

Convex Relaxations for Consensus and Non-Minimal Problems in 3D Vision

Thomas Probst, Danda Pani Paudel, Ajad Chhatkuli, Luc Van Gool

Computer Vision Laboratory, ETH Zurich, Switzerland

Abstract

In this paper, we formulate a generic non-minimal solver using the existing tools of Polynomials Optimization Problems (POP) from computational algebraic geometry. The proposed method exploits the well known Shor's or Lasserre's relaxations, whose theoretical aspects are also discussed. Notably, we further exploit the POP formulation of non-minimal solver also for the generic consensus maximization problems in 3D vision. Our framework is simple and straightforward to implement, which is also supported by three diverse applications in 3D vision, namely rigid body transformation estimation, Non-Rigid Structure-from-Motion (NRSfM), and camera autocalibration. In all three cases, both non-minimal and consensus maximization are tested, which are also compared against the state-of-the-art methods. Our results are competitive to the compared methods, and are also coherent with our theoretical analysis.

The main contribution of this paper is the claim that a good approximate solution for many polynomial problems involved in 3D vision can be obtained using the existing theory of numerical computational algebra. This claim leads us to reason about why many relaxed methods in 3D vision behave so well? And also allows us to offer a generic relaxed solver in a rather straightforward way. We further show that the convex relaxation of these polynomials can easily be used for maximizing consensus in a deterministic manner. We support our claim using several experiments for aforementioned three diverse problems in 3D vision.

1. Introduction

Robust model estimation in 3D vision usually involves two sub-problems: inlier detection and model computation. Faithful methods, already existing in many cases, have offered an unparalleled success of 3D vision applications from 3D scene reconstruction to 6DoF camera localization. In general, the task of inlier detection is performed by maximizing the consensus among measurements. Measurements that agree with some parameter of the known model are

treated as inliers. Customarily, model parameters are estimated from randomly selected minimal number of measurements – using the so-called minimal solver methods. These parameters then seek for the maximum consensus, within the framework of Random Sampling and Consensus (RANSAC). Parameters that maximize the consensus are further refined with respect to all the inlier measurements – using the so-called non-minimal solver methods.

Global methods for consensus maximization have gathered significant attention [27, 5, 16, 10, 37, 46, 33], mainly because RANSAC is non-deterministic and provides no guarantee for optimality. Although the consensus maximization problem has proven to be NP-hard [14], global methods with convex model representations have shown promising results both in terms of speed and optimality [42, 16]. Therefore, existing methods offer satisfactory solutions only for convex models. In the case of non-convex models, some methods introduce linearized auxiliary model parameters with additional convex constraints [43, 42]. There do exist methods for non-convex models which explore the exact parameters [46, 37, 33]. However, all methods designed for non-convex problems are specific to the problems at hand. These methods therefore do not generalize for other non-convex models. Consequently, a generic framework for deterministic consensus maximization of non-convex models, beyond simple linearization, is highly desirable.

Unlike consensus maximization, non-minimal solvers have gathered comparatively little attention. In practice, the need of optimal non-minimal solvers is usually overlooked. Instead, a local refinement on some geometric cost derived for all inliers is performed as a surrogate method, starting from the model parameters obtained by maximizing consensus. This however entails the risk of falling into local minima traps of the global cost function resulting in inadequate solutions, which has been demonstrated for various problems [8, 32, 22, 9, 39, 18, 23]. More specifically, non-minimal solvers have been devised separately for 3D rotations [18], 3D rigid body transformation in [8, 32, 39, 23], perspective-n-point in [22], and two-view relative pose in [9]. On the one hand, these solvers are very

specific to their own addressed problems, and only [32] provides the theoretical optimality guarantee owing to the used Branch-and-Bound (BnB) search paradigm. On the other hand, methods in [22, 8, 9] are fast and also achieve a posteriori certificate, with an open question for theoretical proof, of the global optimality. One notable work [21], provides convex relaxations for multiple computer vision problems while primarily focusing on the task of geometric error minimization, where the addressed problems are often linear with respect to the algebraic error minimization.

In fact, minimization of cost functions defined by polynomials, the global cost function in the non-minimal case, is a long standing problem in algebraic geometry. Most notably, one can use the Sum-of-Squares (SoS) methods to minimize such polynomials using Semi-Definite Programming (SDP). The theoretical optimality in this process can be guaranteed, when a hierarchy of SoS polynomials as constraints [36] is enforced. In this context, many methods in 3D vision make use of the SoS hierarchy; among which are non-rigid 3D reconstruction [34, 4], camera auto-calibration [11, 33], 3D-3D registration [33], and 3D modeling [1], to name a few. The SoS hierarchy is also known as the Lasserre’s hierarchy of relaxations in SDP [25]. The first term of Lasserre’s relaxations is indeed the so-called Shor’s relaxation, a well known result from the 80’s [40], which is also known to be tight for quadratic polynomial optimization [6]. Therefore, it is not very surprising that the non-minimal solvers of [8, 9], which use Shor’s relaxation with additional problem specific constraints, are fast and accurate. Moreover, in the context of non-minimal solvers, where one is offered sufficiently many polynomials, the tightness obtained using only the Shor’s relaxation may not be an issue. Motivated by these observations, we are interested to study the behaviour of SoS hierarchy for both consensus maximization and non-minimal problems.

In this paper, we formulate a generic non-minimal solver using the existing tools of Polynomials Optimization Problems (POP) from computational algebraic geometry. We also provide theoretical insights for using Shor’s or Lasserre’s relaxation, and support them by synthetic and real data experiments. More precisely, we suggest that Shor’s relaxation can very often offer satisfactory solutions when the models can be expressed as quadratic polynomials. Furthermore, we show that higher order of Lasserre’s relaxation is required for non-minimal solvers for higher degree polynomials in the case of challenging non-rigid 3D reconstruction. More interestingly, we successfully apply the same formulation used for non-minimal solvers also in the consensus maximization context. It then becomes obvious to make these relaxations tighter by adding more terms of Lasserre’s relaxation, as one desires. However, we argue that one may not always need to use higher order relaxation when seeking for consensus among polynomials.

Our framework is simple and straightforward to implement, which is supported by three diverse applications in 3D vision, namely rigid body transformation estimation, non-rigid reconstruction, and camera autocalibration. In all three cases, both non-minimal and consensus maximization are tested. We further show that the suggested relaxations work satisfactorily even for the minimal problem setup. All of our experiments are compared against the leading state-of-the-art methods. Our results are competitive to the compared methods, both in speed and accuracy, which are also coherent with our theoretical analysis.

2. Background and Notations

We denote matrices with upper case letters and their elements by double-indexed lower case letters: $A = (a_{ij})$. Similarly, vectors are indexed: $\mathbf{a} = (a_i)$. We write $A > 0$ (resp. $A \geq 0$) to denote that the symmetric matrix A is positive definite (resp. positive semi-definite). Let $\mathbb{R}[x]$ be the ring of polynomials parametrized by variables $x \in \mathbb{R}^n$, and $\mathbf{z}_d(x)$ be a vector of monomials, on the entries of x , ascending in degree up to d . Any polynomial $p(x) \in \mathbb{R}[x]$ can be represented using $\mathbf{z}_d(x)$ and the Gram matrix.

Definition 2.1 (Gram matrix [38]) For a degree $\leq 2d$ polynomial $p(x) \in \mathbb{R}[x]$, the symmetric matrix G such that $p(x) = \mathbf{z}_d(x)^T G \mathbf{z}_d(x)$ is a Gram matrix of $p(x)$.

One is often interested on solving the following general non-convex polynomial optimization problem.

Problem 2.2 (Polynomial optimization problem [24])

For a set of general non-convex polynomials $p_i(x), i = 0, \dots, m$, the Polynomial Optimization Problem (POP) is given by,

$$\min_x \{p_0(x) \mid p_i(x) \geq 0, i = 1, \dots, m\}. \quad (1)$$

In the following, we first focus on the cases when $p(x)$ is quadratic before diving into higher degree polynomials. Notice that for quadratic polynomials, $\mathbf{z}_1(x)$ is the homogeneous representation of x . In such case, a relaxed convex solution can be obtained using the Shor’s method.

Definition 2.3 (Shor’s relaxation [41]) For general non-convex quadratic polynomials $p_i(x), i = 0, \dots, m$, the POP of (1) is equivalent to the following problem.

$$\min_{Y \geq 0} \{tr(G_0 Y) \mid rk(Y) = 1, tr(G_i Y) \geq 0, i = 1, \dots, m\}, \quad (2)$$

for rank and trace operators $rk(\cdot)$ and $tr(\cdot)$, respectively. A convex relaxation of (2) can be obtained by dropping the rank-1 constraint on $Y := \mathbf{z}_1(x)\mathbf{z}_1(x)^T$.

It is well known that the Shor's relaxation is tight for the setup of (2) [6, 26]. Therefore, the solution of (1) is very well approximated by such relaxation. Although, we develop most of the theory using the quadratic polynomials with Shor's relaxation, we also address the case of higher degree polynomials with tighter relaxations. In the latter, we make use of a relaxation based on Lasserre's hierarchy [25], similar to the Shor's relaxation for quadratic case.

Now, we briefly present the theory behind Lasserre's relaxation. In particular, we discuss the case when the highest order of relaxation is $2n$. For a vector of relaxed variables w , the truncated moment matrix of order $2n$ is $M^{2n}(w) = W = W^T$ such that $w_{\alpha\beta} = w_{\alpha+\beta}$ where $\alpha, \beta \in \mathbb{N}^n$. By construction, any lower order moment matrix is a submatrix of W where $M^n(w) = Y := z_d(x)z_d(x)^T$ and $M^0(w) = 1$. Lasserre's method derives a hierarchy of constraints using the so-called Localizing matrices.

Definition 2.4 (Localizing matrix [25]) *The localizing matrix for polynomial $p(x)$ and relaxation order $s \leq n$ is a matrix $M^s(p(x)w)$ given by,*

$$M^s(p(x)w) = \mathcal{L}(\text{tr}(GY)M^s(w)), \quad (3)$$

where the Riesz function $\mathcal{L}(\cdot)$ maps the bilinear terms on Y and W to the corresponding terms of W . Note that $M^s(w)$ is a submatrix of Y for $s \leq n$. Therefore, $M^s(p(x)w)$ can be expressed linearly on W in this case.

Definition 2.5 (Lasserre's relaxation [25]) *An efficient relaxed solution of Problem 2.2 of general non-convex polynomials can be obtained by solving a Semi-Definite Program (SDP) of a hierarchy of relaxations. For the relaxation order of $2n$ and $i = 1, \dots, m$, it is given by,*

$$\min_{W \geq 0} \left\{ \text{tr}(G_0 M^n(w)) \mid \text{tr}(M^s(p_i(x)w)) \geq 0 \right\}. \quad (4)$$

The Lasserre's relaxation is known to be tighter than Shor's relaxation, which also provides the certificate of finite convergence when $s \geq n$ [25]. In fact, the Lasserre's relaxation includes the Shor's relaxation as a special case when $s = 0$ because $M^0(p_i(x)w) = \text{tr}(G_i Y)$ and $M^n(w) = Y$. Needless to say that the Lasserre's relaxation is also applicable to quadratic polynomials at the cost of a higher computation. However, for the non-minimal problems of quadratic polynomials, our experiments show that the Lasserre's relaxation is not really necessary.

3. Non-minimal Problem of Polynomials

In this section, we formulate the non-minimal problem of polynomials. We first propose a solution to this problem for general quadratic polynomials using the Shor's relaxation method. Later, the proposed solution will be generalized for the case of higher degree polynomials using the Lasserre's method of relaxation.

Problem 3.1 *For $x \in \mathbb{R}^n$ and $m \geq n$, the non-minimal problem of polynomials $p_i(x) \in \mathbb{R}[x], i = 1, \dots, m$, is,*

$$\min_{x, \epsilon_i} \left\{ \sum_i \epsilon_i \mid \epsilon_i \geq |p_i(x)|, i = 1, \dots, m \right\}. \quad (5)$$

The problem of 3.1 is in fact the \mathcal{L}_1 minimization problem over polynomials. One is often interested to minimize such objective because in many 3D vision problems ϵ is a geometric measure (such as point-to-line distance for two-view epipolar constraint using Essential matrix). Depending upon application, one may be interested to minimize \mathcal{L}_2 (or \mathcal{L}_∞ for that matter). This nonetheless, is not really a problem. For the sake of simplicity, we first present the theory using the \mathcal{L}_1 formulation. Its extension to \mathcal{L}_p - norm is discussed in Section 6. Non-minimal problems are the outcome of over-determined systems. The solution we are seeking is the one that agrees with all polynomials with minimum cumulative error, unlike the exact solution of the minimal case. Note when $n = m$, the non-minimal problem becomes a minimal problem.

3.1. Quadratic Polynomials

Proposition 3.2 *For a set of non-convex quadratic polynomials $\{p_i(x)\}_{i=1}^m$, the Shor's relaxation provides a convex relaxation of the non-minimal Problem 3.1 as,*

$$\min_{Y \geq 0, \epsilon_i} \left\{ \sum_i \epsilon_i \mid \epsilon_i \geq |\text{tr}(G_i Y)|, i = 1, \dots, m \right\}. \quad (6)$$

Proof Here, we provide the intuition behind the proof. Note that the problem of (6) is equivalent to (5) for quadratic polynomials $\{p_i(x)\}_i^m$ when $\text{rk}(Y) = 1$. Therefore (6) can directly be obtained by dropping the rank constraint, similar as in the Shor's relaxation 2.3. In fact, a rigorous proof can be obtained using the POP formulation of Section 2. Please, refer to the supplementary for the alternative proof. ■

The relaxed convex problem (6) is an SDP. This can be solved efficiently using the interior point method [7]. Ideally, the rank of Y is expected to be of rank-1. However, this is not usually the case. Therefore, we recover the primal solution $x \in \mathbb{R}^n$ after enforcing the rank-1 constraint using Singular Value Decomposition (SVD). In fact, the principal singular-vector Y is the homogeneous representation of $x \in \mathbb{R}^{n+1}$. Recall that the dual relaxed variable Y encodes the primal solution in the form $Y := z_1(x)z_1(x)^T$.

3.2. General Polynomials

Proposition 3.3 *For a set of non-convex general polynomials $\{p_i(x)\}_{i=1}^m$ of degree $\leq d$, the Lasserre's relaxation with order $s \leq n \in \mathbb{N}$ provides a convex relaxation of the non-minimal Problem 3.1 as,*

$$\min_{W \geq 0, \epsilon_i} \left\{ \sum_i \epsilon_i \mid \epsilon_i \geq |\text{tr}(M^s(p_i(x)w))|, i = 1, \dots, m \right\}. \quad (7)$$

Proof The proof is similar to the Proposition 3.2. Please, refer the supplementary material for the details. ■

The optimal primal solution $x \in \mathbb{R}^n$ can be recovered from W using SVD, similar to that of (6) discussed above.

4. Consensus Maximization Problem

The non-minimal method presented in the previous section assumes that the polynomials may be corrupted by the noise. Therefore, a solution that minimizes the cumulative error of all polynomials is searched. In the presence of noise and outliers, we wish to solve the following Problem.

Problem 4.1 Given a set $\mathcal{S} = \{p_i(x)\}_{i=1}^m$ and a threshold ϵ ,

$$\max_{x, \zeta \in \mathcal{S}} \left\{ |\zeta|, \epsilon \geq |p_i(x)|, \forall p_i(x) \in \zeta \right\}. \quad (8)$$

The consensus maximization problem seeks for a feasible solution that results the largest inlier set – a set of polynomials with residual smaller than ϵ . This problem, however, is difficult to solve and known to be NP-hard [14], even when $p(x)$ is a linear function on x . In this work, we approach this problem using the Branch-and-Bound search paradigm.

4.1. Branch-and-Bound Method

Our Branch-and-Bound(BnB) search is performed by branching on the space of binary assignment variables, one for each member in \mathcal{S} . During the BnB process, we seek for a feasible x using the mixed-integer programming method.

Definition 4.2 (Mixed-Integer Programming, MIP)

For a set of binary variables $\mathcal{Z} \in \{0, 1\}^m$ representing the inlier/outlier assignments, a given set of polynomials $\mathcal{S} = \{p_i(x)\}_{i=1}^m$ and a threshold ϵ , the consensus maximization Problem 4.1 is equivalent to,

$$\min_{x, z_i \in \mathcal{Z}} \left\{ \sum_{i=1}^m z_i |z_i \mathbf{M} + \epsilon \geq |p_i(x)|, \mathcal{Z} \in \{0, 1\}^m \right\}, \quad (9)$$

where \mathbf{M} is a sufficiently large positive scalar constant, commonly used in optimization to ignore invalid constraints [17]. If the binary variable $z_i = 0$, the polynomial $p_i(x)$ is an inlier. Otherwise, it is an outlier.

4.2. Consensus using MI-SDP

We formulate the consensus maximization problem by using the polynomial relaxations within the framework of Mixed-Integer Semi-Definite Programming (MI-SDP), for two different cases. First, we use Shor’s relaxation of (2) with MI-SDP for solving the consensus maximization problem of quadratic polynomials. Later, the Lasserre’s relaxation of (4) is used within the same framework, for the consensus maximization of more general polynomials.

Corollary 4.3 (Quadratic case) Given a set of non-convex quadratic polynomials $\mathcal{S} = \{p_i(x)\}_{i=1}^m$ and a threshold ϵ , the consensus maximization Problem 4.1 can be solved using the following Mixed integer semi-definite program,

$$\min_{Y \geq 0, z_i \in \mathcal{Z}} \left\{ \sum_{i=1}^m z_i |z_i \mathbf{M} + \epsilon \geq |\text{tr}(G_i Y)|, \mathcal{Z} \in \{0, 1\}^m \right\}. \quad (10)$$

The MI-SDP of (10) can be solved efficiently using off-the-self optimization toolboxes. The MI-SDP solution provides us the optimal set of inlier polynomials, together with a feasible x . These inlier polynomials are then used to solve the non-minimal problem of (6), to obtain an optimal solution x . A similar MI-SDP formulation for the consensus maximization among general polynomials is given below.

Corollary 4.4 (General case) Given a set of non-convex general polynomials $\mathcal{S} = \{p_i(x)\}_{i=1}^m$ and a threshold ϵ , the consensus maximization Problem 4.1 for assignments $\mathcal{Z} \in \{0, 1\}^m$ can be solved using the following MI-SDP,

$$\min_{W \geq 0, z_i \in \mathcal{Z}} \left\{ \sum_{i=1}^m z_i |z_i \mathbf{M} + \epsilon \geq |\text{tr}(M^s(p_i(x)w))|\right\}. \quad (11)$$

5. 3D Vision Problems

We present three examples of 3D vision problems for both consensus maximization and non-minimal problems. We start by a generic formulation of consensus maximization for all three problems. The inlier set of consensus maximization are then used to solve a non-minimal problem for the optimal set of parameters. During both of these stages, we also introduce problem specific polynomial constraints.

Let us consider a set of vectorized polynomials $\{P_i(x)\}_{i=1}^m$, possibly from many outlier measurements, is given to us. In the first step, we are interested in solving the following consensus maximization problem,

$$\min_{x, z_i \in \mathcal{Z}} \left\{ \sum_{i=1}^m z_i |z_i \mathbf{M} + \epsilon \geq |P_i(x)|, x \in \mathcal{K}, \mathcal{Z} \in \{0, 1\}^m \right\}, \quad (12)$$

where $x \in \mathcal{K}$ represents problem specific constraints and \mathcal{Z} measures inlier/outlier assignments, and " \geq " represents one-to-many elementwise inequality. We solve (12) using our Corollary 4.3/4.4 for x and \mathcal{Z} . Once the inlier/outlier assignments, i.e. \mathcal{Z} , is obtained, we solve the following non-minimal problem using our Proposition 3.2/3.3.

$$\min_{x, \epsilon_i} \left\{ \sum_{i=1}^m \epsilon_i |z_i \mathbf{M} + \epsilon_i \geq |P_i(x)|, x \in \mathcal{K} \right\}. \quad (13)$$

5.1. Rigid Body Transformation

We consider correspondences between two point clouds that differ by a 3D rigid body transformation. Let $\{u, v\}$ be euclidean coordinates of a pair of corresponding points such

that $\mathbf{v} = \mathbf{R}\mathbf{u} + \mathbf{t}$, for rotation matrix $\mathbf{R} \in SO(3)$ and translation $\mathbf{t} \in \mathbb{R}^3$. For quaternions $\mathbf{q} \in \mathbb{R}^4$, we represent the rotation matrix as $\mathbf{R}(\mathbf{q})$ with entries quadratic in \mathbf{q} . Given a set of correspondences $\{\mathbf{u}_i, \mathbf{v}_i\}_{i=1}^m$, we solve consensus maximization and non-minimal problems, respectively of (12) and (13), for variable $\mathbf{x} = (\mathbf{q}^\top, \mathbf{t}^\top)^\top$ and polynomials $P_i(\mathbf{x}) = \mathbf{v}_i - \mathbf{R}(\mathbf{q})\mathbf{u}_i - \mathbf{t}$, with $\mathcal{K} = \{\mathbf{x} \mid \|\mathbf{q}\|^2 = 1\}$.

5.2. Camera Autocalibration

From a set of given Fundamental matrices $\{\mathbf{F}_i\}_{i=1}^m$, we wish to estimate the camera intrinsic matrix \mathbf{K} . Here, we assume that the camera intrinsic is constant across all the images involved during Fundamental matrix estimation. Let ω be the Dual Image of Absolute Conic (DIAC) expressed in \mathbf{K} as $\omega = \mathbf{K}\mathbf{K}^\top$. The simplified Kruppa's equation [28] allows us to express ω in the form of polynomials using Fundamental matrices \mathbf{F}_i . Let $\mathbf{F}_i = \mathbf{U}_i\mathbf{D}_i\mathbf{V}_i$ be the singular value decomposition, with $\mathbf{D} = \text{diag}([r_i, s_i, 0])$. For $\mathbf{U}_i = [\mathbf{u}_{i1} \mid \mathbf{u}_{i2} \mid \mathbf{u}_{i3}]$ and $\mathbf{V}_i = [\mathbf{v}_{i1} \mid \mathbf{v}_{i2} \mid \mathbf{v}_{i3}]$, two independent polynomials of simplified Kruppa's equations are,

$$\begin{aligned} P_{i1}(\omega) &= (r_i s_i \mathbf{v}_{i1}^\top \omega \mathbf{v}_{i2})(\mathbf{u}_{i2}^\top \omega \mathbf{u}_{i2}) + (r_i^2 \mathbf{v}_{i1}^\top \omega \mathbf{v}_{i1})(\mathbf{u}_{i1}^\top \omega \mathbf{u}_{i2}), \\ P_{i2}(\omega) &= (r_i s_i \mathbf{v}_{i1}^\top \omega \mathbf{v}_{i2})(\mathbf{u}_{i1}^\top \omega \mathbf{u}_{i1}) + (s_i^2 \mathbf{v}_{i2}^\top \omega \mathbf{v}_{i2})(\mathbf{u}_{i1}^\top \omega \mathbf{u}_{i2}). \end{aligned} \quad (14)$$

We parameterize ω using $\mathbf{x} \in \mathbb{R}^5$ because ω is a 3×3 matrix with $\omega = \omega^\top$ and $\omega_{(3,3)} = 1$. Given a set of Fundamental matrices $\{\mathbf{F}_i\}_{i=1}^m$, we solve consensus maximization and non-minimal problems of (12) and (13), for $\mathbf{x} \in \mathbb{R}^5$, $P_i(\mathbf{x}) = (P_{i1}(\omega), P_{i2}(\omega))^\top$, and $\mathcal{K} = \{\mathbf{x} \mid \omega \geq 0\}$. The intrinsic \mathbf{K} is recovered using the Cholesky decomposition on ω .

5.3. Non-Rigid Structure-from-Motion (NRSfM)

Method in [34] proposes a framework of modeling NRSfM as a POP using the geometric prior of isometry. It models the rest shape as a Riemannian manifold and the deformed shapes as isometric mappings of the rest shape. The isometric deformation prior is then expressed by relating the metric tensor, the Christoffel symbols parametrized by $k_1, k_2 \in \mathbb{R}$ and the inter-image registration warps in camera coordinates. We borrow the notations and definitions of the Christoffel symbol parameterization from the original problem formulation [34]. As in [34], we summarize the resulting system of polynomials as,

$$P_i(\mathbf{x}) = \mathcal{P}(\mathbf{x}) - \{\mathcal{P}_{1i}(\mathbf{x})\}_{i=2}^n \quad (15)$$

(15) is in fact a system of $n - 1$ independent quartic polynomials for n images that relates the point-wise inter-image warp measurements $\mathbf{q}_{i=2}^n$ to the Christoffel symbols parameterized by $\mathbf{x} = [k_1 \ k_2]^\top$. Solving (15) amounts to solving the isometric NRSfM problem as k_1, k_2 can be used to obtain the Jacobian of each shape with respect to the reference

image coordinates in the camera frame. One can then compute the surface normal at each corresponding point of the n shapes. We solve consensus maximization and the non-minimal problem of (15) to solve the NRSfM problem.

6. Discussion

The \mathcal{L}_1 constraint of (5) (and the ones that follow) can be extended to \mathcal{L}_∞ by replacing all ϵ_i variables by a single variable ϵ , which is a common practice in convex optimization [7]. Similarly, \mathcal{L}_2 norm can also be minimized by imposing the conic constraints on stacked linear vector of polynomials $\mathbf{v}_i(\mathbf{x}) = [\text{trace}(\mathbf{G}_{1i}\mathbf{Y}), \dots, \text{trace}(\mathbf{G}_{mi}\mathbf{Y})]^\top$ of (6) (resp. of (7)) for m polynomials of the i^{th} measurement such that $\|\mathbf{v}_i(\mathbf{x})\|_2 \leq \epsilon_i$. In this process, we introduce one variable for each measurement, therefore the time complexity of our method is expected to be linear on the number of measurements. Although these auxiliary variables may seem to introduce overhead, they are in fact helpful. These variables allow us to express the non-minimal problems as an SOCP. The SOCP constraints enable us to easily extend our non-minimal formulation to that of consensus maximization, where variables ϵ_i naturally turn into inlier threshold. Alternatively for non-minimal problems, one could choose to minimize $\sum_i \|\mathbf{v}_i(\mathbf{x})\|_2^2$ (or even $\sum_i p_i(\mathbf{x})^2$) directly, using the standard SOS optimization methods. Such formulation not only compromises the flexibility, but also adds burden by increasing the degree of polynomials.

7. Experiments

We conducted several experiments with synthetic and real datasets. The synthetic data was generated using the toolbox of [8], in the very similar setups, whereas, the real datasets used are Fountain and Herz-Jesu [44], and Flag [45], Hulk and Tshirt [34]. In the case of real datasets, only the outliers were synthetically generated for quantitative evaluations. All our results for non-minimal solvers are generated after the consensus maximization, except for the minimal case. Our framework is implemented in MATLAB2015a and all the optimization problems are solved using MOSEK [29]. All experiments were carried out on a 16GB RAM Pentium i7/3.40GHz. Some qualitative results obtained for different applications, discussed later, are shown in Figure 1.

7.1. Rigid Body Transformation

We conducted a large amount of synthetic experiments to explore the behaviours of the proposed framework on rigid body transformation estimation, both for non-minimal and consensus maximization problems. This is mostly because there exists three different non-minimal solvers, namely **Briales'17** [3], **Olsson'08** [31] and **Olsson-BnB** [32], comparison against them is one of our interests. In the first ex-



Figure 1: Non-minimal results obtained using our framework. From left to right: rigid registration (model: green; data: red) on the public dataset [15]; 3D reconstruction of the Fountain sequence obtained using the estimated camera intrinsics; non-rigid 3D reconstruction (ground truth: green; reconstruction: red) of tshirt with Lasserre’s relaxation.

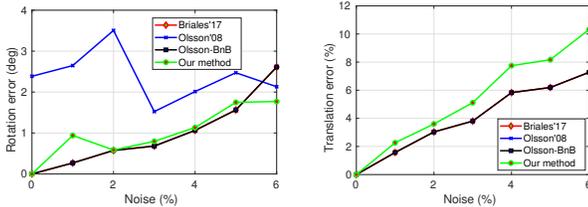


Figure 2: Results for minimal (3 points) rigid body transformation estimation. Left to right: noise vs. errors in rotation angles and translation for four different methods.

periment, we tested performance of different methods in the minimal setup with varying levels of noise. The obtained results are reported in Figure 2. Note that the problem of rigid body transformation is a 3-point problem. One can observe for Figure 2 that our method performs very competitively compared to the globally optimal method **Olsson-BnB** in terms of estimating rotation and translation parameters, whereas, our method is very competitive in terms of time, shown in Figure 3 (left), with **Briaies’17**. Experiments with **Olsson’08** shows that this method does not support the minimal setup.

Experiments for non-minimal setups were also conducted, to report similar measures as that of minimal case, by varying the number of points for fixed noise level. Results of these experiments are reported in Figure 4. Here we can observe that our method still behaves very similar to **Olsson-BnB** in terms of accuracy. Note that the time, shown in Figure 3 (right), for our method increases linearly with the number of points. The reason behind this has already been discussed in Section 6, where we also suggest techniques to speed up by compromising the flexibility. Unlike the minimal case, **Olsson’08** performs very similar to **Olsson-BnB** and **Briaies’17**.

Using the flexibility offered by our method, we conducted experiments with increasing amount of outliers and report the experiments in Figure 5. Our framework provides consistent results even when outliers are present in the measurements, as it passes through the process of consensus

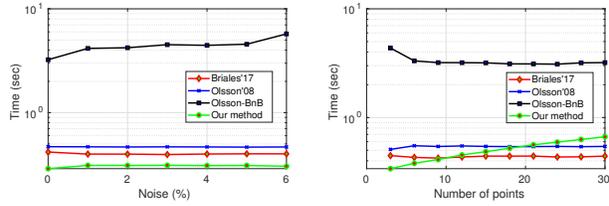


Figure 3: Time taken for minimal case (3 points) with noise and non-minimal case (of fixed 1.0% noise) with number points. Results for rigid body transformation estimation.

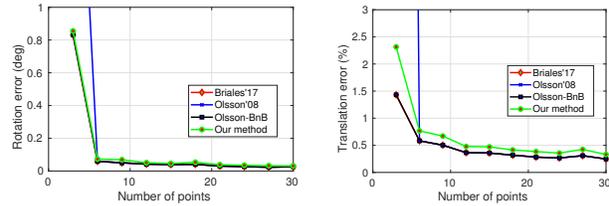


Figure 4: Non-minimal (1.0% noise) results for rigid body transformation estimation. Errors in rotation angles and translation vs. increasing number of points.

maximization. Although it is unfair to compare with non-minimal solver against a framework that offers both consensus maximization and non-minimal solution, Figure 5 shows the impact of outliers when only non-minimal solvers are used. Note that our method performs well, even when the data is contaminated with 90% outliers (10 inliers and 90 outliers). As expected, the non-minimal solvers remain consistent in time with the increase of outliers, of which they are indifferent, as shown in Figure 6 (left). On the other hand, our framework filters outliers prior to solving the non-minimal problem. Our method was also compared with a global consensus maximization **Speciale’17** [42], where both methods perform very similarly in terms of optimality, whereas, **Speciale’17** performs faster due to its problem specific constraints. The time comparison with increasing number of outliers is shown in Figure 6 (right).

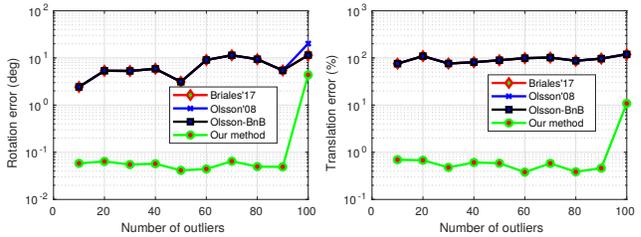


Figure 5: Our framework vs. other non-minimal solvers with increasing outliers (and fixed 10 inlier correspondences) for rigid body transformation estimation. Errors in rotation and translation estimation with increasing outliers.

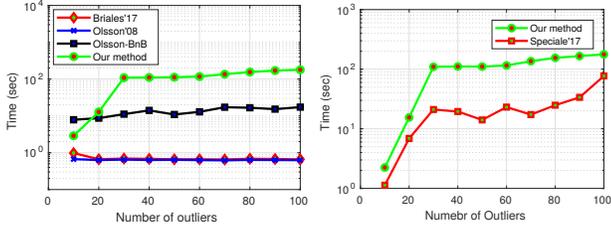


Figure 6: Time comparison with global non-minimal solvers (left) and global consensus maximization method (right), with increasing outliers (and fixed 10 inlier correspondences) for rigid body transformation estimation.

7.2. Camera Autocalibration

We conducted experiments for camera autocalibration on two real datasets: Fountain and Herz-Jesu from [44]. The results obtained by our framework is compared against that of three existing global methods for autocalibration: **LMI direct** method from [20], **Rank-3 direct** method from [12], and a **Stratified** method from [13]. All of these three methods assume that the projective motion required for autocalibration is free of outliers, and minimize the global cost in an optimal manner. Therefore, they can be thought as some kind of non-minimal solvers. The projective reconstruction required for [20], [12] and [13] were obtained using [30]. We also compared our method with two local methods for camera calibration, namely **Practical** form [19] and **Simplified-Kruppa** from [28]. We provide the quantitative results for calibration accuracy, by computing errors on the camera intrinsic parameters. Three different error measurement metrics are used: errors in focal length Δf , principal point Δuv , skew Δs . Obtained results by all six methods are reported in Table 1.

For the experiments of consensus maximization, we synthetically introduced the outlier Kruppa’s equations. The outliers are added in an increasing manner upto 80%. Our consensus maximization method, with Shor’s relaxation, is able to detect all inliers and outliers correctly for both datasets. We tracked the number of pessimistic and optimistic inliers for increasing BnB iterations, which is shown in the Figure 7 (left). The consensus maximization experiments are compared with a global consensus maximization method for autocalibration, namely **Paudel’18** from [37]. For a fair comparison, the constraint on the bounds of the DIAC are chosen as in [37]. We therefore assume that the focal length lies within [1 10] interval relative to the image size, aspect ratio lies between 0.7-1.25, principal points lie around the image center within a radius of $(\frac{1}{4})^{th}$ of the image size, and the skew is close to zero. With these assumptions, we derive bounds on DIAC using the interval analysis arithmetic [2], similarly for the corresponding variable X in (10). Note that these are valid assumption in most cases for camera calibration. The results obtained by our

Dataset	Method	Δf	Δuv	Δs	Time(s)
Fountain (11-views)	Practical [19]	0.0117	0.0149	0.0037	0.36
	Stratified [13]	0.0777	0.0969	0.0125	388.24
	Rank-3 Direct [12]	0.0100	0.0147	0.0044	5.75
	LMI Direct [20]	0.0506	0.0269	0.0024	156.88
	Simplified-Kruppa [28]	2.93e-05	0.0069	3.23e-05	1.88
	Ours	0.0060	0.0061	6.46e-04	0.83
Herz-jesu (8-views)	Practical [19]	0.0017	0.0113	0.0068	0.36
	Stratified [13]	0.7231	0.4462	0.3232	380.72
	Rank-3 Direct [12]	0.0026	0.0096	0.0069	1 6.54
	LMI Direct [20]	0.0138	0.0086	0.005	115.61
	Simplified-Kruppa [28]	4.46e-05	0.0069	3.40e-05	0.53
	Ours	0.0128	0.0220	8.77e-04	0.81

Table 1: Our framework vs. two local [19, 28] and three global [13, 12, 20] methods in non-minimal setup of camera autocalibration on two real datasets. Views-1 are the number of measurements used in all the experiments.

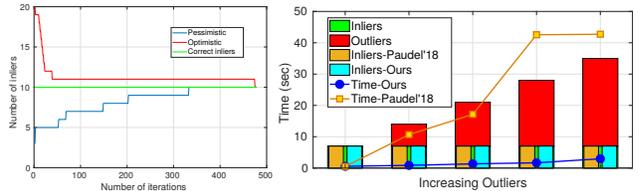


Figure 7: Left: convergence graph for Fountain dataset. Right: time taken and inliers detected by our method and [37], with increasing outliers and fixed inliers on Herz-jesu.

framework as well that of **Paudel’18** are reported in Figure 7 (right) for increasing outliers, showing the detected inliers and time taken by both methods. These results also show that our framework can greatly benefit if the bounds on the sought parameters are also known. In fact, this may very often be the case in many 3D vision problems.

7.3. Non-Rigid Structure-from-Motion

Experiments with NRSfM were conducted in the similar setup of **Parashar’18** [35]. From N images, $2N - 2$ polynomials of 2 variables on Christoffel symbols k_1 and k_2 , i.e. (15), were extracted using the theory presented in Section 5.3. For these polynomials, we first computed the non-minimal solution using the proposed framework with Shor’s relaxation. However, the Shor’s relaxation alone did not provide us satisfactory solutions. Note that (15) are expressed using the second-order measurements of the local correspondences, therefore their coefficients are very sensitive to noise. Moreover, polynomials of (15) are solved for each point independently. In such cases, non-minimal solvers are very valuable, where the over-determined system of quartic polynomials are provided by multiple measurements across views. Therefore, the method of [35] uses the hierarchy of SOS relaxations to obtain the desired solution. In this context, we observed that the Lasserre’s relaxation is indeed necessary, because the non-minimal solutions obtained using Shor’s relaxation were not on par with that of SOS method. Therefore, we show the difference

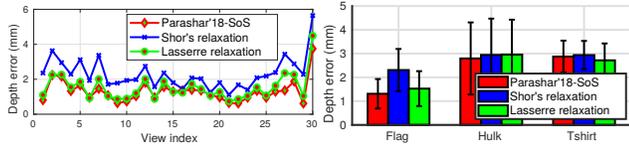


Figure 8: Results for non-minimal isometric non-rigid reconstruction. Left to right: depth and normal errors for flag dataset, depth errors of three methods on different datates.

between Shor’s and Lasserre’s relaxations for NRSfM. In previous two problems, their differences were not very significant in terms of accuracy, but only in terms of speed.

We use the datasets, Flag [45], Hulk and Tshirt [34] to evaluate our non-minimal solver. Our obtained solutions are compared with the baseline Parashar’18 [35] in Figure 8 for the non-minimal case, where the reconstructed depth error is shown across views on the left, and across datasets on the right. Note from Figure 8 that the baseline method performs significantly better compared to only the Shor’s relaxations-based method. However, when we use Lasserre’s relaxation of order one (i.e. $s = 1$ in (4)), the solution becomes very competitive to that of the hierarchy of SOS relaxations used in [35]. This shows that the higher order relaxations do not offer significant improvements at least for this particular formulation of NRSfM. In fact, it is generally agreed that NRSfM is a very challenging problem in 3D vision. Our observation from real dataset shows that higher order relaxations may not be necessary, in a wide range of applications, including the case of isometric NRSfM. Of course, our observation may be biased from only three problems that we have tested. Especially, second degree polynomials for rigid body transformation and camera calibraion, and degree four polynomials that involve only two unknown variables at a time in the case of NRSfM. Nevertheless, one needs to be aware that lower order relaxations could be tried first, before divulging into the higher ones. If the lower order relaxations already offer satisfactory solutions, this not only allows one to obtain non-minimal solutions faster but also allows to use rather non-obvious polynomial problems within the global framework of consensus maximization by using the BnB paradigm.

To support our claim that even the lower degree relaxations are sufficient for consensus maximization, we conducted several experiments with various amount of synthetic outliers on aforementioned real datasets. For the setup of NRSfM we were able to detect almost all outliers when outliers upto 70% were introduced. Results for one such instance of consensus maximization, with 50% outliers views, are reported in Figure 9, for all three datasets. In Figure 9, we show the estimated depth error across inlier views on the left. As expect, results with Lasserre’s relaxations are significantly better than that of the Shor’s relaxation. Nevertheless, Shor’s relaxation still shows its expected behaviour.

Although, it may not be very interesting to compare the results of a non-minimal solver, i.e. Parashar’18, against that of a consensus maximization method in general, it is somehow different in this case. One aspect that we have not yet discussed is the power of SOS methods. It is generally known that the SOS solvers are robust to noise [36], SOS solver appears to be relatively stable even in the presence of outliers. In the first glance, we thought this could be because of the iterative refinement of outliers within the algorithm of Parashar’18. This turns out not to be the case. The SOS method consistently performed reasonably well, even when the iterative refinement process was removed. This is because of two reasons: one is the robustness of SOS methods towards a moderate amount of outliers; other is the failure to obtain a valid solution means in some sense the detection of outliers. Such outlier detection is specific to the problem formation of [35] because the POP in this case is defined point-wise.

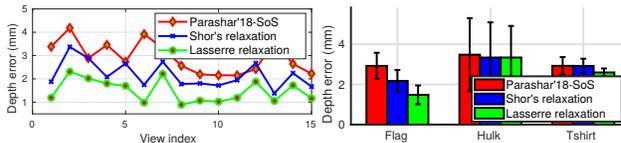


Figure 9: Consensus maximization results for isometric non-rigid reconstruction. Left to right: depth and normal errors for flag dataset, depth errors of three methods on different datasets with 50% synthetically added outliers.

8. Conclusion

In this paper, we demonstrated that a proper usage of the existing tools in numerical algebraic geometric POP can be used in a straightforward way in many 3D vision problems. This is achieved by using the known “good approximate solutions” expressed as convex SDP formulation. On the one hand, we argue that the existing solutions can be used in their current form to solve many 3D vision problems, especially for which the optimal non-minimal solvers have not been devised yet. We also discussed about a good practice to formulate POP. Using our theoretical reasoning, we made suggestions for a generic relaxed non-minimal solver that is suitable for may 3D vision problems. We further argued that the standard method for polynomial relaxations are indeed powerful, which can also be used for consensus maximization in a deterministic manner. We have supported our suggestions/claims using several experiments of three diverse problems in 3D vision. We reach to this conclusion mainly because many polynomials in 3D vision problems are inherently of low degrees with limited variables.

References

- [1] Amir Ali Ahmadi, Georgina Hall, Ameet Makadia, and Vikas Sindhwani. Geometry of 3d environments and sum of squares polynomials. *arXiv preprint arXiv:1611.07369*, 2016. 2
- [2] Götz Alefeld and Günter Mayer. Interval analysis: theory and applications. *Journal of computational and applied mathematics*, 121(1):421–464, 2000. 7
- [3] Daniel Barath, Tekla Toth, and Levente Hajder. A minimal solution for two-view focal-length estimation using two affine correspondences. 2017. 5
- [4] Adrien Bartoli. Generalizing the prediction sum of squares statistic and formula, application to linear fractional image warp and surface fitting. *International Journal of Computer Vision*, 122(1):61–83, 2017. 2
- [5] Jean-Charles Bazin, Yongduek Seo, Richard Hartley, and Marc Pollefeys. Globally optimal inlier set maximization with unknown rotation and focal length. In *European Conference on Computer Vision*, pages 803–817. Springer, 2014. 1
- [6] Stephen Boyd and Lieven Vandenbergh. Semidefinite programming relaxations of non-convex problems in control and combinatorial optimization. In *Communications, Control, and Signal Processing*, pages 279–287. Springer, 1997. 2, 3
- [7] Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004. 3, 5
- [8] Jesus Briales, Javier Gonzalez-Jimenez, et al. Convex global 3d registration with lagrangian duality. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 5
- [9] Jesus Briales, Laurent Kneip, SIST ShanghaiTech, and Javier Gonzalez-Jimenez. A certifiably globally optimal solution to the non-minimal relative pose problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 145–154, 2018. 1, 2
- [10] Zhipeng Cai, Tat-Jun Chin, Huu Le, and David Suter. Deterministic consensus maximization with biconvex programming. In *European Conference on Computer Vision*, 2018. 1
- [11] Manmohan Chandraker, Sameer Agarwal, Fredrik Kahl, David Nister, and David Kriegman. Autocalibration via rank-constrained estimation of the absolute quadric. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007. 2
- [12] Manmohan Chandraker, Sameer Agarwal, Fredrik Kahl, David Nistér, and David Kriegman. Autocalibration via rank-constrained estimation of the absolute quadric. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007. 7
- [13] Manmohan Chandraker, Sameer Agarwal, David Kriegman, and Serge Belongie. Globally optimal affine and metric upgrades in stratified autocalibration. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007. 7
- [14] Tat-Jun Chin, Zhipeng Cai, and Frank Neumann. Robust fitting in computer vision: Easy or hard? In *European Conference on Computer Vision*, 2018. 1, 4
- [15] Tat-Jun Chin, Yang Heng Kee, Anders Eriksson, and Frank Neumann. Guaranteed outlier removal with mixed integer linear programs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5858–5866, 2016. 6
- [16] Tat-Jun Chin, Pulak Purkait, Anders Eriksson, and David Suter. Efficient globally optimal consensus maximisation with tree search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2413–2421, 2015. 1
- [17] John W Chinneck. *Feasibility and Infeasibility in Optimization: Algorithms and Computational Methods*, volume 118. Springer Science & Business Media, 2007. 4
- [18] Anders P. Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging and strong duality. In *CVPR*, 2018. 1
- [19] Riccardo Gherardi and Andrea Fusiello. Practical autocalibration. *Computer Vision—ECCV 2010*, pages 790–801, 2010. 7
- [20] Adlane Habed, Danda Pani Paudel, Cédric Demonceaux, and David Fofi. Efficient pruning lmi conditions for branch-and-prune rank and chirality-constrained estimation of the dual absolute quadric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 493–500, 2014. 7
- [21] Fredrik Kahl and Didier Henrion. Globally optimal estimates for geometric reconstruction problems. *International Journal of Computer Vision*, 74(1):3–15, 2007. 2
- [22] Laurent Kneip, Hongdong Li, and Yongduek Seo. Upnp: An optimal $\mathcal{O}(n)$ solution to the absolute pose problem with universal applicability. In *European Conference on Computer Vision*, pages 127–142. Springer, 2014. 1, 2
- [23] Pierre-Yves Lajoie, Siyi Hu, Giovanni Beltrame, and Luca Carlone. Modeling perceptual aliasing in SLAM via discrete-continuous graphical models. *IEEE Robotics and Automation Letters*, 4(2):1232–1239, 2019. 1
- [24] Jean B Lasserre. Convergent lmi relaxations for nonconvex quadratic programs. In *Decision and Control, 2000. Proceedings of the 39th IEEE Conference on*, volume 5, pages 5041–5046. IEEE, 2000. 2
- [25] Jean B Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on optimization*, 11(3):796–817, 2001. 2, 3
- [26] Jean B Lasserre. Semidefinite programming vs. lp relaxations for polynomial programming. *Mathematics of operations research*, 27(2):347–360, 2002. 3
- [27] Hongdong Li. Consensus set maximization with guaranteed global optimality for robust geometry estimation. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 1074–1080. IEEE, 2009. 1
- [28] Manolis IA Lourakis and Rachid Deriche. *Camera self-calibration using the singular value decomposition of the fundamental matrix: From point correspondences to 3D measurements*. PhD thesis, INRIA, 1999. 5, 7

- [29] ApS MOSEK. The mosek optimization toolbox for matlab manual, version 8.0. *MOSEK ApS, Denmark*, 2015. 5
- [30] John Oliensis and Richard Hartley. Iterative extensions of the sturm/triggs algorithm: Convergence and nonconvergence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2217–2233, 2007. 7
- [31] Carl Olsson and Anders Eriksson. Solving quadratically constrained geometrical problems using lagrangian duality. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, pages 1–5. IEEE, 2008. 5
- [32] Carl Olsson, Fredrik Kahl, and Magnus Oskarsson. Branch-and-bound methods for euclidean registration problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):783–794, 2009. 1, 2, 5
- [33] Danda Pani Paudel, Adlane Habed, Cédric Demonceaux, and Pascal Vasseur. Robust and optimal sum-of-squares-based point-to-plane registration of image sets and structured scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2048–2056, 2015. 1, 2
- [34] Shaifali Parashar, Daniel Pizarro, and Adrien Bartoli. Isometric non-rigid shape-from-motion in linear time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4679–4687, 2016. 2, 5, 8
- [35] Shaifali Parashar, Daniel Pizarro, and Adrien Bartoli. Isometric non-rigid shape-from-motion with riemannian geometry solved in linear time. *IEEE transactions on pattern analysis and machine intelligence*, 40(10):2442–2454, 2018. 7, 8
- [36] Pablo A Parrilo. *Structured semidefinite programs and semi-algebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000. 2, 8
- [37] Danda Pani Paudel and Luc Van Gool. Sampling algebraic varieties for robust camera autocalibration. In *European Conference on Computer Vision*, pages 275–292. Springer, Cham, 2018. 1, 7
- [38] Victoria Powers and Thorsten Wörmann. An algorithm for sums of squares of real polynomials. 1998. 2
- [39] David M. Rosen, Luca Carlone, Afonso S. Bandeira, and John J. Leonard. Se-sync: A certifiably correct algorithm for synchronization over the special euclidean group. *I. J. Robotics Res.*, 38(2-3), 2019. 1
- [40] Naum Z Shor. Quadratic optimization problems. *Soviet Journal of Computer and Systems Sciences*, 25:1–11, 1987. 2
- [41] Naum Z Shor. Quadratic optimization problems. *Soviet Journal of Circuits and Systems Sciences*, 25(6):1–11, 1987. 2
- [42] Pablo Speciale, Danda Pani Paudel, Martin R Oswald, Till Kroeger, Luc Van Gool, and Marc Pollefeys. Consensus maximization with linear matrix inequality constraints. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5048–5056. IEEE, 2017. 1, 6
- [43] Pablo Speciale, Danda P Paudel, Martin R Oswald, Hayko Riemenschneider, Luc V Gool, and Marc Pollefeys. Consensus maximization for semantic region correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7317–7326, 2018. 1
- [44] C. Strecha, W. von Hansen, L. Van Gool, P. Fua, and U. Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008. 5, 7
- [45] Ryan White, Keenan Crane, and David A. Forsyth. Capturing and animating occluded cloth. In *SIGGRAPH 2007*, 2007. 5, 8
- [46] Jiaolong Yang, Hongdong Li, and Yunde Jia. Optimal essential matrix estimation via inlier-set maximization. In *European Conference on Computer Vision*, pages 111–126. Springer, 2014. 1