns

$$B_k = [1 \quad \cos(x_k) \quad \cdots \quad \cos(\frac{d}{2}x_k) \quad \sin(x_k) \quad \cdots \quad \sin(\frac{d}{2}x_k)]^\top,$$

so that $B_k^\top U_i = u_i(x_k)$. Let \bar{p} be the vector of coefficients of $p(x)$, so that $B_k^\top \bar{p} = p(x_k)$. Then we can write

$$f_p(U) = \frac{1}{2d+1} \sum_{k=1}^{2d+1} \left(\|U^\top B_k\|^2 - p(x_k) \right)^2,$$

$$\nabla f_p(U) = \frac{4}{2d+1} U^\top B \text{Diag} \left(\|U^\top B_k\|^2 - p(x_k) \right) B^\top,$$

where $\text{Diag} \left(\|U^\top B_k\|^2 - p(x_k) \right)$ is a diagonal matrix with $\|U^\top B_k\|^2 - p(x_k)$ as the k th diagonal entry. Since matrix-vector multiplication by B is equivalent to a discrete Fourier transform, $\nabla f_p(U)$ can be computed in $O(rd \log d)$ time using the FFT. Theorem 1.1 shows that spurious local minima do not exist when $r \geq 2$, so we can pick r to be a constant and obtain a near-linear iteration complexity. This is in contrast to other SDP-based algorithms and custom interior point methods for solving this problem, which run into computational difficulties even when $2d = 10,000$.

We implemented our algorithm for finding the sum of squares decomposition of trigonometric polynomials in Julia,⁹ using the `FFTW.jl` [19] package for FFTs to compute $\nabla f_p(U)$ and the `NLOpt.jl` [25] package to minimize $f_p(U)$ using a first-order algorithm (L-BFGS). We performed the timing experiments on Intel Xeon Platinum 8260 processors, allocating at least $r+1$ cores to each run, using polynomials of degree- $2d$ ranging from 2,000 to 1,000,000. The test polynomials are generated with coefficients drawn from a standard normal distribution, with a constant coefficient added so that they all have a small positive minimum value. U is initialized with a small random value; its magnitude depends on the size of the problem. The algorithm is terminated when the relative error for each entry of U is on the order of 10^{-7} . Although $r = 2$ is sufficient, the results in Table 1(b) shows that $r = 4$ minimizes the total computational cost, as measured by the total number of FFT calls (by far

⁹The code and data required to reproduce the results in this section can be found in [28].




FIG. 2. Value of $f_p(\mathbf{u})$ against iterations of L -BFGS for computing the sum of squares decomposition of 20 random nonnegative trigonometric polynomials ($2d = 100,000$), showing a linear convergence rate.

the most expensive operation) needed for convergence. In addition, since the matrix-vector products can be easily parallelized across multiple threads, increasing r does not incur a significant per-iteration cost if sufficient threads are used. Thus we choose $r = 4$ for our large-scale experiments in Table 1a. Figure 2 plots the convergence rate of 20 instances, and we can see that they achieve a linear convergence rate. This is in contrast to grid-based methods [41, 17] which scale sublinearly in accuracy.

8. Conclusion. When does it make sense to solve nonconvex formulations of convex problems? In this paper we addressed this question for sum of squares decomposition and optimization of univariate polynomials, showing that solving the nonconvex formulation can provide a large computational speedup while still maintaining provable guarantees on the convergence to the global optima. Key to our approach is retaining polynomial structure in the nonconvex formulation. This enables us to use algebraic methods to construct a certificate showing that all SOCPs are global minima.

Our approach for finding sum of squares decompositions generalizes to multivariate polynomials, although we do not have guarantees for the rank needed to exclude spurious second-order critical points. On the other hand, results for low-rank matrix factorization tell us that this rank is equal to the Pythagoras number for quadratic forms. Thus we conjecture that a version of Theorem 1.1 is true for ternary quartics and matrix polynomials, cases where nonnegativity is equivalent to the existence of a sum of squares decomposition. Some similarities between our conditions and a new characterization of these cases in terms of varieties of minimal degree [6] also suggest that Theorem 1.1 could be generalized to these cases. In particular, the case where the syzygy module only contains the *Koszul syzygies* [16, p. 581] could generalize the coprime case studied in subsection 4.1.

Another direction for future work is to apply our methods to other structured semidefinite programs or polynomial-valued objectives such as symmetric tensor decomposition.

Acknowledgment. We thank the two anonymous referees for their valuable feedback and comments.

REFERENCES

- [1] A. S. BANDEIRA, N. BOUMAL, AND V. VORONINSKI, *On the low-rank approach for semidefinite programs arising in synchronization and community detection*, in Proceedings of the 29th Conference on Learning Theory, Proc. Mach. Learn. Res. 49, 2016, pp. 361–382.
- [2] A. BARVINOK, *A remark on the rank of positive semidefinite matrices subject to affine constraints*, Discrete Comput. Geom., 25 (2001), pp. 23–31, <https://doi.org/10.1007/s004540010074>.
- [3] S. BHOJANAPALLI, N. BOUMAL, P. JAIN, AND P. NETRAPALLI, *Smoothed analysis for low-rank solutions to semidefinite programs in quadratic penalty form*, in Proceedings of the 31st Conference On Learning Theory, Proc. Mach. Learn. Res. 75, S. Bubeck, V. Perchet, and P. Rigollet, eds., 2018, pp. 3243–3270.
- [4] S. BHOJANAPALLI, B. NEYSHABUR, AND N. SREBRO, *Global optimality of local search for low rank matrix recovery*, in Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16, 2016, pp. 3880–3888.
- [5] G. BLEKHERMAN, P. A. PARRILO, AND R. R. THOMAS, EDs., *Semidefinite Optimization and Convex Algebraic Geometry*, MOS-SIAM Ser. Optim. 13, SIAM, Philadelphia, 2013.
- [6] G. BLEKHERMAN, D. PLAUMANN, R. SINN, AND C. VINZANT, *Low-rank sum-of-squares representations on varieties of minimal degree*, Int. Math. Res. Not. IMRN, 2019, pp. 33–54.
- [7] N. BOUMAL, V. VORONINSKI, AND A. BANDEIRA, *The non-convex Burer-Monteiro approach works on smooth semidefinite programs*, in Proceedings of Advances in Neural Information Processing Systems, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, eds., 2016; also available online from <https://proceedings.neurips.cc/paper/2016/file/3de2334a314a7a72721f1f74a6cb4cee-Paper.pdf>.
- [8] N. BOUMAL, V. VORONINSKI, AND A. S. BANDEIRA, *Deterministic guarantees for Burer-Monteiro factorizations of smooth semidefinite programs*, Comm. Pure Appl. Math., 73 (2020), pp. 581–608.
- [9] S. BURER AND R. D. MONTEIRO, *A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization*, Math. Program., 95 (2003), pp. 329–357, <https://doi.org/10.1007/s10107-002-0352-8>.
- [10] S. BURER AND R. D. MONTEIRO, *Local minima and convergence in low-rank semidefinite programming*, Math. Program., 103 (2005), pp. 427–444, <https://doi.org/10.1007/s10107-004-0564-1>.
- [11] Y. CHI, Y. M. LU, AND Y. CHEN, *Nonconvex optimization meets low-rank matrix factorization: An Overview*, IEEE Trans. Signal Process., 67 (2019), pp. 5239–5269, <https://doi.org/10.1109/TSP.2019.2937282>.
- [12] M.-D. CHOI, T. Y. LAM, AND B. REZNICK, *Sums of squares of real polynomials*, in K-Theory and Algebraic Geometry: Connections with Quadratic Forms and Division Algebras, Part 2, Proc. Sympos. Pure Math. 58, American Mathematical Society, Providence, RI, 1995, pp. 103–126, <https://doi.org/10.1090/pspum/058.2/1327293>.
- [13] D. CIFUENTES, *On the Burer-Monteiro method for general semidefinite programs*, Optim. Lett., 15 (2021), pp. 2299–2309.
- [14] D. CIFUENTES AND A. MOITRA, *Polynomial Time Guarantees for the Burer-Monteiro Method*, <https://arxiv.org/abs/1912.01745>, 2019.
- [15] D. CIFUENTES AND P. A. PARRILO, *Sampling algebraic varieties for sum of squares programs*, SIAM J. Optim., 27 (2017), pp. 2381–2404, <https://doi.org/10.1137/15M1052548>.
- [16] D. A. COX, J. LITTLE, AND D. O’SHEA, *Ideals, Varieties, and Algorithms*, Springer, New York, 2015.
- [17] B. DUMITRESCU, *Positive Trigonometric Polynomials and Signal Processing Applications*, Signals Commun. Technol. 103, Springer, New York, 2007.
- [18] D. DUMMIT AND R. FOOTE, *Abstract Algebra*, Wiley, New York, 2003.
- [19] M. FRIGO AND S. G. JOHNSON, *The design and implementation of FFTW3*, Proc. IEEE, 93 (2005), pp. 216–231.
- [20] R. GE, F. HUANG, C. JIN, AND Y. YUAN, *Escaping from saddle points—Online stochastic gradient for tensor decomposition*, in Proceedings of the Conference on Learning Theory, 2015, pp. 797–842.
- [21] R. GE, C. JIN, AND Y. ZHENG, *No spurious local minima in nonconvex low rank problems: A unified geometric analysis*, in Proceedings of the 34th International Conference on Machine Learning, J. Mach. Learn. Res. 70, 2017, pp. 1233–1242.

- [22] R. GE, J. D. LEE, AND T. MA, *Matrix completion has no spurious local minimum*, in Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16, 2016, pp. 2981–2989.
- [23] R. W. HAMMING, *Numerical Methods for Scientists and Engineers*, McGraw-Hill, New York, 1973.
- [24] C. JIN, R. GE, P. NETRAPALLI, S. M. KAKADE, AND M. I. JORDAN, *How to escape saddle points efficiently*, in Proceedings of the International Conference on Machine Learning, Proc. Mach. Learn. Res. 70, 2017, pp. 1724–1732.
- [25] S. G. JOHNSON, *The NLopt Nonlinear-Optimization Package*, 2014.
- [26] L. KAPELEVICH, C. COEY, AND J. P. VIELMA, *Sum of squares generalizations for conic sets*, Math. Program., 199 (2023), pp. 1417–1429, <https://doi.org/10.1007/s10107-022-01831-6>.
- [27] M. LAURENT, *Sums of squares, moment matrices and optimization over polynomials*, in Emerging Applications of Algebraic Geometry, Springer, New York, 2009, pp. 157–270.
- [28] B. LEGAT, C. YUAN, AND P. A. PARRILO, *Low-rank univariate sum of squares has no spurious local minima*, 2023, <https://doi.org/10.24433/CO.7364738.v1>.
- [29] Q. LI, Z. ZHU, AND G. TANG, *The non-convex geometry of low-rank matrix optimization*, Inf. Inference, 8 (2019), pp. 51–96.
- [30] J. LOFBERG AND P. PARRILO, *From coefficients to samples: A new approach to SOS optimization*, in Proceedings of the 43rd IEEE Conference on Decision and Control (CDC), IEEE, 2004, pp. 3154–3159, <https://doi.org/10.1109/CDC.2004.1428957>.
- [31] D. PAPP, *Univariate polynomial optimization with sum-of-squares interpolants*, in Modeling and Optimization: Theory and Applications, Springer, New York, 2017, pp. 143–162.
- [32] P. A. PARRILO, *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*, Ph.D. thesis, California Institute of Technology, 2000.
- [33] P. A. PARRILO, *Polynomial games and sum of squares optimization*, in Proceedings of the 45th IEEE Conference on Decision and Control, IEEE, 2006, pp. 2855–2860.
- [34] G. PATAKI, *On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues*, Math. Oper. Res., 23 (1998), pp. 257–511 <https://doi.org/10.1287/moor.23.2.339>.
- [35] T. ROH AND L. VANDENBERGHE, *Discrete transforms, semidefinite programming, and sum-of-squares representations of nonnegative polynomials*, SIAM J. Optim., 16 (2006), pp. 939–964.
- [36] C. SCHEIDERER, *Extreme points of Gram spectrahedra of binary forms*, Discrete Comput. Geom., 67 (2022), pp. 1174–1190.
- [37] K. SCHMÜDGEN, *The K -moment problem for compact semi-algebraic sets*, Math. Ann., 289 (1991), pp. 203–206.
- [38] B. STURMFELS, *Solving Systems of Polynomial Equations*, CBMS Reg. Conf. Ser. Math. 97, American Mathematical Society, Providence, RI, 2002.
- [39] J. SUN, Q. QU, AND J. WRIGHT, *A geometric analysis of phase retrieval*, Found. Comput. Math., 18 (2018), pp. 1131–1198.
- [40] I. WALDSPURGER AND A. WATERS, *Rank optimality for the Burer-Monteiro factorization*, SIAM J. Optim., 30 (2020), pp. 2577–2602.
- [41] S.-P. WU, S. BOYD, AND L. VANDENBERGHE, *FIR filter design via spectral factorization and convex optimization*, in Applied and Computational Control, Signals, and Circuits, Springer, New York, 1999, pp. 215–245.