

A CONVEX RELAXATION TO COMPUTE THE NEAREST STRUCTURED RANK DEFICIENT MATRIX

DIEGO CIFUENTES

ABSTRACT. Given an affine space of matrices \mathcal{L} and a matrix $\hat{\theta} \in \mathcal{L}$, consider the problem of finding the closest rank deficient matrix to $\hat{\theta}$ on \mathcal{L} with respect to the Frobenius norm. This is a nonconvex problem with several applications in estimation problems. We introduce a novel semidefinite programming (SDP) relaxation, and we show that the SDP solves the problem exactly in the low noise regime, i.e., when $\hat{\theta}$ is close to be rank deficient. We evaluate the performance of the SDP relaxation in applications from control theory, computer algebra, and computer vision. Our relaxation reliably obtains the global minimizer in all cases for non-adversarial noise.

1. INTRODUCTION

Let k, m, n be positive integers with $m \geq k$, $n \geq m - k + 1$. Given an injective affine map $\mathcal{P} : \mathbb{R}^n \rightarrow \mathbb{R}^{k \times m}$ and a vector $\theta \in \mathbb{R}^n$, we consider the nearest point problem

$$(1) \quad \min_{u \in U} \|u - \theta\|^2, \quad \text{where } U = \{u \in \mathbb{R}^n : \mathcal{P}(u) \text{ is rank deficient}\},$$

and where we use the Euclidean norm in \mathbb{R}^n (or a scaled Euclidean norm). Equivalently, given the affine space $\mathcal{L} := \text{Im } \mathcal{P}$ and the matrix $\hat{\theta} = \mathcal{P}(\theta)$, we consider

$$(1') \quad \min_{\hat{u} \in \mathbb{R}^{k \times m}} \|\hat{u} - \hat{\theta}\|^2, \quad \text{such that } \hat{u} \in \mathcal{L} \quad \text{and} \quad \hat{u} \text{ is rank deficient},$$

where we use the (possibly scaled) Frobenius norm in $\mathbb{R}^{k \times m}$.

The above problem has several applications in control, signal processing, systems theory, computer algebra, and computer vision [4, 7, 13, 15, 16, 18]. Some common choices for the affine space \mathcal{L} include the spaces of Hankel, Toeplitz, and Sylvester matrices. Problem (1) is sometimes known as structured total least squares [7, 18, 23]. The name comes from the total least squares method for the linear regression problem $Ax \approx b$, given by minimizing $\|(\Delta A \mid \Delta b)\|^2$ subject to $(A + \Delta A \mid b + \Delta b)$ being rank deficient. By further imposing an affine constraint on the matrix, we arrive to (1').

Problem (1) is nonconvex and computationally hard. It is usually solved in practice with local optimization methods, that give no guarantees of convergence to the global optimum. A practical approach to certify global optimality in nonlinear programming is by constructing a convex relaxation. We follow this approach in this paper, and propose a novel *semidefinite programming* (SDP) relaxation for (1). The optimal value of the SDP relaxation is always a lower bound on the optimal value of (1). If the

Date: April 21, 2019.

Key words and phrases. Convex relaxation, Low rank approximation, Semidefinite programming, Total least squares.

optimal solution of the SDP is rank-one, then the relaxation is *exact* (or *tight*), and we can recover the global minimizer of (1) from the SDP.

The main technical contribution of this paper is to show that our SDP is always exact in the *low noise* regime, i.e., when θ is close enough to U . More precisely, we show that if $\bar{\theta} \in U$ satisfies a mild transversality condition, then the SDP relaxation is exact as $\theta \rightarrow \bar{\theta}$. We point out that the same transversality condition appears when analyzing the rate of convergence of local optimization methods [24].

Assumption 1. *The affine space \mathcal{L} meets the manifold $\{A \in \mathbb{R}^{k \times m} : \text{rank } A = k-1\}$ transversally at $\mathcal{P}(\bar{\theta})$. This means that $\text{codim}(\mathcal{L} \cap T) = \text{codim } \mathcal{L} + \text{codim } T$, where T denotes the tangent space of the manifold at $\mathcal{P}(\bar{\theta})$.*

Theorem 1.1. *Let $\bar{\theta} \in U$ such that Assumption 1 holds. Then the SDP relaxation (6) correctly recovers the minimizer of (1) whenever θ is close enough to $\bar{\theta}$.*

The structure of the paper is as follows. In Section 2 we introduce our SDP relaxation for (1). In Section 3 we prove our main result, Theorem 1.1. Finally, in Section 4 we evaluate the performance of our SDP relaxation in applications from control theory, computer algebra, and computer vision. We illustrate that SDP reliably solves all these problems, and its noise tolerance is notably better than state of the art methods.

Related work. Problem (1') is a special instance of structured low rank approximation (SLRA), which consists in minimizing $\|\hat{u} - \bar{\theta}\|$ subject to $\hat{u} \in \mathcal{L}$ and $\text{rank } \hat{u} \leq r$, for a given $r \in \mathbb{N}$. The SLRA problem has been studied extensively [4, 7, 20, 23, 24], notably by Markovsky [15, 16]. Most practical algorithms rely on local optimization, see e.g., [16, §3] and [24]. But due to nonconvexity, there are no guarantees to get the global minimum.

A heuristic method for SLRA consists in replacing the rank constraint by the convex constraint $\|\hat{u}\|_* \leq t$, where $\|\cdot\|_*$ is the nuclear norm and t is a tunable parameter. This leads to a tractable SDP, see [9]. However, the resulting matrix \hat{u} is not necessarily optimal, and may not even satisfy the rank constraint. A closely related problem to SLRA is that of finding the matrix $\hat{u} \in \mathcal{L}$ with the smallest rank. The nuclear norm heuristic has been widely studied in this context, see e.g., [22].

To the best of our knowledge, our proposed SDP is the first convex relaxation for problem (1) with provable guarantees. Special instances from computer vision have received more attention [1, 12], as we will elaborate in Section 4.3.

Since problem (1) is a polynomial optimization problem, the sum-of-squares (SOS) method [2, 14] provides a hierarchy of SDP relaxations with increasing power and complexity. Our proposed relaxation involves a positive semidefinite (PSD) matrix of size $(n+1)k$, and lies in between the first and the second level of the SOS hierarchy, whose sizes are $n+k+1$ and $\binom{n+k+2}{2}$. Though SOS relaxations often perform well in practice, there are very few theoretical results explaining this behavior. This paper contributes in this direction, by showing that our relaxation is always exact in the low noise regime.

2. SEMIDEFINITE PROGRAMMING RELAXATION

We proceed to derive the SDP relaxation of problem (1). Our strategy consists in first phrasing the problem as a *quadratically constrained quadratic program* (QCQP), and then consider the associated Lagrangian dual, which is always an SDP.

For convenience, we will center the problem at zero with the change variables $v = u - \theta$. Observe that (1) is equivalent to:

$$(2) \quad \min_{z \in \mathbb{R}^k, v \in \mathbb{R}^n} v^T v \quad \text{such that} \quad z^T \mathcal{P}_\theta(v) = 0, \quad z^T z = 1,$$

where $\mathcal{P}_\theta(v) := \mathcal{P}(v + \theta)$. The above is a QCQP, and hence its Lagrangian dual is an SDP relaxation of (2). Unfortunately, this simple relaxation does not reveal any information about the original problem, as shown next.

Lemma 2.1. *The optimal value of the Lagrangian dual of (2) is zero for any θ .*

Proof. Denoting H_i the Hessian of the i -th entry of $z^T \mathcal{P}_\theta(v)$, the Lagrangian dual is

$$\max_{\gamma \in \mathbb{R}, \lambda \in \mathbb{R}^m} \gamma \quad \text{such that} \quad \begin{pmatrix} 0 & 0 \\ 0 & \mathbf{1}_n \end{pmatrix} - \frac{1}{2} \sum_i \lambda_i H_i - \gamma \begin{pmatrix} \mathbf{1}_k & 0 \\ 0 & 0 \end{pmatrix} \succeq 0.$$

Since the diagonal entries of H_i are zero, then the first entry of the PSD matrix from above is $-\gamma$, and it must be nonnegative. It follows that $\gamma=0, \lambda=0$ is optimal. \square

We will obtain another QCQP reformulation of (1) which gives raise to a better SDP relaxation. Let $\tilde{v} := (1, v) \in \mathbb{R}^{n+1}$ be the vector obtained by prepending a one, and consider the Kronecker products

$$y := v \otimes z = \text{vec}(zv^T) \in \mathbb{R}^{nk}, \quad x := (z, y) = \tilde{v} \otimes z \in \mathbb{R}^N, \quad N := (n+1)k.$$

We will restate problem (2) in terms of the *rank-one* tensor x . Recall that this rank-one condition corresponds to the vanishing of certain minors. Indexing x by pairs (i, j) with $0 \leq i \leq n, 1 \leq j \leq k$, then

$$(3) \quad \begin{aligned} \text{rank}(x) = 1 & \iff x_{l_1} x_{l_2} = x_{l_3} x_{l_4} \text{ for all } l = (l_1, l_2, l_3, l_4) \in L, \\ L := \{l = ((i_1, j_1), (i_2, j_2), (i_1, j_2), (i_2, j_1)) : 0 \leq i_1 < i_2 \leq n, 1 \leq j_1 < j_2 \leq k\}. \end{aligned}$$

Let us express the constraints $z^T \mathcal{P}_\theta(v) = 0$ and $z^T z = 1$ in terms of x . We can write

$$\mathcal{P}_\theta(v) = (P_1 \tilde{v} \ P_2 \tilde{v} \ \cdots \ P_m \tilde{v}) \quad \text{for some} \quad P_i \in \mathbb{R}^{k \times (n+1)}.$$

Observe that P_i depends on the parameter θ , but we omit the subindex to simplify the notation. Let

$$p_i := \text{vec}(P_i) \in \mathbb{R}^N, \quad G := \begin{pmatrix} 0 & 0 \\ 0 & \mathbf{1}_{nk} \end{pmatrix} \in \mathbb{S}^N, \quad E := \begin{pmatrix} \mathbf{1}_k & 0 \\ 0 & 0 \end{pmatrix} \in \mathbb{S}^N,$$

where \mathbb{S}^N denotes the space of symmetric $N \times N$ matrices. Properties of the Kronecker product imply that

$$x^T G x = \|y\|^2 = \|v\|^2, \quad p_i^T x = p_i^T (\tilde{v} \otimes z) = z^T P_i \tilde{v}, \quad x^T E x = \|z\|^2.$$

It follows that problem (2) is equivalent to

$$(4) \quad \begin{aligned} \min_{x \in \mathbb{R}^N} \quad & x^T G x \\ & p_i^T x = 0 \quad \text{for } i = 1, \dots, m \\ & x_{l_1} x_{l_2} = x_{l_3} x_{l_4} \quad \text{for } l \in L \\ & x^T E x = 1 \end{aligned}$$

Example 2.2. Let $n = 1, k = m = 2$, and consider the map $\mathcal{P} : \mathbb{R} \rightarrow \mathbb{R}^{2 \times 2}$, $u \mapsto \begin{pmatrix} 1 & u \\ u & u \end{pmatrix}$. Then $N = 4$ and the rank-one tensor $x = (x_{01}, x_{02}, x_{11}, x_{12})$ satisfies $x_{01}x_{12} = x_{02}x_{11}$. Since $\mathcal{P}_\theta(v) = \begin{pmatrix} 1 & v+\theta \\ v+\theta & v+\theta \end{pmatrix}$ then $p_1 = \text{vec}\begin{pmatrix} 1 & 0 \\ \theta & 1 \end{pmatrix}$, $p_2 = \text{vec}\begin{pmatrix} \theta & 1 \\ \theta & 1 \end{pmatrix}$.

We will make two more transformations to (4). First, we replace the linear equations by some quadratics:

$$p_i^T x = 0 \quad \iff \quad (p_i^T x)x_j = x^T p_i e_j^T x = 0 \text{ for all } 1 \leq j \leq N$$

where $\{e_j\}_j \subseteq \mathbb{R}^N$ denotes the canonical basis. For the second transformation we need to introduce new notation.

Definition 2.3 (Block symmetrization). Given $A \in \mathbb{S}^N$, we let $\text{bSym}(A) \in \mathbb{S}^N$ be the matrix constructed as follows: divide A into $k \times k$ blocks and replace each of these blocks $B \in \mathbb{R}^{k \times k}$ by its symmetric part $\frac{1}{2}(B + B^T)$. For a non-symmetric matrix $A \in \mathbb{R}^{N \times N}$, we may define $\text{bSym}(A) := \frac{1}{2} \text{bSym}(A + A^T) \in \mathbb{S}^N$.

It can be checked that $x^T (p_i e_j^T) x = x^T \text{bSym}(p_i e_j^T) x$ when x is rank-one. We now present our final QCQP formulation of problem (1):

$$(5) \quad \begin{array}{l} \min_{x \in \mathbb{R}^N} \quad x^T G x \\ \quad x^T \text{bSym}(p_i e_j^T) x = 0 \quad \text{for } i = 1, \dots, m \text{ and } j = 1, \dots, N \\ \quad x_{l_1} x_{l_2} = x_{l_3} x_{l_4} \quad \text{for } l \in L \\ \quad x^T E x = 1 \end{array}$$

The corresponding primal/dual SDP relaxations are:

$$(6) \quad \begin{array}{l} \min_{Y \in \mathbb{S}^N} \quad G \bullet Y \\ \quad \text{bSym}(p_i e_j^T) \bullet Y = 0, \quad ij \in [m] \times [N] \\ \quad E_l \bullet Y = 0, \quad l \in L \\ \quad E \bullet Y = 1 \\ \quad Y \succeq 0 \end{array} \quad \begin{array}{l} \max_{\mu_{ij}, \nu_l, \gamma \in \mathbb{R}} \quad \gamma \\ \quad \mathcal{Q}(\mu, \nu, \gamma) \succeq 0 \\ \quad \mathcal{Q}(\mu, \nu, \gamma) := G + \sum_{ij \in [m] \times [N]} \mu_{ij} \text{bSym}(p_i e_j^T) + \sum_{l \in L} \nu_l E_l - \gamma E \end{array}$$

where \bullet denotes the trace inner product in \mathbb{S}^N , and where $E_l \in \mathbb{S}^N$ is the matrix such that $E_l \bullet Y = Y_{l_1 l_2} - Y_{l_3 l_4}$.

The above pair of SDP's is our proposed relaxation for (1). Both SDP's achieve the same value since Slater's condition is satisfied, and this value is always a lower bound for (1). The relaxation is *exact* if the minimizer Y^* of (6) has rank one. In such a case, the optimal values of (1) and (6) agree, and furthermore, the minimizer of (1) can be recovered from Y^* . In the next section we will show that the relaxation is always exact in the low noise regime.

Example 2.4. Retake the case from Example 2.2. There is a matrix $\text{bSym}(p_i e_j^T)$ for each $i \in \{1, 2\}$ and $j \in \{1, 2, 3, 4\}$. The first and the last are:

$$\begin{aligned} \text{bSym}(p_1 e_1^T) &= \text{bSym} \begin{pmatrix} 1 & 0 & 0 & 0 \\ \theta & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} = \frac{1}{2} \text{bSym} \begin{pmatrix} 2 & \theta & 0 & 1 \\ \theta & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} = \frac{1}{4} \begin{pmatrix} 4 & 2\theta & 0 & 1 \\ 2\theta & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}, \\ \text{bSym}(p_2 e_4^T) &= \text{bSym} \begin{pmatrix} 0 & 0 & 0 & \theta \\ 0 & 0 & 0 & \theta \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \frac{1}{2} \text{bSym} \begin{pmatrix} 0 & 0 & 0 & \theta \\ 0 & 0 & 0 & \theta \\ 0 & 0 & 0 & 1 \\ \theta & \theta & 1 & 2 \end{pmatrix}. \end{aligned}$$

Besides the eight matrices $\text{bSym}(p_i e_j^T)$, the SDP relaxation involves three more matrices:

$$G = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad E = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad E_l = \frac{1}{2} \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

There is a single E_l in this case, corresponding to the tuple $l = ((0, 1), (1, 2), (0, 2), (1, 1))$.

Remark (Scaled norms). The SDP relaxation (6) can be easily adapted to scaled Euclidean norms in \mathbb{R}^n . If the objective in (2) is replaced by $v^T C v$, where $C \succ 0$, then we simply need to set $G = \begin{pmatrix} 0 & 0 \\ 0 & C \otimes \mathbf{1}_k \end{pmatrix}$. Theorem 1.1 remains valid.

Remark (Removing constraints). The equations $E_l \bullet Y = 0$ might be removed from the SDP (6) while preserving the objective value. This might be convenient for computational purposes. To see why, let $l \in L$, and consider the group action $\sigma_l : \mathbb{Z}_2 \rightarrow \mathbb{S}^N$ of switching the $l_1 l_2$ and $l_3 l_4$ coordinates. Note that σ_l fixes the matrices $G, E, \text{bSym}(p_i e_j^T)$. The constraint $E_l \bullet Y = 0$ defines the invariant subspace of σ_l . This constraint can be dropped, as there is always an optimal solution in the invariant subspace [10, 21].

3. EXACTNESS UNDER LOW NOISE

In [5] we proposed a general framework to analyze the stability of SDP relaxations. In this section we will apply these ideas to the QCQP (5), and prove that the SDP relaxation (6) is exact under low noise (Theorem 1.1).

We now introduce the stability result we will use. Given a family of quadratic equations $\mathbf{h}_\theta = (h_\theta^1, h_\theta^2, \dots, h_\theta^\ell)$ parametrized by $\theta \in \Theta$, consider the QCQP:

$$(P_\theta) \quad \min_{z \in \mathbb{R}^k, y \in \mathbb{R}^{N-k}} \|y\|^2 \quad \text{such that} \quad h_\theta^i(z, y) = 0 \quad \text{for } i = 1, \dots, \ell.$$

Let $\bar{\theta} \in \Theta$ be a parameter such that the optimal value of $(P_{\bar{\theta}})$ is zero. The next theorem analyzes the Lagrangian relaxation of (P_θ) as $\theta \rightarrow \bar{\theta}$.

Theorem 3.1 (SDP stability in ED problems). *Let $\bar{x} = (\bar{z}, 0)$ be the minimizer of $(P_{\bar{\theta}})$. Assume that the following conditions hold:*

- (smoothness) *The set $\{(\theta, x) : \mathbf{h}_\theta(x) = 0\}$ is a smooth manifold nearby $(\bar{\theta}, \bar{x})$.*
- (Abadie constraint qualification) *The tangent space of $X_{\bar{\theta}} := \{x : \mathbf{h}_{\bar{\theta}}(x) = 0\}$ at \bar{x} is given by the kernel of $\nabla \mathbf{h}_{\bar{\theta}}(\bar{x})$.*
- (not a branch point) *The right kernel of $\nabla_z \mathbf{h}_{\bar{\theta}}(\bar{x})$ is trivial.*
- (restricted Slater) *There exists $\lambda \in \Lambda$ such that $\mathcal{A}(\lambda)|_{(\bar{z})^\perp} \succ 0$, where*

$$(7) \quad \Lambda := \{\lambda \in \mathbb{R}^\ell : \lambda^T \nabla \mathbf{h}_{\bar{\theta}}(\bar{x}) = 0\} \quad \text{and} \quad \mathcal{A}(\lambda) := \sum_{i=1}^{\ell} \lambda_i \nabla_{zz}^2 h_{\bar{\theta}}^i \in \mathbb{S}^k.$$

Then the Lagrangian relaxation of (P_θ) solves the problem exactly whenever θ is close enough to $\bar{\theta}$.

Proof. This is a special instance of [5, Thm. 5.1] in which the branch point condition is specialized to the case of nearest point problems (see [5, Ex. 5.12]). \square

The QCQP (5) is a special case of (P_θ) , where the constraints are

$$(8) \quad h_0(x) := x^T E x - 1, \quad h_{ij}(x) := x^T \text{bSym}(p_i e_j^T) x, \quad h_l(x) := x_{l_1} x_{l_2} - x_{l_3} x_{l_4},$$

and where $x = (z, y)$ with $z \in \mathbb{R}^k, y \in \mathbb{R}^{km}$. Recall that each p_i depends affinely on θ . The proof of Theorem 1.1 amounts to verifying that the conditions from Theorem 3.1 are satisfied. We proceed to analyze each of these conditions.

3.1. Smoothness. Let $X_\theta \subseteq \mathbb{R}^N$ be the zero set of the equations in (8). We need to show that $\{(\theta, x) : x \in X_\theta\} \subseteq \Theta \times \mathbb{R}^N$ is a manifold. Consider also the feasible set of (2):

$$W_\theta := \{(z, v) : z^T \mathcal{P}_\theta(v) = 0, z^T z = 1\} \subseteq \mathbb{R}^k \times \mathbb{R}^n.$$

It will be shown in Lemma 3.2 that $W_{\bar{\theta}}$ is a manifold nearby $(\bar{z}, 0)$. Since $\mathcal{P}_\theta(v) = \mathcal{P}(v + \theta)$ then W_θ is a translation of $W_{\bar{\theta}}$, and hence $\{(\theta, w) : w \in W_\theta\}$ is a manifold at $(\bar{\theta}, (\bar{z}, 0))$. Observe that the Segre embedding $(z, v) \mapsto z \otimes (1, v)$ gives an isomorphism $W_\theta \cong X_\theta$. It follows that $\{(\theta, x) : x \in X_\theta\}$ is also a manifold nearby $(\bar{\theta}, \bar{x})$.

3.2. Constraint qualification. From now on we will only focus on the nominal parameter $\bar{\theta}$, and hence we will ignore the dependence on θ . Let $X \subseteq \mathbb{R}^N$ and $W \subseteq \mathbb{R}^k \times \mathbb{R}^n$ be as above. We need to show that the Abadie constraint qualification (ACQ) holds for X . We will use the fact that ACQ agrees with the notion of *regularity* of schemes in algebraic geometry; see e.g., [8, §16.6]. In particular, it is preserved by isomorphisms. Since $W \cong X$, it suffices to show that ACQ holds for W .

Lemma 3.2. *Under Assumption 1, then ACQ holds for W at $(\bar{z}, 0)$. As a consequence, W is a smooth manifold nearby $(\bar{z}, 0)$.*

Proof. Let $V := \{(z, A) \in \mathbb{R}^m \times \mathbb{R}^{k \times m} : Az = 0, z^T z = 1\}$. It is known that V is regular everywhere (it is a resolution of $\{A : \text{rank } A \leq k-1\}$). Observe that W is the preimage of V under the affine injection $\phi : (z, v) \mapsto (z, \mathcal{P}_{\bar{\theta}}(v))$. Also note that V meets the image of ϕ transversally at $(\bar{z}, 0)$ by Assumption 1. It follows that W is regular at $(\bar{z}, 0)$. \square

3.3. Branch point. In order to show that \bar{x} is not a branch point, we will first provide an explicit formula for the Jacobian $\nabla \mathbf{h}(\bar{x})$. We need new notation. Consider the following bases of \mathbb{R}^{n+1} , \mathbb{R}^k , and \mathbb{R}^N :

$$\begin{aligned} \{f_{j_1}\}_{j_1=0}^n &\subseteq \mathbb{R}^{n+1}, \quad \{d_{j_2}\}_{j_2=1}^k \subseteq \mathbb{R}^k \text{ are the canonical bases,} \\ \{e_j\}_j &:= \{f_{j_1} \otimes d_{j_2}\}_{j=(j_1, j_2)} \subseteq \mathbb{R}^N = \mathbb{R}^{n+1} \otimes \mathbb{R}^k. \end{aligned}$$

Let $P_i \in \mathbb{R}^{k \times (n+1)}$ be obtained by reshaping $p_i \in \mathbb{R}^N$ into a matrix, and let $a_i, b_{i1}, \dots, b_{in}$ denote the $n+1$ columns of P_i . Equivalently,

$$(9) \quad a_i, b_{i1}, \dots, b_{in} \in \mathbb{R}^k \quad \text{are such that} \quad p_i = f_0 \otimes a_i + \sum_{s=1}^n f_s \otimes b_{is} \in \mathbb{R}^N.$$

The next lemma gives a formula for (some of) the rows of $\nabla \mathbf{h}(\bar{x})$.

Lemma 3.3. *Let $\bar{x} = (\bar{z}, 0)$, and consider the constraints from (8). Then*

$$\nabla h_0(\bar{x}) = 2f_0 \otimes \bar{z},$$

and for any i and $j = (j_1, j_2)$ we have

$$\nabla h_{ij}(\bar{x}) = \begin{cases} \frac{1}{2} \bar{z}_{j_2} (f_{j_1} \otimes a_i), & \text{if } j_1 \neq 0 \\ \bar{z}_{j_2} (f_0 \otimes a_i) + \sum_{s=1}^n f_s \otimes \text{sym}(b_{is} d_{j_2}^T) \bar{z}, & \text{if } j_1 = 0 \end{cases}$$

where sym stands for the symmetric part, i.e., $\text{sym}(S) = \frac{1}{2}(S + S^T)$.

Proof. Since $\bar{x} = f_0 \otimes \bar{z}$, then $\nabla h_0(\bar{x}) = 2E\bar{x} = 2f_0 \otimes \bar{z}$. To compute ∇h_{ij} we make use of the following identity:

$$\text{bSym}((f \otimes a)^T (g \otimes b)) (u \otimes z) = \text{sym}(fg^T)u \otimes \text{sym}(ab^T)z, \quad \text{for } a, b, z \in \mathbb{R}^k, f, g, u \in \mathbb{R}^{n+1}.$$

The above identity can be checked by straightforward manipulation. It follows that

$$\begin{aligned} 2\nabla h_{ij}(\bar{x}) &= 4 \text{sym}(f_0 f_{j_1}^T) f_0 \otimes \text{sym}(a_i d_{j_2}^T) \bar{z} + 4 \sum_{s=1}^n \text{sym}(f_s f_{j_1}^T) f_0 \otimes \text{sym}(b_{is} d_{j_2}^T) \bar{z} \\ &= \bar{z}_{j_2} (f_{j_1} + f_0 (f_{j_1}^T f_0)) \otimes a_i + 2 \sum_{s=1}^n (f_{j_1}^T f_0) f_s \otimes \text{sym}(b_{is} d_{j_2}^T) \bar{z}. \end{aligned}$$

Reducing the above expression we get the formula in the lemma. \square

In order to make the following proofs more explicit, we will introduce a new assumption, which is *weaker* than Assumption 1.

Assumption 2. *The vectors $\{a_i\}_i$ as in (9) span $(\bar{z})^\perp$.*

Remark (Assum1 \implies Assum2). If Assumption 1 holds, then $\mathcal{P}_{\bar{\theta}}(0) = (a_1 \ a_2 \ \dots \ a_m)$ has rank $k - 1$, and its left kernel is spanned by \bar{z} . So Assumption 2 also holds.

We now verify the branch point condition.

Lemma 3.4. *Under Assumption 2, the right kernel of $\nabla_z \mathbf{h}(\bar{x})$ is trivial.*

Proof. Let $\zeta \in \mathbb{R}^k$ be in the right kernel. Then $\nabla_z h_0(\bar{x})\zeta = 0$ and $\nabla_z h_{ij}(\bar{x})\zeta = 0$ for all i, j . By Lemma 3.3, we have $\bar{z}^T \zeta = 0$ and $\bar{z}_{j_2} (a_i^T \zeta) = 0$ for all i, j_2 . Since $\|\bar{z}\| = 1$, some \bar{z}_{j_2} is nonzero. Then ζ is orthogonal to each of \bar{z}, a_1, \dots, a_m , so it must be zero. \square

3.4. Restricted Slater. In order to show the restricted Slater condition, we will first establish an easier condition that implies it. Consider the linear space $\Lambda \subseteq \mathbb{R}^\ell$ and the linear map $\mathcal{A} : \Lambda \rightarrow \mathbb{S}^k$ from (7). Note that Λ is the set of Lagrange multipliers at \bar{x} . The map \mathcal{A} may not depend on all the multipliers λ_i . We partition the multipliers $\lambda = (\lambda^I, \lambda^{II})$ into two groups: $\lambda^I \in \mathbb{R}^{\ell_1}$ are the multipliers that appear in $\mathcal{A}(\lambda)$, and $\lambda^{II} \in \mathbb{R}^{\ell_2}$ are the ones that do not appear. We let $\alpha : \mathbb{R}^{\ell_1} \rightarrow \mathbb{S}^k$ be the restriction of \mathcal{A} to \mathbb{R}^{ℓ_1} . We can similarly partition the Jacobian $\mathbf{J} := \nabla \mathbf{h}(\bar{x})$ in the form $\mathbf{J} = \begin{pmatrix} \mathbf{J}_I \\ \mathbf{J}_{II} \end{pmatrix}$, with $\mathbf{J}_I \in \mathbb{R}^{\ell_1 \times N}$, $\mathbf{J}_{II} \in \mathbb{R}^{\ell_2 \times N}$. As before, sym stands for the symmetric part of a matrix.

Lemma 3.5. *Let $\alpha, \mathbf{J}_I, \mathbf{J}_{II}$ as above. Consider the linear spaces $K := \mathbf{J}_I(\ker \mathbf{J}_{II}) \subseteq \mathbb{R}^{\ell_1}$ and $V := \text{sym}(\bar{z} \otimes \mathbb{R}^k) \subseteq \mathbb{S}^k$, and assume that*

$$V^\perp \subseteq \text{Im } \alpha \quad \text{and} \quad \alpha(K) \subseteq V.$$

Then the restricted Slater condition holds.

Proof. Let $\pi : \Lambda \rightarrow \mathbb{R}^{\ell_1}$ be the projection $(\lambda^I, \lambda^{II}) \mapsto \lambda^I$. Note that $\mathcal{A} = \alpha \circ \pi$ by definition of α . Let $\mathcal{A}^*, \alpha^*, \pi^*$ be the adjoint operators of \mathcal{A}, α, π . We claim that

$$(10) \quad \ker \mathcal{A}^* \subseteq \ker \alpha^* \oplus \alpha(\ker \pi^*) \quad \text{and} \quad \ker \pi^* = K.$$

By taking the orthogonal complement of the first equation, we conclude that

$$\text{Im } \mathcal{A} \supseteq \text{Im } \alpha \cap (\alpha(K))^\perp \supseteq V^\perp = \{S \in \mathbb{S}^k : S\bar{z} = 0\}.$$

The above relation implies the restricted Slater condition. Hence, it suffices to show (10).

Since $\mathcal{A}^* = \pi^* \circ \alpha^*$, the first equation in (10) is a consequence of the following identity:

$$(11) \quad \ker(L_2 \circ L_1) = \ker L_1 \oplus L_1^*(\text{Im } L_1 \cap \ker L_2) \quad \text{for linear maps } L_1, L_2.$$

This identity follows by applying the rank-nullity theorem (domain $L = \ker L \oplus L^*(\text{Im } L)$) to the linear maps $L_1, L_2|_{\text{Im } L_1}$, and $L_2 \circ L_1$.

It remains to show that $\ker \pi^* = \mathbf{J}_I(\ker \mathbf{J}_{II})$. Let $i : \Lambda \rightarrow \mathbb{R}^\ell$ be the inclusion and $\rho : \mathbb{R}^\ell \rightarrow \mathbb{R}^{\ell_1}$ be the projection $\lambda \mapsto \lambda^I$, so that $\pi = \rho \circ i$. Notice that $\ker \rho^* = 0$ since ρ is surjective. Also note that $\text{Im } i = \ker \mathbf{J}^*$ since Λ is defined by the equation $\lambda^T \mathbf{J} = 0$, and hence $\ker i^* = \text{Im } \mathbf{J}$. As $\pi^* = i^* \circ \rho^*$, the identity (11) gives

$$\ker \pi^* = \ker \rho^* \oplus \rho(\text{Im } \rho^* \cap \ker i^*) = \rho(\text{Im } \rho^* \cap \text{Im } \mathbf{J}),$$

Observe that $\text{Im } \rho^* = \mathbb{R}^{\ell_1} \times \{0\}$. It follows that $w \in \ker \pi^*$ if and only if there is some x such that $w = \mathbf{J}_I x$ and $\mathbf{J}_{II} x = 0$, as wanted. \square

Let us now show that the conditions from Lemma 3.5 are satisfied in our situation. We first need a formula for the linear map $\mathcal{A}(\lambda)$, and we need to know which multipliers does it depend on. Table 1 partitions the multipliers $\lambda = (\lambda^I, \lambda^{II})$ into two groups. The map $\mathcal{A}(\lambda)$ only depends on the $\ell_1 = m \times k$ multipliers in the first group:

$$\mathcal{A}(\lambda) = \sum_{i=1}^m \sum_{j \in J_0} \frac{\mu_{ij}}{2} \text{sym}(a_i d_{j_2}^T), \quad \text{where} \quad J_0 := \{(0, j_2) : 1 \leq j_2 \leq k\}.$$

The above formula follows from Lemma 3.3, and from the observation¹ that $\gamma = 0$ for any multiplier $\lambda \in \Lambda$.

TABLE 1. Partition of the multipliers into two groups.

group	I	II	II	II
multiplier	μ_{ij}	μ_{ij}	ν_l	γ
equation	$x^T \text{bSym}(p_i e_j^T) x$	$x^T \text{bSym}(p_i e_j^T) x$	$x_{l_1} x_{l_2} - x_{l_3} x_{l_4}$	$x^T E x - 1$
indices	$i \in [m], j \in \{0\} \times [k]$	$i \in [m], j \in [n] \times [k]$	as in (3)	

¹ The linear space Λ is defined by $\{\mathcal{Q}(\mu, \nu, \gamma)\bar{x} = 0\}$, where \mathcal{Q} is as in (6). Expanding the equation $\bar{x}^T \mathcal{Q}(\mu, \nu, \gamma)\bar{x} = 0$ gives $\gamma = 0$.

A very important step toward verifying the conditions of Lemma 3.5 is to compute the linear space $K := \mathbf{J}_I(\ker \mathbf{J}_{II}) \subseteq \mathbb{R}^{m \times k}$. For the remaining of this section we let

$$\{f_{j_1}\}_{j_1=0}^n \subseteq \mathbb{R}^{n+1}, \quad \{d_{j_2}\}_{j_2=1}^k \subseteq \mathbb{R}^k, \quad \{g_i\}_{i=1}^m \subseteq \mathbb{R}^m \quad \text{the canonical bases.}$$

Lemma 3.6. *Under Assumption 2, then $\mathbf{J}_I(\ker \mathbf{J}_{II}) \subseteq \mathbb{R}^m \otimes \{\bar{z}\}$.*

Proof. Let $L_1 := \mathbb{R}^{n+1} \otimes \{\bar{z}\}$, $L_2 := \{f_0\} \otimes \mathbb{R}^k$. We will first show that

$$(12) \quad \ker \mathbf{J}_{II} \subseteq (f_0^\perp \otimes \{\bar{z}\}) \oplus (\{f_0\} \otimes \bar{z}^\perp) \subseteq L_1 \oplus L_2.$$

The second containment is clear. We proceed to prove the first one. Given $w \in \ker \mathbf{J}_{II}$, we can write it in the form

$$w = \ell_1 \otimes \bar{z} + f_0 \otimes \ell_2 + r f_0 \otimes \bar{z} + f \otimes \zeta, \quad \text{for some } r \in \mathbb{R}, \ell_1, f \in f_0^\perp, \ell_2, \zeta \in \bar{z}^\perp.$$

It suffices to show that $r = 0$ and $f \otimes \zeta = 0$. Since $w \in \ker \mathbf{J}_{II}$, then $\nabla h_0(\bar{x})w = 0$ and also $\nabla h_{ij}(\bar{x})w = 0$ for $j = (j_1, j_2)$ with $j_1 > 0$. Lemma 3.3 gives formulas for these gradients. The equation $\nabla h_0(\bar{x})w = 0$ says that $r = 0$. Assume by contradiction that $f \otimes \zeta \neq 0$. The equations $\nabla h_{ij}(\bar{x})w = 0$ imply

$$\bar{z}_{j_2}(f^T f_{j_1})(\zeta^T a_i) = 0, \quad \text{for all } i, j_1, j_2.$$

Since $\|\bar{z}\| = 1$ then some \bar{z}_{j_2} is nonzero. Also note that some $f^T f_{j_1} \neq 0$ since $f \in f_0^\perp \setminus \{0\}$. Similarly, some $\zeta^T a_i \neq 0$ since $\bar{z}^\perp = \text{span}\{a_i\}_i$. This is a contradiction. Hence, (12) holds.

It remains to show that $\mathbf{J}_I(L_1) \subseteq H$, $\mathbf{J}_I(L_2) \subseteq H$, where $H := \mathbb{R}^m \otimes \{\bar{z}\}$. The rows of \mathbf{J}_I are $\nabla h_{ij}(\bar{w})$, where $j = (j_1, j_2)$, $j_1 = 0$. An explicit formula is given in Lemma 3.3. Let $w = f' \otimes \bar{z} \in L_1$, and let us see that $\mathbf{J}_I w \in H$. The ij -th entry of $\mathbf{J}_I w$ is

$$\begin{aligned} (f' \otimes \bar{z})^T (\bar{z}_{j_2}(f_0 \otimes a_i) + \sum_{s=1}^n f_s \otimes \text{sym}(b_{is} d_{j_2}^T) \bar{z}) \\ = \sum_{s=1}^n (f'^T f_s) (\bar{z}^T \text{sym}(b_{is} d_{j_2}^T) \bar{z}) = \sum_{s=1}^n \bar{z}_{j_2} (f'^T f_s) (b_{is}^T \bar{z}) \end{aligned}$$

where we used that $a_i^T \bar{z} = 0$. Then,

$$\mathbf{J}_I w = \sum_{i, j_2, s} \bar{z}_{j_2} (f'^T f_s) (b_{is}^T \bar{z}) (g_i \otimes d_{j_2}) = \sum_{i, s} (f'^T f_s) (b_{is}^T \bar{z}) (g_i \otimes \bar{z}) \in H.$$

Consider now $w = f_0 \otimes \zeta \in L_2$. The ij -th entry of $\mathbf{J}_I w$ is

$$(f_0 \otimes \zeta)^T (\bar{z}_{j_2}(f_0 \otimes a_i) + \sum_{s=1}^n f_s \otimes \text{sym}(b_{is} d_{j_2}^T) \bar{z}) = \bar{z}_{j_2} (a_i^T \zeta)$$

and thus $\mathbf{J}_I w = \sum_i (a_i^T \zeta) g_i \otimes \bar{z} \in H$. □

We are ready to show the restricted Slater condition.

Lemma 3.7. *Under Assumption 2, then the conditions from Lemma 3.5 are satisfied. Hence, the restricted Slater condition holds.*

Proof. Recall that α is the restriction of \mathcal{A} to $\mathbb{R}^{m \times k}$. We may view α as follows:

$$\alpha : \mathbb{R}^m \otimes \mathbb{R}^k \rightarrow \text{sym}(\mathbb{R}^k \otimes \mathbb{R}^k), \quad g_i \otimes d \mapsto \frac{1}{2} \text{sym}(a_i \otimes d) \quad \text{for any } i \in [m], d \in \mathbb{R}^k.$$

We conclude that $\text{Im } \alpha = \text{sym}((\bar{z})^\perp \otimes \mathbb{R}^k)$, and hence the first condition holds. For the second one, Lemma 3.6 shows that $K \subseteq \mathbb{R}^m \otimes \{\bar{z}\}$, and the above formula of α implies that $\alpha(\mathbb{R}^m \otimes \{\bar{z}\}) \subseteq \text{sym}(\bar{z} \otimes (\bar{z})^\perp) \subseteq V$. \square

4. APPLICATIONS

Our SDP relaxation is implemented in CVX [11], and is available in <http://www.mit.edu/~diegcif/>. We proceed to illustrate its performance. In all experiments we use the SDP solver Mosek with the default parameters.

4.1. Hankel structure. Consider the problem of finding the nearest rank deficient Hankel matrix. This problem appears in several applications from systems theory and control. Some concrete applications are approximate realization, system identification, noisy deconvolution, and stochastic realization [15, 16].

The left of Table 2 illustrates the performance of our SDP relaxation for small values of the parameters k, m . The numbers in the table indicate the percentage of experiments for which SDP solved the problem exactly, i.e., the optimal solution is rank-one. For each k, m we used 2000 random instances sampled uniformly among those of norm one. Observe that the SDP was exact for all instances with $m \leq 6$.

For comparison, the right of Table 2 illustrates the performance of a local optimization method using the software `slra` [17]. As expected, the accuracy of local optimization is lower than that of SDP. Moreover, even when the local method is exact, it is hard to certify that it indeed converged to the global optimal. In particular, we can only verify that the local method succeeded for the instances for which SDP also succeeds.

TABLE 2. Nearest rank deficient Hankel matrix for random instances.

$k \setminus m$	SDP succeeds								BFGS and SDP succeed							
	3	4	5	6	7	8	9	10	3	4	5	6	7	8	9	10
3	100%	100%	100%	100%	100%	100%	99%	99%	94%	90%	84%	80%	78%	75%	71%	69%
4		100%	100%	100%	97%	93%	86%	79%		93%	86%	79%	71%	65%	57%	50%
5			100%	100%	99%	96%	88%	79%			92%	84%	76%	68%	62%	55%
6				100%	100%	98%	90%	78%				90%	80%	72%	63%	53%
7					100%	100%	98%	88%					88%	79%	68%	61%

The fraction of instances for which SDP is exact should diminish as $k, m \rightarrow \infty$. Nevertheless, SDP always behaves well in the low noise regime, as we showed in Theorem 1.1. The following example illustrates this behavior.

Example 4.1 (Approximate realization). Consider the discrete linear time invariant (LTI) system with transfer function $\frac{z-1}{z^2-1.6z+0.8}$. Let $y = (y_1, y_2, y_3, \dots)$ be the impulse response of the system, and let $H_m(y)$ be the $3 \times m$ Hankel matrix with the first entries of y . The rank of $H_m(y)$ is two (the order of the system). Assume now that the signal y is corrupted by a Gaussian noise, and hence $H_m(y)$ is full rank. The approximate

realization problem consists in finding the nearest \hat{y} such that $H_m(\hat{y})$ has rank two. Table 3 shows the performance of SDP for $m = 40$ and for different standard variations of the noise. The numbers are the percentage of experiments which are solved exactly over 400 repetitions. Observe that SDP solves all instances below a certain noise level, as predicted by Theorem 1.1. For comparison we also show the performance of the local optimization methods BFGS and LMA. As before, we can only verify that the local methods are exact for instances for which SDP is exact.

TABLE 3. Approximate realization problem.

noise	0.0	0.1	0.2	0.3	0.4	0.5
SDP	100%	100%	100%	98%	97%	94%
BFGS	100%	100%	91%	68%	53%	37%
LMA	100%	100%	90%	66%	49%	34%

Remark (Computational complexity). It is known that SDP relaxations are more expensive than local optimization. For instance, for the case $(k, m) = (3, 40)$ the SDP solver `Mosek` takes 11.4s, whereas `slra` only takes 1.4ms. Nonetheless, we point out that this higher complexity has been overcome in several other applications, by taking advantage of problem structure [6, 10, 21, 25], and by using heuristic methods such as that of Burer and Monteiro [3]. It is left for future work to investigate how to best take advantage of these techniques in solving the SDP (6).

4.2. Approximate GCD. Let $d, n_1, n_2 \in \mathbb{N}$ and let $\hat{f} \in \mathbb{R}[t]_{n_1}, \hat{g} \in \mathbb{R}[t]_{n_2}$ be univariate polynomials of degrees n_1, n_2 . The degree- d approximate GCD problem is:

$$\min_{f \in \mathbb{R}[t]_{n_1}, g \in \mathbb{R}[t]_{n_2}} \|f - \hat{f}\|^2 + \|g - \hat{g}\|^2, \quad \text{such that} \quad \deg(\gcd(f, g)) \geq d,$$

where $\|\cdot\|$ denotes the norm of the coefficient vector. The approximate GCD problem can be restated as in (1). Indeed, it is known that $\deg(\gcd(f, g)) \geq d$ if and only if $\text{Syl}_d(f, g)$ is rank deficient, where $\text{Syl}_d(f, g)$ is the $(n-2d) \times (n-d-1)$ Sylvester matrix, which is filled with the coefficients of f, g ; see e.g., [13].

Example 4.2. Consider the polynomials

$$f := (t^2 - 2)(t^4 + 2), \quad g := (t^2 - 2)(t^3 - 1), \quad \gcd(f, g) = t^2 - 2.$$

We normalize the coefficients of f, g and then corrupt them with Gaussian noise. Table 4 illustrates the performance of SDP in computing the degree-2 approximate GCD. The numbers indicate the percentage of successful experiments over 400 repetitions. Note that SDP solves all instances below a certain noise level. For comparison, we also show the performance of two local optimization methods.

TABLE 4. Approximate GCD problem.

noise	0.0	0.05	0.1	0.15	0.2	0.25
SDP	100%	100%	99%	93%	91%	92%
BFGS	100%	99%	91%	73%	62%	58%
LMA	100%	99%	91%	72%	63%	57%

4.3. Multiple view geometry. Given integers $n \geq k$ and vectors $\theta \in \mathbb{R}^k$, $a_i, b_i \in \mathbb{R}^k$, consider the *fractional programming* problem

$$(13) \quad \min_{z \in \mathbb{R}^k, u \in \mathbb{R}^n} \|u - \theta\|^2, \quad \text{such that} \quad u_i = \frac{a_i^T z}{b_i^T z} \quad \text{for } i = 1, \dots, n.$$

After clearing denominators, this takes the form of (1) with $m = n$. The above problem appears in several applications from computer vision.

Example 4.3 (Triangulation). Given ℓ projective cameras $P_j : \mathbb{P}^3 \dashrightarrow \mathbb{P}^2$ and noisy images $\theta_j \in \mathbb{R}^2$ of an unknown point $x \in \mathbb{P}^3$, the triangulation problem is

$$\min_{x \in \mathbb{R}^4, u_j \in \mathbb{R}^2} \sum_{j=1}^{\ell} \|u_j - \theta_j\|^2, \quad \text{such that} \quad u_j = \Pi P_j x \quad \text{for } j = 1, \dots, \ell,$$

where Π is the dehomogenization map $(y_1, y_2, y_3) \mapsto (y_1/y_3, y_2/y_3)$. This is a special case of (13) with $k = 4$ and $n = 2\ell$. Aholt, Agarwal and Thomas proposed an SDP relaxation for this problem in [1], and they showed that it is exact under low noise when the camera centers are not coplanar. The SDP from [1] is smaller than ours: its PSD matrix is $(n+1) \times (n+1)$, as opposed to $4(n+1) \times 4(n+1)$ for ours. But, as we will see, ours is more precise, succeeding in almost all instances we considered.

Table 5 compares both SDP relaxations, indicating the percentage of experiments that are solved exactly (the solution is rank-one) over 400 repetitions. The left of Table 5 considers the synthetic data set described in [1, 19]: the cameras are placed uniformly at random on the sphere of radius two pointing to the origin, and the points are generated uniformly at random inside the unit cube. For the right of Table 5 the cameras are placed uniformly on the line segment $(2, 0, 0) - (2, 0, 1)$, and the points are as before. In both cases the image size is approximately 2×2 units. The second configuration is very problematic for the SDP from [1], since the camera centers are coplanar. On the other hand, our SDP (6) behaves equally well.

TABLE 5. Camera triangulation problem.

		cameras on a sphere							cameras on a line					
noise		0.0	0.1	0.2	0.3	0.4	0.5	0.0	0.1	0.2	0.3	0.4	0.5	
$\ell = 3$	our SDP	100%	100%	100%	100%	100%	99%	100%	100%	100%	100%	100%	100%	
	AAT [1]	100%	93%	87%	81%	74%	68%	100%	0%	0%	0%	0%	0%	
$\ell = 7$	our SDP	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	99%	
	AAT [1]	100%	100%	99%	96%	88%	79%	100%	0%	0%	0%	0%	0%	

Example 4.4 (Resectioning). Given ℓ points $z_j \in \mathbb{P}^3$ and noisy images $\theta_j \in \mathbb{R}^2$ under an unknown projective camera $P : \mathbb{P}^3 \dashrightarrow \mathbb{P}^2$, the resectioning problem is

$$\min_{P \in \mathbb{R}^{3 \times 4}, u_j \in \mathbb{R}^2} \sum_{j=1}^{\ell} \|u_j - \theta_j\|^2, \quad \text{such that} \quad u_j = \Pi P z_j \quad \text{for } j = 1, \dots, \ell.$$

This is a special case of (13) with $k = 12$ and $n = 2\ell$. Table 6 illustrates the performance of our SDP (6) on the synthetic data set from [19]: the cameras are randomly placed on the sphere of radius two, and the points inside the unit cube. Our SDP is exact for most instances below a noise level of 0.1. For comparison, we refer to [19, Fig 2], where it is observed that their method cannot certify global optimality for any instance above a noise level of 0.03.

Kahl and Henrion proposed an SDP relaxation for this problem in [12], though without any guarantees. Table 6 also shows the performance of their SDP (the Schur formulation of order 3). We use the heuristic method explained in [12] of including a small multiple of the trace in the objective function ($\epsilon = 0.001$). This promotes a rank-one solution, at the expense of a loss in optimality. As shown in Table 6, the solution still had higher rank in most of the experiments performed, even with this heuristic.

TABLE 6. Uncalibrated camera resectioning.

		noise		0.0	0.05	0.1	0.15	0.2	0.25
$\ell = 6$	our SDP			100%	82%	80%	80%	83%	84%
	KH [12]			100%	26%	7%	4%	2%	2%
$\ell = 15$	our SDP			100%	100%	85%	59%	37%	25%
	KH [12]			100%	0%	0%	0%	0%	0%

Remark. The homography estimation problem is also a special instance of (13), so we may use our SDP relaxation. Another relaxation was proposed in [12], but with no guarantees. A related problem is the estimation of the essential matrix of two views, for which an SDP relaxation that is exact under low noise was recently proposed [26].

Acknowledgments. The author thanks Sameer Agarwal, Pablo Parrilo, Bernd Sturmfels, Rekha Thomas, and André Uschmajew for helpful discussions and comments. This work was partially done while the author was in the Max Planck Institute in Leipzig.

REFERENCES

- [1] C. Aholt, S. Agarwal, and R. Thomas. A QCQP approach to triangulation. In *ECCV (1)*, volume 7572 of *Lect. Notes Comput. Sci.*, pages 654–667. Springer, 2012.
- [2] G. Blekherman, P. Parrilo, and R. Thomas. *Semidefinite optimization and convex algebraic geometry*, volume 13. MOS-SIAM Series Optim., 2013.
- [3] S. Burer and R. D. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program.*, 95(2):329–357, 2003.
- [4] M. T. Chu, R. E. Funderlic, and R. J. Plemmons. Structured low rank approximation. *Linear Algebra Appl.*, 366:157–172, 2003.

- [5] D. Cifuentes, S. Agarwal, P. A. Parrilo, and R. R. Thomas. On the local stability of semidefinite relaxations. *arXiv:1710.04287*, 2017.
- [6] D. Cifuentes and P. A. Parrilo. Sampling algebraic varieties for sum of squares programs. *SIAM J. Optim.*, 27(4):2381–2404, 2017.
- [7] B. De Moor. Structured total least squares and L2 approximation problems. *Linear Algebra Appl.*, 188:163–205, 1993.
- [8] D. Eisenbud. *Commutative algebra: with a view toward algebraic geometry*, volume 150. Springer Science & Business Media, 2013.
- [9] M. Fazel. *Matrix rank minimization with applications*. PhD thesis, Stanford University, 2002.
- [10] K. Gatermann and P. A. Parrilo. Symmetry groups, semidefinite programs, and sums of squares. *J. Pure Appl. Algebra*, 192(1):95–128, 2004.
- [11] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, Mar. 2014.
- [12] F. Kahl and D. Henrion. Globally optimal estimates for geometric reconstruction problems. *Int. J. Comput. Vision*, 74(1):3–15, 2007.
- [13] E. Kaltofen, Z. Yang, and L. Zhi. Approximate greatest common divisors of several polynomials with linearly constrained coefficients and singular polynomials. In *Proc. Int. Symp. Symbolic Algebraic Comput. (ISSAC)*, pages 169–176, 2006.
- [14] J. B. Lasserre. *Moments, positive polynomials and their applications*, volume 1 of *Optim. Series*. Imperial College Press, 2009.
- [15] I. Markovsky. Structured low-rank approximation and its applications. *Automatica*, 44(4):891–909, 2008.
- [16] I. Markovsky. *Low rank approximation: algorithms, implementation, applications*. Springer Science & Business Media, 2011.
- [17] I. Markovsky and K. Usevich. Software for weighted structured low-rank approximation. *J. Comput. Appl. Math.*, 256:278–292, 2014.
- [18] I. Markovsky and S. Van Huffel. Overview of total least-squares methods. *Signal Process.*, 87(10):2283–2302, 2007.
- [19] C. Olsson, F. Kahl, and R. Hartley. Projective least-squares: Global solutions with local optimization. In *Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 1216–1223, 2009.
- [20] G. Ottaviani, P.-J. Spaenlehauer, and B. Sturmfels. Exact solutions in structured low-rank approximation. *SIAM J. Matrix Anal. Appl.*, 35(4):1521–1542, 2014.
- [21] F. Permenter. *Reduction methods in semidefinite and conic optimization*. PhD thesis, Massachusetts Institute of Technology, 2017.
- [22] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501, 2010.
- [23] J. B. Rosen, H. Park, and J. Glick. Total least norm formulation and solution for structured problems. *SIAM J. Matrix Anal. Appl.*, 17(1):110–126, 1996.
- [24] É. Schost and P.-J. Spaenlehauer. A quadratically convergent algorithm for structured low-rank approximation. *Found. Comput. Math.*, 16(2):457–492, 2016.
- [25] L. Vandenberghe and M. S. Andersen. Chordal graphs and semidefinite optimization. *Found. Trends Optim.*, 1(4):241–433, 2014.
- [26] J. Zhao. An efficient solution to non-minimal case essential matrix estimation. *arXiv:1903.09067*, 2019.

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE MA 02139, USA

E-mail address: diegcif@mit.edu