Fast certifiable algorithm for the absolute pose estimation of a camera

Mercedes Garcia-Salguero, Elijs Dima, André Mateus and Javier Gonzalez-Jimenez

Abstract

Estimating the absolute pose of a camera given a set of N points and their observations is known as the resectioning or Perspective-n-Point (PnP) problem. It is at the core of most computer vision applications and it can be stated as an instance of 3D registration with point-line distances, making the error quadratic in the unknown pose. The PnP problem, though, is nonconvex due to the constraints associated with the rotation, and iterative algorithms may get trapped into any suboptimal solutions without notice. This work proposes an efficient certification algorithm for central and noncentral cameras that either confirms the optimality of a solution or is inconclusive. We exploit different sets of constraints for the rotation to assess their performance in terms of certification. Two of the formulations lack the Linear Independence Constraint Qualification (LICQ) while one of them has more constraints than variables. This hinders the usage of the 'standard' procedure which estimates the Lagrange multipliers in closed-form. To overcome that, we formulate the certification as an eigenvalue optimization and solve it through a line-search method. Our evaluation on synthetic and real data shows that minimal formulations certify most solutions (more than 90% on real data) whereas redundant formulations are able to certify all of them and even random problem instances. The proposed algorithm runs in microseconds for all these formulations.

Index Terms

Certifiable algorithms, Point-Line registration, resectioning problem, dual problem, Eigenvalue optimization

I. Introduction

Estimating the absolute pose (rotation and translation) of a camera given a set of N world points and their observations is known as the resectioning or Perspective-n-Point (PnP) problem. It appears in different 3D computer vision applications, and is one of the main steps in reconstruction pipelines, including Structure-from-Motion (SfM) and Simultaneous Localization and Mapping (SLAM).

The common approach, known as Bundle Adjustment (BA), estimates the pose and 3D world points jointly. However, BA turns out to be highly nonconvex and thus suffers from multiple local minima that are not easily detected. A good initialization may avoid these minima and may improve the convergence of the algorithm, although even the initialization algorithm may be nonconvex, suffering also from local minima. Therefore, being able to assess the quality of the PnP solution from these initial algorithms is of special importance for the stability and correct performance of the application.

In this work we tackle the optimality certification of the solution to the resectioning problem. Whereas an intuitive metric to be minimized is the reprojection error as Euclidean distance on the image plane, it presents some issues when the camera model does not have an actual image plane, e.g.,for generalized cameras [16]. For these general cases, observations are defined as 3D unit-norm vectors and other metrics that capture this information should be employed instead. An alternative approach minimizes the error between the 3D point and the ray originated from the camera center towards the observation, thus posing the problem as a special case of Point-Line registration, for which extensive literature is available, see [22] and references therein. The problem, though, is nonconvex due to the constraints associated with the rotation, implying that the quality of the solution cannot be guaranteed for noisy data. Iterative solvers, however, tend to return good solutions and even the global optimum despite being at risk of getting trapped in local minima. This empirical observation motivates the use of iterative solvers in real-world applications, but also the development of algorithms that are able to certify that the estimated solution is the optimum. Whereas some of these approaches also estimate the solution, it is usually more efficient to just certify a provided solution, as it has been done for other problems, as described in Section II.

Contributions In this work we tackle the certification of the resectioning problem for central [15] and noncentral cameras [16] and our contribution is threefold. First, we propose two closed-form certifiers based on minimal formulations. Second, we propose two more certifiers that rely on redundant formulations and are solved via an eigenvalue optimization. Both approaches can be implemented with any linear algebra library. Our third and last contribution compares the performance of said certifiers on synthetic and real data. Even for random, synthetic data the redundant formulations under the proposed algorithm certify most solutions. For synthetic data with moderate number of correspondences and low noise the minimal formulations certify most solutions. The redundant formulations, on the other hand, are able to detect optimal solutions even for extreme cases with low number of correspondences and highly noisy data. On real data all the formulations certify more than 90% of solutions, although only the fully redundant formulation certifies all of them.

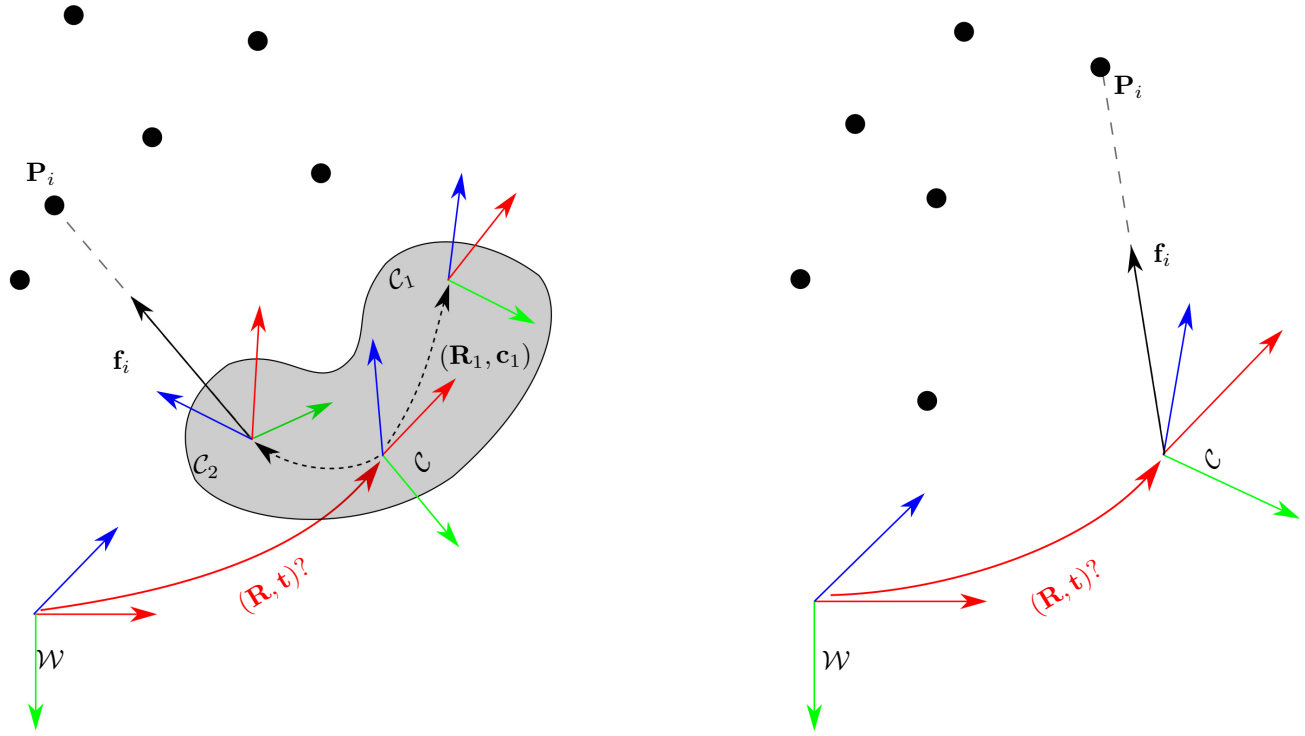

(a) Generalized camera model.

(b) Central camera model.

Fig. 1: (a) Generalized; and (b) central resectioning problem where, given a set of N observations $\mathbf{f}_\mathrm{i}$ of N 3D points $\mathbf{P}_\mathrm{i}$ for $i = 1, \ldots, N$, we aim to estimate the absolute pose $\mathbf{R}, \mathbf{t}$ of the camera $\mathcal{C}$ wrt the world frame $\mathcal{W}$.

II. Related work

In this section we review recent works that address the resectioning problem with a focus on optimality certification.

Hesch and Roumeliotis [18] tackled the central case including unknown depth of the observations. By removing these scales and the translation, the authors reduced the number of unknowns to three through the Cayley-Gibbs-Rodriguez parameterization of the rotation. This makes the problem an instance of an unconstrained least-squares, that was solved using the Macaulay matrix. Sweeney et al. [32] generalized [18] to accommodate more than one camera (the noncentral case, see [16]) and the overall scale. They also solved the problem through the Macaulay matrix following a similar procedure. Zheng et al. [37] solved the central problem by parameterizing the rotation through the non-unit quaternion. The unconstrained problem was solved through a Gröbner-basis method, reportedly in milliseconds for any number of correspondences. Recently, Wu et al. [35] unified the formulations for some computer vision problems, including the resectioning problem, that involved rotations under the same quadratic framework, and posed the general, quadratic problem in terms of quaternions. The authors also rely on a Gröbner-basis method and report computational time of milliseconds.

Despite being optimal in theory, Gröbner-basis methods may suffer from numerical instabilities [9]. Another approach that allows to estimate and certify, under some conditions, the optimality of the solution uses convex relaxations of the original problem. Schweighofer and Pinz [29] formulated the noncentral problem with the quaternion parameterization and relied on a Sum-of-squares approach (Lasserre's hierarchy [19]) to solve the problem with optimality guarantees. The derived SemiDefinite Positive (SDP) problem was solved with off-the-shelf tools. Sun and Deng [31] included line correspondences into the central problem, resorted to a quaternion parameterization for the rotation and leveraged also Lasserre's hierarchy to solve the problem with optimality guarantees. Agostinho et al. [2] further improved the above-mentioned approach by using the rotation matrix, obtaining an SDP solver. All these approaches employ the point-line distance as error on the optimization problems, which makes the central and noncentral resectioning problem instances of 3D point-line registration. Previous approaches for general registration rely on Gröbner-basis methods e.g., [22], [34], [39], Branch-and-Bound e.g., [23] or convex relaxations as SDP solvers e.g., [6]. We refer the interested readers to those works and references therein for further information.

Nonetheless, all the above-mentioned approaches suffer from high computational times, numerical errors and/or the need of specific tools for their resolution. Instead, iterative algorithms are more common in practice and have good performance, provided the initial guess is good enough. Since the rotation space is well-known and minimal

parameterizations are available, it is possible to use (unconstrained) Newton's method, e.g., [38]. Other approaches opt for optimizing directly on the space of rotations, which avoids the conversion to this minimal parameterization, e.g., [10]. Regardless of the representation of the search space, the resectioning problem is nonconvex and in general presents more than one local minimum [28]. Therefore, there is no guarantee that the iterative algorithms return the global optimum.

An alternative approach invokes these iterative algorithms but introduces an optimality certifier for the final solution, thus benefiting of the efficiency of the former and the optimality guarantees of the latter. We are interested in certifiers that are derived from convex relaxations, concretely the dual problem [4]. These certifiers can be computed efficiently provided some conditions are fulfilled by the problem, among them, having more variables than constraints and the Jacobian of the constraints being full-rank at the solution. These conditions were enforced in previous works by relaxing the domain of 3D rotations $\mathcal{SO}(3)$ to 3D orthogonal matrices $\mathcal{O}(3)$, e.g.,pose graph optimization [5], [26], rotation averaging [11], etc. Our evaluation in Section V shows that whereas this relaxation can work well, it fails for some problem instances. Adding redundant constraints and eventually enforcing the determinant constraints increases the percentage of detected optimal solutions, see also [33]. However, these formulations lack the above-mentioned conditions, and alternative approaches must be adopted to achieve efficient certifiers. This was the reason behind the proposal in [14] where a low-rank, iterative certifier was devised for the relative pose problem.

III. Problem formulation

The resectioning or PnP problem aims to obtain the camera pose as rotation $\mathbf{R} \in \mathbb{R}^{3\times3}$ and translation $\mathbf{t} \in \mathbb{R}^3$ given a set of N world points $\mathbf{P}_\mathrm{i}, i = 1, \dots, N$ and their corresponding observations $\mathbf{f}_\mathrm{i}$ and without loss of generality and to accommodate all camera models we consider the observations $\mathbf{f}_\mathrm{i}$ to have unit norm, that is, $\mathbf{f}_\mathrm{i} \in \mathcal{S}^2 \subset \mathbb{R}^3$. These observations may have been captured by M rigidly attached cameras in the form of a camera rig, see Figure 1a. For one camera ($M = 1$) the problem is reduced to the central configuration (Figure 1b). We develop the proposal for the noncentral configuration and provide, when needed, the formulation for the central case.

We state the problem as the minimization of the squared Euclidean distance between the transformed 3D point $\mathbf{Q}_i$ and the ray $\mathbf{q}_i \in \mathcal{S}^2$ passing through the observation, making the resectioning problem a special case of the Point-Line registration with squared error $\epsilon_i^2 = \|(\mathbf{I}_3 - \mathbf{q}_i\mathbf{q}_i^T)\mathbf{Q}_i\|_2^2$. We consider that the i-th observation $\mathbf{f}_\mathrm{i}$ is observed by the j-th camera with known relative pose $(\mathbf{R}_\mathrm{j}, \mathbf{c}_\mathrm{j})$ wrt the center of the rig. Therefore, the transformed 3D point considering the unknown pose $(\mathbf{R}, \mathbf{t})$ is $\mathbf{R}\mathbf{P}_\mathrm{i} + \mathbf{t} - \mathbf{c}_\mathrm{j}$ and the ray $\mathbf{q}_i$ associated with the observation is $\mathbf{R}_\mathrm{j}\mathbf{f}_\mathrm{i}$. Under this metric, the problem reads

$$h^\star = \min_{\mathbf{R}\in\mathbb{R}^{3\times3},\mathbf{t}\in\mathbb{R}^3} \sum_{i=1}^{N} \|\big(\mathbf{I}_3 - (\mathbf{R}_\mathrm{j}\mathbf{f}_\mathrm{i})(\mathbf{R}_\mathrm{j}\mathbf{f}_\mathrm{i})^T\big)(\mathbf{R}\mathbf{P}_\mathrm{i} + \mathbf{t} - \mathbf{c}_\mathrm{j})\|_2^2 \quad \text{subject to } \mathbf{R} \in \mathcal{SO}(3), \tag{O}$$

where $\mathcal{SO}(3)$ is the space of 3D rotations. If the set $\mathcal{SO}(3)$ is defined by quadratic expressions, then the problem becomes an instance of a Quadratically Constrained Quadratic Program (QCQP). These types of problems are, in their general form, nonconvex and NP-hard, although we can derive convex relaxations that allow us to approximate their solutions and even solve them when the relaxation is tight. Whereas previous works have followed this approach, see Section II, in this work we aim to solve the problem in a more efficient way. We first obtain a solution for O through an iterative algorithm and then try to certify it as the global optimum. The certification is the key step and we devote the remaining of this manuscript to it.

To derive the certifier we first re-write O as a QCQP in its standard form for which a convex relaxation known as the dual problem can be derived through standard methods, see [4]. As commented in previous works, see Section II, the set of rotations $\mathcal{SO}(3)$ is defined by quadratic, nonhomogeneous expressions that depend only on the rotation $\mathbf{R}$ whose entries we collect column-wise in the vector $\mathbf{r} = \mathrm{vec}(\mathbf{R}) \in \mathbb{R}^9$. The cost is written as quadratic forms in $\mathbf{r}$ and $\mathbf{t}$ through algebraic manipulation (see A), taking the form: $h(\mathbf{R}, \mathbf{t}) = \mathbf{r}^T\tilde{\mathbf{C}}_\mathrm{R}\mathbf{r} + 2\mathbf{t}^T\mathbf{C}_\mathrm{tR}\mathbf{r} + \mathbf{t}^T\mathbf{C}_\mathrm{t}\mathbf{t} - 2\mathbf{c}_t^T\mathbf{t} - 2\tilde{\mathbf{c}_R}^T\mathbf{r} + \tilde{c}$. Since the translation $\mathbf{t}$ is an unconstrained variable, we can marginalize it (see B), obtaining a problem that only depends on the rotation $\mathbf{R}$ as $h(\mathbf{R}, \mathbf{t}) = f(\mathbf{R}) = \mathbf{r}^T\mathbf{C}_\mathrm{R}\mathbf{r} - 2\mathbf{c}_R^T\mathbf{r} + c$, where $\mathbf{C}_\mathrm{R} \in \mathbb{R}^{9\times9}, \mathbf{c}_R \in \mathbb{R}^{9\times1}$ and $c \in \mathbb{R}$. For simplicity we collect the entries of the rotation into the vector $\mathbf{x} \doteq [\mathbf{r}^T, y]^T \in \mathbb{R}^{10}$, where the scalar $y = \pm1$ is included to make the cost and constraints homogeneous in $\mathbf{x}$. We create the 10×10 cost matrix $\mathbf{C}$

$$\mathbf{C} \doteq \begin{pmatrix} \mathbf{C}_\mathrm{R} & -\mathbf{c}_R \\ -\mathbf{c}_R & c \end{pmatrix}, \tag{1}$$

and similarly the constraints are given by $\mathbf{x}^T\mathbf{A}_i\mathbf{x} = 0, i = 1, \dots, K$, where the explicit form of $\mathbf{A}_i \in \mathbb{S}^{10}$ and the number of constraints K depend on which definition of the rotation space we choose, to be defined later in Section IV-B. We also need to include the constraint for the homogeneous variable y, which is also quadratic in $\mathbf{x}$, as $y^2 = 1 \Leftrightarrow \mathbf{x}^T\mathbf{L}\mathbf{x} = 1$. Problem O is then expressed in its general QCQP form as

$$f^\star = \min_{\mathbf{x}\in\mathbb{R}^{10}} \mathbf{x}^T\mathbf{C}\mathbf{x} \quad \text{subject to } \mathbf{x}^T\mathbf{L}\mathbf{x} = 1,\ \mathbf{x}^T\mathbf{A}_i\mathbf{x} = 0, i = 1, \dots, K. \tag{QCQP}$$

The central case with $M = 1$ has $\mathbf{c}_R = \mathbf{0}_{9\times 1}$ and $c = 0$.

Iterative solver: Before continuing, we include a brief note about the resolution of problem O using the the Riemannian machinery, that we leverage in this proposal, as it has been shown to perform well for other similar problems, e.g.,pose graph optimization [5], [26] or relative pose [14]. We provide here only the main aspects and refer the reader to those references for further information about this approach. As solver we employ the Trust region method [1] which comes with good convergence properties. The optimization takes place on the direct product of the rotation and the 3D Euclidean spaces, and the associated operators can be found in e.g., [1].

With respect to the model (the cost) the re-formulation in problem QCQP allows us to derive the second-order information (Hessian-product vector) in a simple form: (1) the cost is $f(\mathbf{R},\mathbf{t}) = \mathbf{r}^T\tilde{\mathbf{C}}_{\mathrm{R}}\mathbf{r} + 2\mathbf{t}^T\mathbf{C}_{\mathrm{tR}}\mathbf{r} + \mathbf{t}^T\mathbf{C}_{\mathrm{t}}\mathbf{t} - 2\mathbf{c}_t^T\mathbf{t} - 2\tilde{\mathbf{c}_R}^T\mathbf{r} + \tilde{c}$; (2) the gradient is $\mathbf{G}_{\mathrm{f}}(\mathbf{R},\mathbf{t}) = [2\tilde{\mathbf{C}}_{\mathrm{R}}\mathbf{r} + 2\mathbf{C}_{\mathrm{tR}}^T\mathbf{t} - 2\tilde{\mathbf{c}_R} \mid 2\mathbf{C}_{\mathrm{t}}\mathbf{t} + 2\mathbf{C}_{\mathrm{tR}}\mathbf{r} - 2\mathbf{c}_t]$; and (3) the Hessian-product vector is $\mathbf{H}_{\mathrm{f}}(\mathbf{R},\mathbf{t})[\mathbf{u}] = [2\mathbf{C}_{\mathrm{R}}\mathbf{u}_r + 2\mathbf{C}_{\mathrm{tR}}^T\mathbf{u}_t \mid 2\mathbf{C}_{\mathrm{t}}\mathbf{u}_t + 2\mathbf{C}_{\mathrm{tR}}\mathbf{u}_r]$. As initialization, we employ the solution from the DLT solver, see [17], as the eigenvector associated with the smallest eigenvalue of the coefficient matrix $\mathbf{C}$. For this solution to be unique, we require at least 6 points in correspondence.

IV. Optimality Certificate

The next development does not depend on the structure of the coefficient matrix $\mathbf{C}$, and it is general for any problem with quadratic cost on the rotation matrix.

A. Dual problem

Given a potential optimal solution for Problem QCQP, we aim to obtain a certificate of its optimality. An efficient approach can be built upon a convex relaxation of the problem, for example the dual problem [4], which is the one we use here since it has been shown to perform well [7]. Following the standard procedure, the dual problem of Problem QCQP reads:

$$d^\star = \min_{\boldsymbol{\lambda}\in\mathbb{R}^K,\rho\in\mathbb{R}} \rho \quad \text{subject to } \mathbf{C} - \sum_{i=1}^{K}\lambda_i\mathbf{A}_i - \rho\mathbf{L} \succeq 0 \tag{DUAL}$$

where λ_i and ρ are the Lagrange multipliers associated to the constraints $\mathbf{A}_i$ for $i = 1,\dots,K$ and $\mathbf{L}$, respectively. The problem in DUAL is a relaxation of the primal QCQP, meaning that the optimal cost $d^\star$ is not greater than the cost for the original $f^\star$ and the next inequality holds for all feasible points, i.e.,points that fulfill the constraints of the problems

$$d(\boldsymbol{\lambda},\rho) \leq d^\star \leq f^\star \leq f(\mathbf{x}). \tag{2}$$

Assuming strong duality holds, that is $d^\star = f^\star$, and that the provided solution $\mathbf{x}$ is the global optimum, we can derive the conditions that form the certifier (formal development is included in D). Two steps are required: (1) estimate the candidate to Lagrange multipliers; and (2) check whether this candidate is an actual dual feasible point. Since we assume strong duality, that is, $d^\star = f^\star$, from the primal (QCQP) and dual problem (DUAL) we have that $f^\star = \mathbf{x}^T\mathbf{C}\mathbf{x}$ and $d^\star = \rho$ and so $d^\star = \rho = \mathbf{x}^T\mathbf{C}\mathbf{x} = f^\star$. Therefore, we can eliminate ρ from both steps of the certifier by substituting it with the value $\mathbf{x}^T\mathbf{C}\mathbf{x}$. The first step relates the data of the problem $\mathbf{C}$, the primal solution $\mathbf{x}$ and the Lagrange multipliers $\boldsymbol{\lambda}(\mathbf{x})$ through the linear system:

$$\underbrace{\mathbf{J}(\mathbf{x})}_{10\times K}\boldsymbol{\lambda}(\mathbf{x}) = \underbrace{\mathbf{C}\mathbf{x}}_{10\times 1} \tag{3}$$

where $\mathbf{J}(\mathbf{x}) \in \mathbb{R}^{10\times K}$ is the Jacobian of the constraints evaluated at the primal point $\mathbf{x}$. Notice that the Jacobian does not depend on the number of points N, number of cameras M, nor the form of the cost $\mathbf{C}$, but only on the constraints $\mathbf{A}_i$ that define the space. Checking whether the solution $\boldsymbol{\lambda}(\mathbf{x})$ of (3) is a dual point requires evaluating the Hessian of the Lagrangian $\mathbf{H} \doteq \mathbf{C} - \sum_{i=1}^{K}\lambda_i\mathbf{A}_i - \rho\mathbf{L} \in \mathbb{S}^{10}$. For $\boldsymbol{\lambda}(\mathbf{x})$ to be feasible, the Hessian $\mathbf{H}$ should be positive semidefinite (PSD), that is, all its eigenvalues should be nonnegative (up to some tolerance to account for numerical errors). Notice that whereas the second step is usually the most time consuming one, the main limitations come from the first step and the conditions that the problem formulation imposes through the Jacobian.

Most of the certifiers proposed in previous works estimates the Lagrange multipliers in closed-form, which made them efficient [1]. In general, two conditions are required for the candidate to Lagrange multipliers to be estimated in closed-form: (1) the number of constraints is not larger than the number of variables, hence never making the system in Eq. (3) underdetermined; and (2) the Jacobian $\mathbf{J}(\mathbf{x})$ is full-rank for all feasible primal solutions. The last condition, known as Linear Independence Constraint Qualification (LICQ), holds generally. Nevertheless, we can find formulations in which both of these two conditions fail, even for the problem at hand. Notice that failure of either condition implies that, at first, there could exist a family of optimality certificates or none at all (the primal

[1]If the system becomes overdetermined, the solution is obtained in the least-squares sense, minimizing $\|\mathbf{J}(\mathbf{x})\boldsymbol{\lambda}(\mathbf{x}) - \mathbf{C}\mathbf{x}\|_2$.

solution is not the global optimum and/or strong duality does not hold, i.e.,, $d^\star < f^\star$ with strict inequality). LICQ is associated with the dimension of the space and the number of constraints required to fully define this space. If the goal is to provide a closed-form certifier, failure of this condition usually requires a change in the formulation, except when other options are available. A similar case arises when the number of constraints is larger than the number of variables. However, the latter case is common when redundant constraints are being added to the problem to improve performance, i.e.,tightness, since the certifier will succeed on those cases in which the convex relaxation is tight, provided the solution is the global optimum. Empirically, different formulations for the same exact problem will perform differently and remain tight, where other do not. Further, whereas minimal formulations provide efficient certifiers, redundant formulations empirically remain tight for more difficult problem instances, with low number of data and/or large noise, see e.g., [14]. Therefore, certifiers built upon redundant formulations will be able to certify more solutions, see e.g., [14], and our evaluation in Section V further supports this observation.

B. Constraints for the rotation group

Although redundant constraints are rare, rotation matrices are a well-known counterexample. As it has been shown previously in the literature, most problems that involve rotation matrices can be relaxed to orthogonal matrices, hence including reflections in the search space. The set is then relaxed from $\mathcal{SO}(3)$ to $\mathcal{O}(3)$ (the positive determinant constraint is dropped), and the set of constraints is defined as

$$\text{ROWS} \equiv \{\mathbf{R}\mathbf{R}^T = \mathbf{I}_3\}, \tag{4}$$

which provides 6 linear independent constraints. However, since the rotation matrix is squared, we can alternatively work with the columns instead of the rows as

$$\text{COLS} \equiv \{\mathbf{R}^T\mathbf{R} = \mathbf{I}_3\}, \tag{5}$$

which again provides 6 linearly independent constraints As ROWS, COLS only defines the orthogonal group. We enforce the positive determinant on the rotation matrix by including an additional set of constraints, denoted by DET. The determinant restriction can be expressed as quadratic functions, as shown in [33]. Equivalently, these constraints are obtained from the adjugate of the rotation since $\text{Adj}(\mathbf{R})\mathbf{R} = \text{Det}(\mathbf{R})\mathbf{I}_3$ and so $\text{Adj}(\mathbf{R}) = +\mathbf{R}^T$,

$$\text{DET} \equiv \{\text{Adj}(\mathbf{R}) = +\mathbf{R}^T\}, \tag{6}$$

which provides 9 linearly independent constraints. The explicit forms of these sets are available in E.

In this manuscript we derive and compare certifiers for the combinations of constraints: ROWS; COLS; ROWS + COLS (denoted by BOTH); and ROWS + COLS + DET (denoted by ALL), with 6, 6, 11 and 20 linear independent constraints, respectively[2]. ROWS and COLS are the only formulations that fulfill LICQ (empirically we observe the Jacobian $\mathbf{J}(\mathbf{x})$ has rank six for all the formulations, even for random rotation matrices), and therefore are the only ones that allow us to derive closed-form certifiers. Our goal, however, is to obtain certifiers also for BOTH and ALL. Before continue, notice that the formulations ROWS, COLS and BOTH define orthogonal matrices, hence including matrices with positive and negative determinant, making the problem under these constraints a relaxation of the original in (O) as the global optimum may not be rotation matrix. Although in the literature it has been shown that these formulations perform well for some range of problems, and empirically we observe a similar behavior in Section V, in this work we are not interested on estimating the solution to (O) from scratch with these formulations, but certifying a feasible solution that has been estimated by another algorithm, which is by construction a valid rotation with positive determinant. Hence, the effect of the relaxation of ROWS, COLS and BOTH appears if their global solution is not a rotation matrix, as the certification will fail for the given solution even if it is the optimum for the problem.

C. Iterative certifier

We run this problem with the above-mentioned formulations on the proposal in [14] and find that for random coefficient matrices $\mathbf{C}$: (1) the minimum-norm Lagrange multipliers for BOTH certify 75.96% of the solutions and for ALL only 9.96%; (2) the Hessian matrix with size 10×10 attains rank 6 in 1.81% of the certified solutions with BOTH and rank 7 in 0.69% with ALL for matrices with original size 12 (see Figure 2); and (3) the minimum-norm solution for ALL lays far from the optimum, making the iterative algorithm slow to be used in practice even for non-random coefficient matrices. Our first observation implies that a large number of problem instances are certified through a linear system for BOTH, despite the Jacobian being rank deficient. Our second observation implies the existence of

[2]The linear independence of the constraints can be verified by vectorizing the matrices and forming a matrix with the resulting vectors as columns. The singular Value Decomposition, or any rank revealing decomposition will show the dependent constraints. Notice that the constraints may be linearly independent even when the space they define is the same.

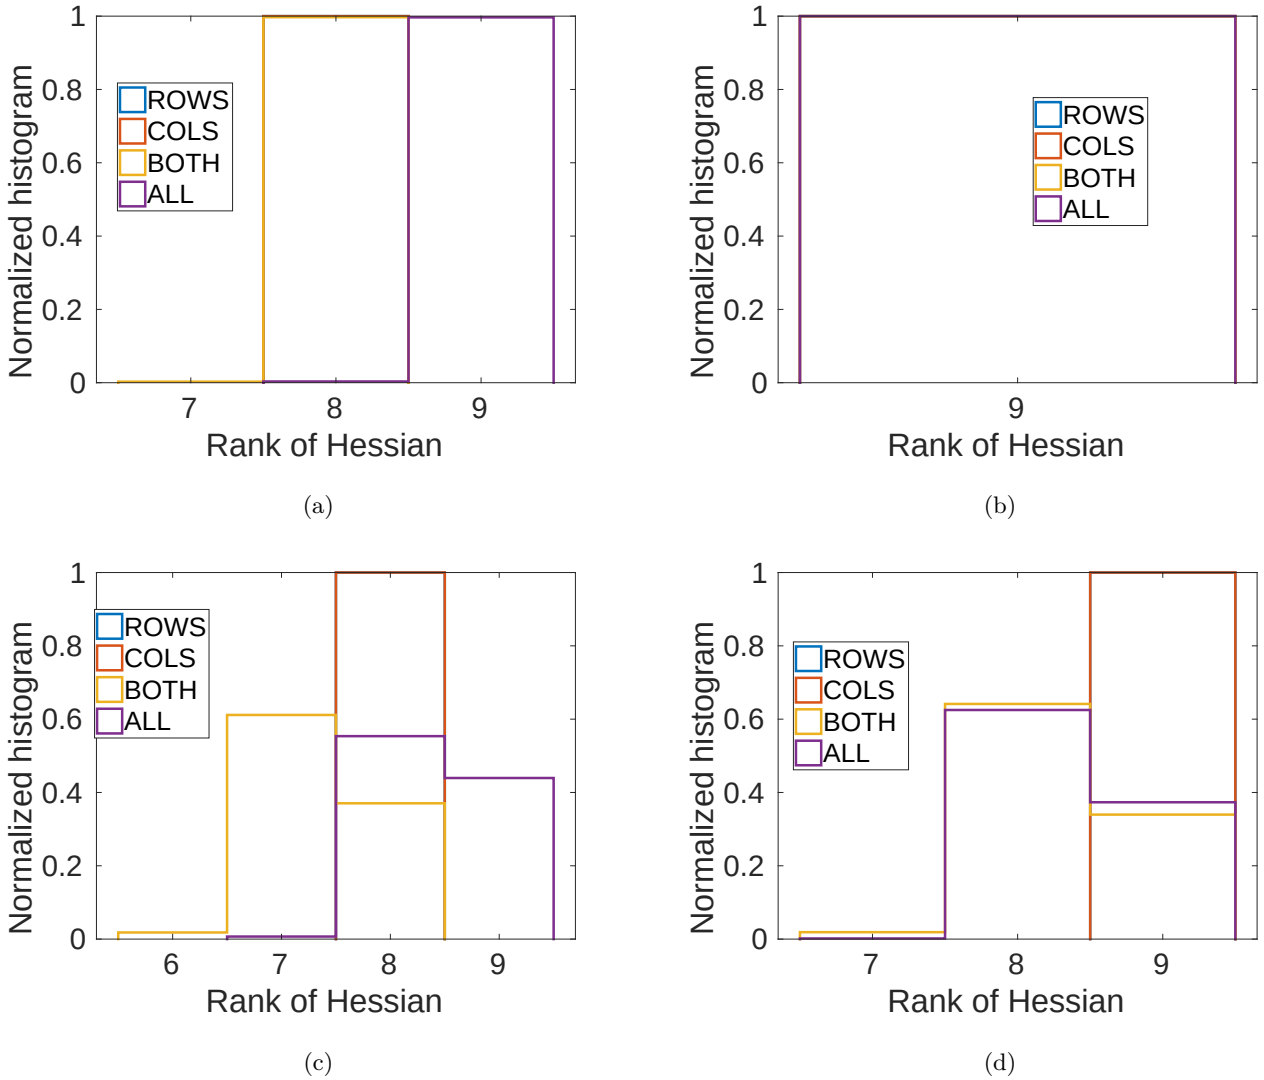

Fig. 2: Normalized histogram of the rank of the 10×10 Hessian for (left column) the central; and (right column) noncentral configurations with random coefficient matrices $\mathbf{C}$. First row shows results for the same problems considering the scale $s = 2$ and second row for $s = 1$.

more than one dual optimal solution for some problem instances, since the SDP solver returns rank-9 dual solutions. Our third and last observation motivates the proposed algorithm.

To ease notation, we drop the dependence of the multipliers $\boldsymbol{\lambda}(\mathbf{x})$ on the primal solution $\mathbf{x}$ and write it as $\boldsymbol{\lambda}$. Consider the minimum-norm solution $\boldsymbol{\lambda}^0 \in \mathbb{R}^K$ that minimizes $\|\mathbf{J}(\mathbf{x})\boldsymbol{\lambda}^0 - \mathbf{C}\mathbf{x}\|_2$ and let the right nullspace of the Jacobian be defined by $\mathbf{N} \in \mathbb{R}^{K\times(K-6)}$. The family of solutions to the system in (3) is $\boldsymbol{\lambda} = \boldsymbol{\lambda}^0 + \mathbf{N}\boldsymbol{\phi}$, with $\boldsymbol{\phi} \in \mathbb{R}^{K-6}$. We can write the Hessian in terms of this vector

$$\mathbf{H} = \mathbf{C} - \sum_{i=1}^{K} \lambda_i \mathbf{A}_i - \rho\mathbf{L} = \mathbf{C} - \rho\mathbf{L} - \sum_{i=1}^{K} (\boldsymbol{\lambda}_i^0 + \mathbf{e}_i^T \mathbf{N}\boldsymbol{\phi})\mathbf{A}_i \tag{7}$$

$$= \mathbf{C} - \rho\mathbf{L} - \sum_{i=1}^{K} \boldsymbol{\lambda}_i^0 \mathbf{A}_i - \sum_{j=1}^{K-6} \boldsymbol{\phi}_j \sum_{i=1}^{K} \mathbf{N}_{i,j}\mathbf{A}_i = \mathbf{H}^0 + \sum_{j=1}^{K-6} \boldsymbol{\phi}_i \mathbf{Z}_j, \tag{8}$$

where $\mathbf{H}^0 \doteq \mathbf{C} - \rho\mathbf{L} - \sum_{i=1}^{K} \boldsymbol{\lambda}_i^0 \mathbf{A}_i \in \mathbb{S}^{10}$, $\mathbf{e}_i \in \mathbb{R}^K$ is a vector of zeros except for its i-th entry equal to one and $\mathbf{Z}_j \doteq \sum_{j=1}^{K-6} \sum_{i=1}^{K} \mathbf{N}_{i,j}\mathbf{A}_i \in \mathbb{S}^{10}$. Our goal is to find the unconstrained $(K-6)$-D vector $\boldsymbol{\phi}$ such that the associated Hessian is PSD.

Instead of formulating the certification in its low-rank form as in [14], we opt for expressing the problem as an eigenvalue optimization, and thus we maximize the minimum eigenvalue of $\mathbf{H}$ as defined in (8). Note that this problem

is equivalent to minimizing the maximum eigenvalue of $-\mathbf{H}$, which has a vast literature, e.g., [20], [30] and references therein. Before continuing, we want to remark some aspects about this problem. First, the problem is convex since the minimum (or maximum) eigenvalue of the symmetric matrix $\mathbf{H}$ is a convex function and the Hessian is an affine function on $\boldsymbol{\phi}$ [4, Sec. 3.2], [21, Sec. 10.7], [24]. Second, the eigenvalues are known to coalesce and the multiplicity of the least eigenvalue would be, in general, larger than one. Third, it is a nonsmooth optimization in a neighborhood of the solution in general, unless the (unknown) multiplicity of the eigenvalue is one [25], [30]. The second point implies the existence of dual optimal solutions that makes the strict complementary condition fail. We empirically observe this behavior with the multiplicity of the smallest eigenvalue increasing up to four. Figure 2c and Figure 2d show the normalized histogram of the rank of the Hessian for each formulation under random coefficient matrices. We notice that this multiplicity depends also on the form of the cost matrix and constraints. We provide further details and comments in Section V-A. With these points in mind we write the certification problem as

$$g^\star = \max_{\boldsymbol{\phi}\in\mathbb{R}^{K-6}} \sum_{i=1}^{r} \mu_i\big(\mathbf{H}^0 + \sum_{j=1}^{K-6} \boldsymbol{\phi}_j \mathbf{Z}_j\big), \tag{9}$$

where $r \geq 1$ is the multiplicity of the least eigenvalue, not necessary one, and $\mu_k(\mathbf{A}) \in \mathbb{R}$ is the k-th eigenvalue of $\mathbf{A}$ (increasing order). For generic and symmetric $\mathbf{H}^0, \mathbf{Z}_j, j = 1, \ldots, K-6$, Problem (9) does not have a finite optimal cost. However, recall that these matrices are obtained from the original Hessian in eq (8) and the relation with the primal solution in eq (3), the latter derived from the assumption of strong duality and $\mathbf{H}\mathbf{x} = \mathbf{0}_{10\times 1}$. Problem (9) forces the least eigenvalue(s) to be nonnegative, yet the vector $\mathbf{x}$ is the eigenvector (one of them) associated to the zero eigenvalue, that is, $g^\star \leq 0$.

Algorithm 1: Gradient descent with Line-search based certifier

INPUT: Matrices $\mathbf{H}^0, \{\mathbf{Z}_j\}_{j=1}^{K-6}$.
OUTPUT: Optimal cost $g^\star$
$k \leftarrow 1$,
$\boldsymbol{\phi}^k \leftarrow \mathbf{0}_{K-6}$,
$\mathbf{H}^k \leftarrow \mathbf{H}^0$;
repeat
  $\mathbf{H}^{k+1} \leftarrow \mathbf{H}^k + \sum_{j=1}^{K-6} \boldsymbol{\phi}_j^k \mathbf{Z}_j$;
  Form Hessian
  $(\mathbf{U}, \mathbf{D}) \leftarrow \text{eigenDecom}(\mathbf{H}^{k+1})$;
  Nonpositive eigenvalues and eigenvectors
  $g_{k+1}^1 \leftarrow \min(\mathbf{D})$;
  Minimum eigenvalue
  $\partial g(\mathbf{U}) = [\text{tr}(\mathbf{U}^T\mathbf{Z}_1\mathbf{U}), \ldots, \text{tr}(\mathbf{U}^T\mathbf{Z}_{K-6}\mathbf{U})]^T$;
  Check convergence
  if $g_{k+1}^1 < \tau$ or $\|\partial g(\mathbf{U})\|_2 < \tau_{grad}$ then
    $g^\star \leftarrow \text{tr}(\mathbf{D})$;
    Solution is optimal
  else
    if $g_{k+1}^1 < g_k^1$ then
      $\mathbf{H}_{\text{psd}} \leftarrow \text{PSD}(\mathbf{H}^{k+1})$;
      Use PSD apprroximation
      $\boldsymbol{\phi}^{k+1} \leftarrow \text{linSolver}(\mathbf{H}_{\text{psd}}^{k+1}, \mathbf{Z}_1, \ldots, \mathbf{Z}_{K-6})$;
    else
      $\boldsymbol{\phi}^{k+1} \leftarrow s|\text{tr}(\mathbf{D})|/\|\partial g(\mathbf{U})\|_2^2 \; \partial g(\mathbf{U})$;
      Compute next step
    end if
  end if
until convergence

To solve the nonsmooth problem (9) we rely on a gradient descent algorithm with exact line-search, see e.g., [4, Sec. 9.3], similar to the one used in e.g., [12], that does not require the Hessian of the cost (nor its approximation). The line-search algorithm requires: (1) the search direction; and (2) the step size. Since the problem is nonsmooth, we work with the subdifferential instead of the gradient, which for Problem (9) has the form [30]

$$\partial g(\mathbf{U}) = \left\{\text{conv}(\mathbf{U}\mathbf{U}^T) \mid \mathbf{U} \in \mathcal{O}(10\times h), \mathbf{H}\mathbf{U} = \mathbf{0}_{10\times h}\right\}, \tag{10}$$

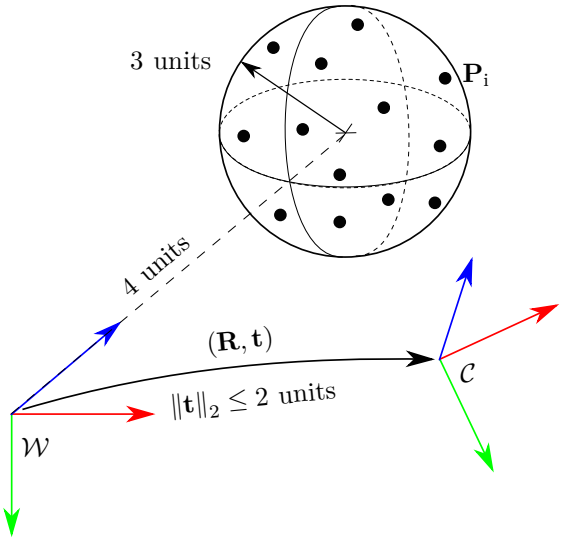

Fig. 3: Configuration of the points and cameras for the synthetic problem instances.

where $\mathcal{O}(10 \times h)$ is the space of $10 \times h$ orthogonal matrices and conv($\bullet$) is the convex hull of the parameter. As search direction we choose the subgradient of the cost: $\partial g(\mathbf{U}) = [\text{tr}(\mathbf{U}^T \mathbf{Z}_1 \mathbf{U}), \ldots, \text{tr}(\mathbf{U}^T \mathbf{Z}_{K-6} \mathbf{U})]^T \in \mathbb{R}^{K-6}$. Here we denote with $h \in \mathbb{R}$ the number of nonpositive eigenvalues of the current Hessian $\mathbf{H}$, which may be larger than one. If the certification is positive, at the end of the algorithm this number equals the multiplicity of the smallest eigenvalue of the Hessian. The step size is computed exactly as $g(\phi_{k-1} + \alpha_k \partial g(\mathbf{U}_{k-1})) = 0$, where ϕ_{k-1} is the previous solution with associated basis of eigenvectors $\mathbf{U}_{k-1}$ and eigenvalues in diagonal form $\mathbf{D}_{k-1} \in \mathbb{R}^{h \times h}$. The step size is given by $\alpha_k = |\text{tr}(\mathbf{D}_{k-1})| / \|\partial g(\mathbf{U}_{k-1})\|_2^2 \in \mathbb{R}_+$ which updates the vector as $\phi_k = |\text{tr}(\mathbf{D}_{k-1})| / \|\partial g(\mathbf{U}_{k-1})\|_2^2 \ \partial g(\mathbf{U}_{k-1})$. We then compute the eigendecomposition of the updated Hessian and repeat the process until convergence. Since the problem is convex, the norm of the subgradient tends to zero as the solution approaches the global optimum.

Convergence speed-up: In order to accelerate the convergence, we introduce a factor $s \in \mathbb{R}_+$ that scales the step size and leads to $s\phi$ being the next solution. We observe that a scale of $s = 2$ reduces the number of iterations almost by half for the random problem instances. However, this scale comes with a side effect that we depict in Figure 2. The first row shows the rank of the Hessian matrices (with size 10×10) for $s = 2$ and the second row for $s = 1$ for the same problem instances and primal solutions. First, the formulations without the determinant constraints and with the central-like coefficient matrix (Figure 2a) exhibit a block-diagonal structure with two blocks, and even the SDP solver returns Hessians with rank $10 - 2(= 8)$. We observe this behavior for ROWS and COLS which always attain rank-8 solutions. For the noncentral-like case (Figure 2b), where the cost is a 10×10 dense matrix, ROWS and COLS return rank-9 Hessians. Second, the scale $s = 2$ makes the ranks close to the maximum value for each case, except for some notable cases where we obtain matrices with rank 7. On the other hand, for $s = 1$ almost half of the problem instances end up with lower ranks for central- and noncentral-like problems, even reaching rank-6 solutions (see Figure 2c and d). To the best of our knowledge, this is the first time that two dual optimal solutions for the same problem have been reported in the literature.

In practice we let $s = 2$ to reduce the number of iterations and if the new solution does not improve the previous one, we estimate ϕ from the closest PSD matrix $\mathbf{H}_{\text{psd}}^{k+1}$ to the current Hessian $\mathbf{H}^{k+1}$ by zeroing the nonpositive eigenvalues. The vector ϕ is obtained from the linear system $\mathbf{G}\phi = \text{vec}(\mathbf{H}_{\text{psd}}^{k+1} - \mathbf{H}^0)$, where $\mathbf{G} \in \mathbb{R}^{55 \times (K-6)}$ has as columns the column-wise, symmetric vectorization of the matrices $\{\mathbf{Z}\}_{K-6}$. The system is overdetermined and we take the solution that minimizes the residual error. This correction calls for a linear solver in $m - 6$ variables and empirically it is required rarely (only for 0.5% of the iterations). Last, we define convergence when (a) the smallest eigenvalue normalized by the largest one g^1 is above the threshold $\tau = -1e-06$ to account for numerical errors; (b) the norm of the subgradient $\|\partial g(\mathbf{U})\|_2$ is below the positive threshold $\tau_{grad} = 1e-09$; or (c) the maximum number of iterations (50) is reached. Algorithm 1 summarizes the iterative certifier including the scale for the step size s in step 17.

	$\mathbf{12 \times 12}$		$\mathbf{13 \times 13}$	
	DLT (332)	SDP (600)	DLT (283)	SDP (600)
ROWS	23	25	7	9
COLS	31	32	6	10
BOTH	233	359	64	90
ALL	332	598	283	599

TABLE I: Number of certified solutions for each formulation (left column) for different sizes of random coefficient matrices (top row). We initialize the on-manifold refinement with the DLT and the SDP, and the numbers in brackets indicate the number of solutions with dual gap below $1e-07$ wrt ALL.

	BOTH		ALL	
	Time (μs)	Iterations	Time (μs)	Iterations
12×12	83.71 (61)	3.19 (2)	147.7 (99)	4.79 (3)
13×13	76.92 (63)	2.467 (2)	103.15 (83.5)	3.1 (2)
central	49.29 (42)	0.488 (0)	90.69 (79)	1.72 (1)
M = 2	47.25 (39.5)	0.46 (0)	79.35 (70)	1.469 (1)
M = 3	49.27 (41)	0.39(0)	81.78 (71)	1.379 (1)

TABLE II: Computational time as mean (median) in μs and number of iterations for the redundant formulations for random coefficient matrices (first and second rows) and actual resectioning problems with central and noncentral cameras (third to last rows).

V. Evaluation

A. Evaluation on random data

Motivated by [7], we briefly include here the performance of the certifier for random data, that is, problems instances without a specific structure. We generate data matrices $\mathbf{C}$ with entries in the range $[-1; 1]$, which are then made PSD by multiplying them with their transpose. We consider matrices with size 12×12 for the central case and 13×13 for the noncentral configuration, and obtain the 10×10 coefficient matrices by taking the Schur's complement [4, App. C.4]. We run the on-manifold optimization, initialized with (a) the DLT; and (b) the output from ALL. The solution from the iterative solver is then certified with the proposed algorithm and we consider the certification positive if the ratio betwen the minimum and maximum eigenvalues of the Hessian is above $-1e-06$. Table I reports the number of certified solutions for each case. We notice that the certified solutions by ROWS and COLS were not the same, whereas the noncertified solutions by ALL reached minimum eigenvalues circa $-1e-04$. In addition, Table II reports the computational times and number of iterations required by BOTH and ALL for these problem instances. The minimal ROWS and COLS require $30\mu s$ per problem instance, independently of the size of the coefficient matrix. For comparison, the SDP solver requires an averaged of 4ms (median 2ms) to solve the same problems using sdpa [36]. Last, Figure 2 shows the normalized histogram of the rank of the 10×10 Hessian for certified solutions and scale values $s = 1, 2$.

B. Evaluation on synthetic data

We generate a random set of N points within a ball of radius 3 units at a distance of 4 units along the Z-axis wrt the world frame. We generate the camera's pose relative to the world frame with an angle of rotation below 0.5 rad and random translation with norm bounded above by 2 units. A scheme of this setup is shown in Figure 3. We perform additional experiments on planar and linear motions, but did not observe any substantial difference, and thus we provide only the results for the general motion. To generate the observations we assume a perspective camera,

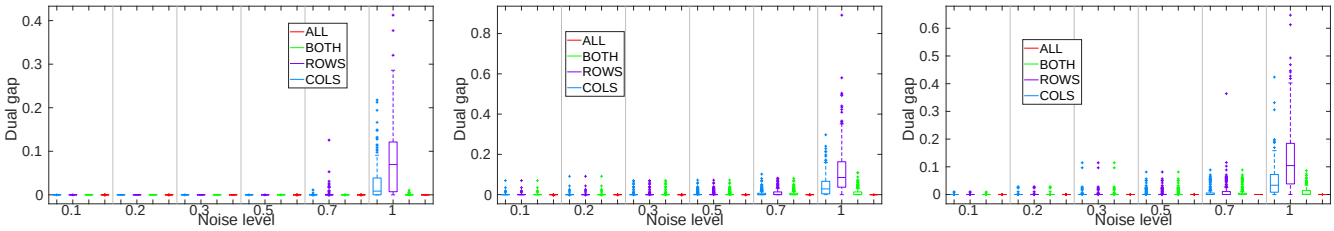

(a) (b) (c)

Fig. 4: Relative dual gap for the different formulations (see legend) on problem instances with $N = 50$ and noise σ (X-axis) for (a) the central case; (b) the noncentral case with $M = 2$; and (c) for $M = 3$ cameras.

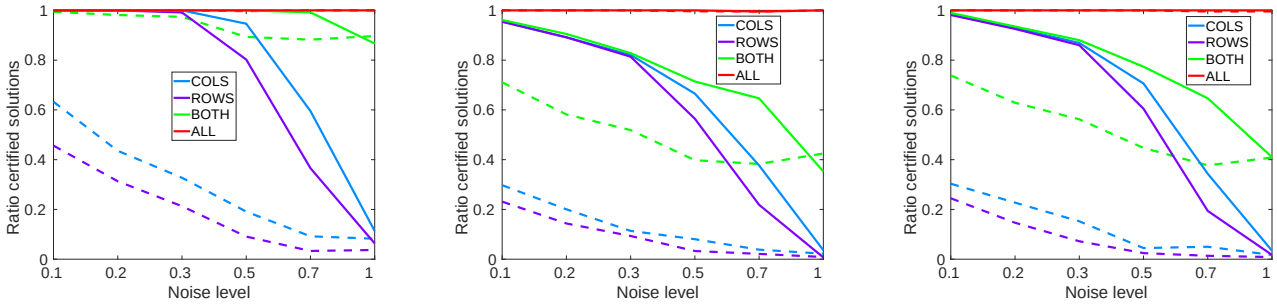

(a) (b) (c)

Fig. 5: Ratio of certified solutions for $N = 6$ (dashed lines) and $N = 20$ (solid lines) for the different formulations (legend) and increasing level of noise (X-axis) for (a) the central case; (b) noncentral with $M = 2$; and (c) with $M = 3$ cameras.

normalize the observations $\mathbf{f}_i$ to have unit norm and introduce noise by drawing a 2D vector from a normal distribution on the tangent plane of the observation with $\ell_\infty = \sigma$. The noise is applied through the exponential map for the $\mathcal{S}^2$. This makes the angle of deviation in rads wrt the noiseless observation equal to the ℓ_2 of the noise vector. We consider problem instances with number of points $N \in \{6, 7, 8, 15, 20, 50, 100\}$ and noise levels $\sigma \in \{0.1, 0.2, 0.3, 0.5, 0.7, 1.0\}$. For the noncentral configuration, we consider rigs formed by $M = 2, 3$ cameras whose relative pose wrt the frame of the rig has random rotation (upper bound of 1 rad) and translation (norm bounded above by 1 unit). For each combination of noise-number of correspondences, we generate 600 random problem instances. We are mainly interested in the performance of the certifier, which depends on: (1) the tightness of the formulations; and (2) the accuracy of the on-manifold refinement. The latter is dependent also on the initial guess and can be substituted by any iterative method and thus we reduce our attention here to the tightness of the relaxation.

Tightness for the different formulations: We measure the quality of the formulation by the dual gap, i.e.,the difference between the dual cost and the cost attained by the closest primal feasible solution to the solution of the SDP. Figure 4 shows this metric for: 4a the central case; 4b the noncentral case with $M = 2$; and for 4c for $M = 3$ cameras. We plot the results for problem instances with $N = 50$ and increasing level of noise (X-axis). For the depicted problem instances ALL has dual gap close to zero, whereas BOTH fails for some cases, mostly with large noise and low number of correspondences. On the other hand, the minimal formulations ROWS and COLS perform well for problem instances with low noise and $N > 50$. These cases are common in real-world applications, therefore our interest on these formulations. For more challenging scenarios with lower number of correspondences and/or highly noisy data, the redundant formulations are preferred. With respect to computational time, all the formulations range form 1 to 1.5 milliseconds and for some problems they require up to 4 milliseconds.

Performance of the certifiers: Figure 5 shows the ratio of certified solutions for each formulation (see legend) for problem instances with $N = 6$ (dashed lines) and $N = 20$ (solid lines) for increasing level of noise (X-axis) for: 5a the central case; 5b the noncentral case with $M = 2$; and 5c for $M = 3$ cameras. We consider the certification positive if the ratio between the minimum and maximum eigenvalue of the Hessian is larger than $-1e-06$. The behavior is similar to the one in Figure 4 and shows that: (1) the minimal formulations ROWS and COLS can certify some of the solutions (more than 80% for all configurations with more than $N = 20$ correspondences and noise under 0.3); while (2) the redundant formulations BOTH and ALL perform better for complicated problem instances with fewer points and larger noise. Computational time and number of iterations are reported in Table II, third to last rows. Finally,

	central				$M=2$			
	ROWS	COLS	BOTH	ALL	ROWS	COLS	BOTH	ALL
V1-E	0.96	0.97	1.00	1.00	0.98	0.97	1.00	1.00
V1-M	0.94	0.96	1.00	1.00	0.97	0.98	1.00	1.00
V1-D	0.93	0.97	1.00	1.00	0.95	0.99	1.00	1.00
V2-E	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00
V2-M	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
V2-D	0.98	0.98	1.00	1.00	0.99	0.98	1.00	1.00
delivery-45	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
electro-45	0.98	0.94	1.00	1.00	1.00	0.94	1.00	1.00
forest-45	0.99	0.95	0.99	1.00	0.98	0.94	1.00	1.00
playground-45	0.99	0.97	1.00	1.00	0.98	0.98	1.00	1.00
terrains-45	1.00	0.99	1.00	1.00	1.00	1.00	1.00	1.00
delivery-67	1.00	0.99	1.00	1.00	1.00	1.00	1.00	1.00
electro-67	0.98	0.92	1.00	1.00	0.98	0.94	1.00	1.00
forest-67	1.0	0.95	1.00	1.00	1.00	0.96	1.00	1.00
playground-67	0.98	0.94	1.00	1.00	0.98	0.93	1.00	1.00
terrains-67	0.97	0.98	1.00	1.00	1.00	0.99	1.00	1.00

TABLE III: Ratio of certified solutions for each sequence (left column) for different camera configurations.

the rank of the final Hessian follows the tendency shown in Figure 2 for $s = 2$.

C. Evaluation on real data

We evaluate our proposal in two datasets: (a) EUROC [8] with the sequences Vicon Room 1: 01, 02, 03 and Vicon Room 2: 01, 02, 03, denoted by V1-easy, V1-med and V1-diff and V2-easy, V2-med and V2-diff; and (b) ETH3D [27] with sequences: delivery area, electro, forest, playground and terrains, denoted by delivery-45, delivery-67, electro-45, etc. These sequences provide ground-truth poses and have one and two stereo cameras with known baselines, respectively.

We obtain correspondences with the next procedure. First we extract and match SURF features [3] between a pair of stereo images (same timestamp) and then triangulate the 3D points considering the ground-truth baseline between the cameras. An additional image is selected such that the translation norm between its pose and the one associated with the point cloud is larger than 0.05 m. We then extract and match SURF features between this image and the two that originated the point cloud, filtering those correspondences with the provided ground-truth pose, as we do not consider outliers nor their detection during the optimization. We keep the image iff there are at least 20 correspondences with at least one of the images. For the central case, we run our algorithm considering only one image and for the non-central we consider the matches in both stereo images. We run the on-manifold optimization and then try to certify the returned solution. We also estimate the solution with the SDP and ALL formulation using sdpa as solver. As in previous section, we consider the solution is the optimum if the dual gap is below $1e-06$ and the ratio of certified solutions that we report in this section considers only these solutions. In total EUROC provides with 1855 problem instances for central (with 1854 optimal solutions) and noncentral configurations (with 1854 optimal solutions) and ETH3D with 2264 for central (with 2249) and 2314 for noncentral (with 2300). All the sequences have more than 120 problems instances and the maximum was 582.

Table III reports the ratio of certified solutions for each sequence for the central and noncentral configurations. More than 90% of the solutions are certified by all formulations, whereas only the redundant ALL certifies all problem instances for central and noncentral configurations. In terms of number of iterations, BOTH requires 0.0149 (median 0) for central and 0.0119 (median 0) for noncentral, whereas ALL goes to 0.5358 (median 0.7) for central and 0.0204 (median 0) for noncentral. We observe a similar tendency for the computational time. The minimal ROWS and COLS require less than $25\mu s$ per problem instance; BOTH requires $41\mu s$ for central and $38\mu s$ for noncentral; and ALL goes to $68\mu s$ for central and $44\mu s$ for noncentral. Mean and median values differ in less than $5\mu s$.

VI. Conclusion and future work

In this work we proposed a certifiable solver for the resectioning problem (PnP) for central and noncentral cameras. The proposal first estimated the solution to the PnP problem via an iterative solver that cannot guarantee the optimality of the solution, and then tried to certify if this solution was the global optimum. Empirically, the certification was shown to detect the optimality of a subset of solutions, but it failed for some problem instances.

We stated the resectioning problem as the minimization of the point-line distance between the 3D points and the observations (as bearing vectors). The resectioning problem was nonconvex due to the rotation variable, and thus we relied on a convex relaxation (the dual problem) for the certification. We exploited different formulations with increasing number of constraints to assess their effect on the certification process. Out of the four formulations considered, two of them lacked LICQ and one of them had more constraints than variables. Since the Lagrange multipliers could not be computed in closed-form for these cases, we re-wrote the certification problem as an eigenvalue optimization. This problem was convex but nonsmooth and we leveraged a line-search algorithm to solve it efficiently. Despite being an iterative algorithm, the certification as here stated was solved in microseconds for both synthetic (including random) and real data. Minimal formulations were able to detect most optimal solutions whereas the redundant formulations certified even those solutions with low number of correspondences and highly noisy data. On real data, all formulations certified more than 90% of the solutions, but only the fully redundant formulation ALL certified all of them.

As future work, we envision the extension of this type of certifiers to other problems involving rotations, for example, registration problems involving points and/or planes.

Acknowledgments

We thank Dr. Volodya Grancharov (Ericsson Research) for his support during this project. This work has been supported by the grant program 600 FPU18/01526 and the research projects ARPEGGIO (PID2020-117057) and HOUNDBOT (UMA20-FEDERJA- 056) funded by the Spanish and Andalusian Governments, respectively, and European Regional Development Fund 605 (ERDF).

References

[1] P.-A. Absil, R. Mahony, and R. Sepulchre, Optimization algorithms on matrix manifolds, Princeton University Press, 2009.
[2] S. Agostinho, J. Gomes, and A. Del Bue, Cvxpnpl: A unified convex solution to the absolute pose estimation problem from point and line correspondences, Journal of Mathematical Imaging and Vision, (2022), pp. 1–21.
[3] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, Speeded-up robust features (surf), Computer Vision and Image Understanding (CVIU), 110 (2008), pp. 346–359.
[4] S. Boyd, S. P. Boyd, and L. Vandenberghe, Convex optimization, Cambridge university press, 2004.
[5] J. Briales and J. Gonzalez-Jimenez, Cartan-sync: Fast and global se (d)-synchronization, IEEE Robotis and Automation Letters (RA-L), 2 (2017), pp. 2127–2134.
[6] J. Briales and J. Gonzalez-Jimenez, Convex global 3d registration with lagrangian duality, in IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4960–4969.
[7] L. Brynte, V. Larsson, J. P. Iglesias, C. Olsson, and F. Kahl, On the tightness of semidefinite relaxations for rotation estimation, Journal of Mathematical Imaging and Vision, (2022), pp. 1–11.
[8] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, The euroc micro aerial vehicle datasets, The International Journal of Robotics Research (IJRR), (2016), https://doi.org/10.1177/0278364915620033.
[9] M. Byröd, K. Josephson, and K. Åström, Fast and stable polynomial equation solving and its application to computer vision, Int'l J. Computer Vision (IJCV), 84 (2009), pp. 237–256.
[10] J. Campos, J. R. Cardoso, and P. Miraldo, Poseamm: A unified framework for solving pose problems using an alternating minimization method, in IEEE Int'l Conf. Robotics and Automation (ICRA), 2019, pp. 3493–3499.
[11] A. Eriksson, C. Olsson, F. Kahl, and T.-J. Chin, Rotation averaging and strong duality, in IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2018, pp. 127–135.
[12] S. Friedland, J. Nocedal, and M. L. Overton, The formulation and analysis of numerical methods for inverse eigenvalue problems, SIAM Journal on Numerical Analysis, 24 (1987), pp. 634–667.
[13] M. Garcia-Salguero, J. Briales, and J. Gonzalez-Jimenez, A tighter relaxation for the relative pose problem between cameras, Journal of Mathematical Imaging and Vision, 64 (2022), pp. 493–505.
[14] M. Garcia-Salguero and J. Gonzalez-Jimenez, Fast certifiable relative pose estimation with gravity prior, Artificial Intelligence, (2023), https://doi.org/https://doi.org/10.1016/j.artint.2023.103862.
[15] C. Geyer and K. Daniilidis, A unifying theory for central panoramic systems and practical implications, in European Conf. Computer Vision (ECCV), 2000, pp. 445–461.
[16] M. D. Grossberg and S. K. Nayar, A general imaging model and a method for finding its parameters, in IEEE Int'l Conf. Computer Vision (ICCV), vol. 2, 2001, pp. 108–115.
[17] R. Hartley and A. Zisserman, Multiple view geometry in computer vision, Cambridge university press, 2003.
[18] J. A. Hesch and S. I. Roumeliotis, A direct least-squares (dls) method for pnp, in IEEE Int'l Conf. Computer Vision (ICCV), 2011, pp. 383–390.
[19] J. B. Lasserre, Global optimization with polynomials and the problem of moments, SIAM Journal on optimization, 11 (2001), pp. 796–817.
[20] A. Lewis and C. Wylie, A simple newton method for local nonsmooth optimization, arXiv preprint arXiv:1907.11742, (2019).
[21] A. S. Lewis and M. L. Overton, Eigenvalue optimization, Acta numerica, 5 (1996), pp. 149–190.
[22] E. Malis, Complete closed-form and accurate solution to pose estimation from 3d correspondences, IEEE Robotis and Automation Letters (RA-L), 8 (2023), pp. 1786–1793.
[23] C. Olsson and A. Eriksson, Solving quadratically constrained geometrical problems using lagrangian duality, in IEEE Int'l Conf. Pattern Recognition (ICPR), IEEE, 2008, pp. 1–5.

[24] M. L. Overton, On minimizing the maximum eigenvalue of a symmetric matrix, SIAM Journal on Matrix Analysis and Applications, 9 (1988), pp. 256–268.
[25] M. L. Overton and R. S. Womersley, Second derivatives for optimizing eigenvalues of symmetric matrices, SIAM Journal on Matrix Analysis and Applications, 16 (1995), pp. 697–718.
[26] D. M. Rosen, L. Carlone, A. S. Bandeira, and J. J. Leonard, Se-sync: A certifiably correct algorithm for synchronization over the special euclidean group, The International Journal of Robotics Research (IJRR), 38 (2019), pp. 95–125.
[27] T. Schöps, J. L. Schönberger, S. Galliani, T. Sattler, K. Schindler, M. Pollefeys, and A. Geiger, A multi-view stereo benchmark with high-resolution images and multi-camera videos, in IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017.
[28] G. Schweighofer and A. Pinz, Robust pose estimation from a planar target, IEEE Trans. Pattern Analysis and Machine Intelligence (T-PAMI), 28 (2006), pp. 2024–2030.
[29] G. Schweighofer and A. Pinz, Globally optimal o (n) solution to the pnp problem for general camera models., in British Machine Vision Conference (BMVC), 2008, pp. 1–10.
[30] A. Shapiro and M. K. Fan, On eigenvalue optimization, SIAM Journal on Optimization, 5 (1995), pp. 552–569.
[31] L. Sun and Z. Deng, Certifiably optimal and robust camera pose estimation from points and lines, IEEE Access, 8 (2020), pp. 124032–124054.
[32] C. Sweeney, V. Fragoso, T. Höllerer, and M. Turk, gdls: A scalable solution to the generalized pose and scale problem, in European Conf. Computer Vision (ECCV), 2014, pp. 16–31.
[33] R. Tron, D. M. Rosen, and L. Carlone, On the inclusion of determinant constraints in lagrangian duality for 3d slam, in Robotics: Science and Systems (RSS), Workshop "The problem of mobile sensors: Setting future goals and indicators of progress for SLAM, vol. 4, 2015.
[34] F. Wientapper, M. Schmitt, M. Fraissinet-Tachet, and A. Kuijper, A universal, closed-form approach for absolute pose problems, Computer Vision and Image Understanding (CVIU), 173 (2018), pp. 57–75.
[35] J. Wu, Y. Zheng, Z. Gao, Y. Jiang, X. Hu, Y. Zhu, J. Jiao, and M. Liu, Quadratic pose estimation problems: Globally optimal solutions, solvability/observability analysis, and uncertainty description, IEEE Trans. Robotics (T-RO), 38 (2022), pp. 3314–3335.
[36] M. Yamashita, K. Fujisawa, M. Fukuda, K. Kobayashi, K. Nakata, and M. Nakata, Latest developments in the sdpa family for solving large-scale sdps, in Handbook on semidefinite, conic and polynomial optimization, Springer, 2012, pp. 687–713.
[37] Y. Zheng, Y. Kuang, S. Sugimoto, K. Astrom, and M. Okutomi, Revisiting the pnp problem: A fast, general and optimal solution, in IEEE Int'l Conf. Computer Vision (ICCV), 2013, pp. 2344–2351.
[38] L. Zhou and M. Kaess, An efficient and accurate algorithm for the perspecitve-n-point problem, in IEEE/RSJ Int'l Conf. Intelligent Robots and Systems (IROS), 2019, pp. 6245–6252.
[39] L. Zhou, S. Wang, and M. Kaess, A fast and accurate solution for pose estimation from 3d correspondences, in IEEE Int'l Conf. Robotics and Automation (ICRA), 2020, pp. 1308–1314.

Appendix A
Matrix form of the cost function

We provide here the cost in its matricial form and use the non-central configuration as it applies for the central case for $M = 1$. We remind the reader that this re-formulation has been published by previous works, see Section II and it is only included here for completeness.

For noncentral configurations, the feature is observed by a camera with a relative pose $\mathbf{R}_\mathrm{j}, \mathbf{c}_\mathrm{j}$ wrt the 'generalized' camera center. Thus, for the absolute rotation $\mathbf{R}$ and translation $\mathbf{t}$ and considering the point is observed in the j-th camera we have that the transformed point is $\mathbf{R}_\mathrm{j}^T(\mathbf{R}\mathbf{P}_\mathrm{i} + \mathbf{t} - \mathbf{c}_\mathrm{j})$. The observation is given by $\mathbf{f}_\mathrm{i}$ and we minimize the point-line distance between the transformed point and the unit-norm observation. To ease notation, we pre-multiply the observation $\mathbf{f}_\mathrm{i}$ with the rotation $\mathbf{R}_\mathrm{j}$, and denoted it also by $\mathbf{f}_\mathrm{i}$. Thus, the error is given by $(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)(\mathbf{R}\mathbf{P}_\mathrm{i} + \mathbf{t} - \mathbf{c}_\mathrm{j})$ and the total cost is the sum over $i = 1, \ldots, N$ of the terms

$$f_i(\mathbf{R}, \mathbf{t}) = \|(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)(\mathbf{R}\mathbf{P}_\mathrm{i} + \mathbf{t} - \mathbf{c}_\mathrm{j})\|_2^2. \tag{11}$$

The term $\mathbf{R}\mathbf{P}_\mathrm{i} + \mathbf{t} - \mathbf{c}_\mathrm{j} \in \mathbb{R}^3$ is expressed in vector form

$$\mathbf{R}\mathbf{P}_\mathrm{i} + \mathbf{t} - \mathbf{c}_\mathrm{j} = (\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)\mathbf{r} + \mathbf{t} - \mathbf{c}_\mathrm{j}, \tag{12}$$

where $\otimes$ is the Kronecker product. Developing the norm, we have

$$f_i(\mathbf{R}, \mathbf{t}) = \|(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)\big((\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)\mathbf{r} + \mathbf{t} - \mathbf{c}_\mathrm{j}\big)\|_2^2 \tag{13}$$

$$= \mathbf{r}^T \underbrace{(\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)^T(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)(\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)}_{9\times 9}\mathbf{r} \tag{14}$$

$$+ \mathbf{t}^T \underbrace{(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)}_{3\times 3}\mathbf{t} + \underbrace{\mathbf{c}_\mathrm{j}^T(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)\mathbf{c}_\mathrm{j}}_{1\times 1} \tag{15}$$

$$+ 2\mathbf{t}^T \underbrace{(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)(\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)}_{3\times 9}\mathbf{r} \tag{16}$$

$$- 2\underbrace{\mathbf{c}_\mathrm{j}^T(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)}_{1\times 3}\mathbf{t} - 2\underbrace{\mathbf{c}_\mathrm{j}^T(\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)(\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)}_{1\times 9}\mathbf{r} \tag{17}$$

with the total cost being

$$f(\mathbf{R},\mathbf{t}) = \mathbf{r}^T \underbrace{\sum_{i=1}^{N} (\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)^T (\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)(\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)}_{\mathbf{C}_\mathrm{R}} \mathbf{r} \tag{18}$$

$$+ \mathbf{t}^T \underbrace{\sum_{i=1}^{N} (\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)}_{\mathbf{C}_\mathrm{t}} \mathbf{t} + \underbrace{\sum_{i=1}^{N} \mathbf{c}_\mathrm{j}^T (\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)\mathbf{c}_\mathrm{j}}_{c} \tag{19}$$

$$+ 2\mathbf{t}^T \underbrace{\sum_{i=1}^{N} (\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)(\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)}_{\mathbf{C}_\mathrm{tR}} \mathbf{r} \tag{20}$$

$$- 2\underbrace{\sum_{i=1}^{N} \mathbf{c}_\mathrm{j}^T (\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)}_{\mathbf{c}_t^T} \mathbf{t} - 2\underbrace{\sum_{i=1}^{N} \mathbf{c}_\mathrm{j}^T (\mathbf{I}_3 - \mathbf{f}_\mathrm{i}\mathbf{f}_\mathrm{i}^T)(\mathbf{P}_\mathrm{i}^T \otimes \mathbf{I}_3)}_{\mathbf{c}_R^T} \mathbf{r}, \tag{21}$$

Stacking the unknowns $\mathbf{r}, \mathbf{t}$ and $y = 1$ into the vector $\mathbf{x} \in \mathbb{R}^{13}$ we obtain the matricial form $\mathbf{G}$:

$$\mathbf{G} \doteq \begin{pmatrix} \mathbf{C}_\mathrm{R} & \mathbf{C}_\mathrm{tR}^T & -\mathbf{c}_R \\ \mathbf{C}_\mathrm{tR} & \mathbf{C}_\mathrm{t} & -\mathbf{c}_t \\ -\mathbf{c}_R^T & -\mathbf{c}_t^T & c \end{pmatrix} \in \mathbb{S}_+^{13}. \tag{22}$$

For the central configuration ($M = 1$), we have that $c = 0, \mathbf{c}_R = \mathbf{0}_{9\times 1}$ and $\mathbf{c}_t = \mathbf{0}_{3\times 1}$.

Appendix B
Marginalization of the translation

We consider the noncentral configuration and provide at the end the form for the translation for $M = 1$ (central case). Notice that the latter can be found in the literature, see Section II.

Recall that the cost has the form $f(\mathbf{R},\mathbf{t}) = \mathbf{r}^T\tilde{\mathbf{C}}_\mathrm{R}\mathbf{r} + 2\mathbf{t}^T\mathbf{C}_\mathrm{tR}\mathbf{r} + \mathbf{t}^T\mathbf{C}_\mathrm{t}\mathbf{t} - 2\mathbf{c}_t^T\mathbf{t} - 2\tilde{\mathbf{c}_R}^T\mathbf{r} + \tilde{c}$. Since $\mathbf{t}$ is unconstrained, we can find its optimal value by considering the (partial) derivative wrt $\mathbf{t}$ as $\mathbf{G}_\mathrm{f}(\mathbf{t}) = 2\mathbf{C}_\mathrm{t}\mathbf{t} + 2\mathbf{C}_\mathrm{tR}\mathbf{r} - 2\mathbf{c}_t$ and making it equal to zero, obtaining $\mathbf{t}^\star = {\mathbf{C}_\mathrm{t}}^{-1}(\mathbf{c}_t - \mathbf{C}_\mathrm{tR}\mathbf{r})$. Substituting this result into the original cost

$$f(\mathbf{R},\mathbf{t}) = g(\mathbf{R}) = \mathbf{r}^T\tilde{\mathbf{C}}_\mathrm{R}\mathbf{r} + 2\mathbf{r}^T\mathbf{C}_\mathrm{tR}^T({\mathbf{C}_\mathrm{t}}^{-1}(\mathbf{c}_t - \mathbf{C}_\mathrm{tR}\mathbf{r})) + \tag{23}$$

$$({\mathbf{C}_\mathrm{t}}^{-1}(\mathbf{c}_t - \mathbf{C}_\mathrm{tR}\mathbf{r}))^T\mathbf{C}_\mathrm{t}({\mathbf{C}_\mathrm{t}}^{-1}(\mathbf{c}_t - \mathbf{C}_\mathrm{tR}\mathbf{r})) \tag{24}$$

$$- 2\mathbf{c}_t^T({\mathbf{C}_\mathrm{t}}^{-1}(\mathbf{c}_t - \mathbf{C}_\mathrm{tR}\mathbf{r})) - 2\tilde{\mathbf{c}_R}^T\mathbf{r} + \tilde{c} = \tag{25}$$

$$= \mathbf{r}^T \underbrace{\left(\tilde{\mathbf{C}}_\mathrm{R} - \mathbf{C}_\mathrm{tR}^T{\mathbf{C}_\mathrm{t}}^{-1}\mathbf{C}_\mathrm{tR}\right)}_{\mathbf{C}_\mathrm{R}}\mathbf{r} \tag{26}$$

$$+ 2\underbrace{\left(\mathbf{C}_\mathrm{tR}^T{\mathbf{C}_\mathrm{t}}^{-1}\mathbf{c}_t - \tilde{\mathbf{c}_R}\right)}_{\mathbf{c}_R}\mathbf{r} + \underbrace{\left(\tilde{c} - \mathbf{c}_t^T{\mathbf{C}_\mathrm{t}}^{-1}\mathbf{c}_t\right)}_{c} \tag{27}$$

For $M = 1$ we have that $\mathbf{R}_\mathrm{j} = \mathbf{I}_3$ and $\mathbf{c}_\mathrm{j} = \mathbf{0}_{3\times 1}$, and so $\mathbf{t}^\star = -{\mathbf{C}_\mathrm{t}}^{-1}\mathbf{C}_\mathrm{tR}\mathbf{r}$. The cost matrices have the form $\mathbf{C}_\mathrm{R} = \tilde{\mathbf{C}}_\mathrm{R} - \mathbf{C}_\mathrm{tR}^T{\mathbf{C}_\mathrm{t}}^{-1}\mathbf{C}_\mathrm{tR}$, $\mathbf{c}_R = \mathbf{0}_{9\times 1}$ and $c = 0$.

Appendix C
General dual problem for QCQP

In this Section we provide the dual problem for the generic QCQP. We follow the standard procedure that can be found e.g.,in [4]. Assuming the dual problem has a finite solution,

$$d^\star = \max_{\rho,\boldsymbol{\lambda}} \min_{\mathbf{x}} \mathcal{L}(\rho, \mathbf{x}, \boldsymbol{\lambda}) \tag{28}$$

where $\boldsymbol{\lambda}, \rho$ are the Lagrange multipliers and $\mathcal{L}(\rho, \mathbf{x}, \boldsymbol{\lambda})$ is known as the Lagrangian,

$$\mathcal{L}(\rho, \mathbf{x}, \boldsymbol{\lambda}) = \mathbf{x}\mathbf{C}\mathbf{x} + \rho - \rho\mathbf{x}^T\mathbf{L}\mathbf{x} - \sum_{i=1}^{K} \mathbf{x}^T\lambda_i\mathbf{A}_i\mathbf{x} = \tag{29}$$

$$= \mathbf{x}^T\left(\mathbf{C} - \rho\mathbf{L} - \sum_{i=1}^{K} \lambda_i\mathbf{A}_i\right)\mathbf{x} + \rho. \tag{30}$$

The minimum of the Lagrangian is finite iff the symmetric matrix $\mathbf{C}-\rho\mathbf{L}-\sum_{i=1}^{K}\lambda_i\mathbf{A}_i$, the so-called Hessian of the Lagrangian, is positive semidefinite (PSD), i.e.,all its eigenvalues are non-negative. Since we only consider finite values, we restrict the study to this case. The minimum value for $\mathbf{x}^T\mathbf{H}\mathbf{x}$ is zero, and thus the cost for the dual problem is ρ. We re-write the dual problem as a constrained problem

$$d^\star = \max_{\rho\in\mathbb{R},\boldsymbol{\lambda}\in\mathbb{R}^K} \rho \qquad (31)$$
$$\text{subject to } \mathbf{C}-\rho\mathbf{L}-\sum_{i=1}^{K}\lambda_i\mathbf{A}_i \succeq 0$$

which is the problem in DUAL.

Appendix D
Derivation of optimality certifier from general dual problem

For this certifier, we assume strong duality holds, i.e.,$d^\star=\rho=f^\star=\mathbf{x}^T\mathbf{C}\mathbf{x}$ and so we are left to obtain the other multipliers $\boldsymbol{\lambda}\in\mathbb{R}^K$. Further, if $\mathbf{x}$ is the primal solution as assumed then $\mathbf{x}^T\mathbf{H}\mathbf{x}=0$ and from the dual solution being feasible, we know as well the associated Hessian must be PSD. Therefore, $\mathbf{x}^T\mathbf{H}\mathbf{x}=0\Leftrightarrow\mathbf{H}\mathbf{x}=\mathbf{0}_{13\times1}$, and so the primal solution lies on the nullspace of the Hessian. We can re-arrange this expression as a linear system on the (candidate to) Lagrange multipliers, $\left(\mathbf{C}-\rho\mathbf{L}-\sum_{i=1}^{K}\lambda_i\mathbf{A}_i\right)\mathbf{x}=\mathbf{0}_{13\times1}\Leftrightarrow[\mathbf{A}_1\mathbf{x}\dots\mathbf{A}_K\mathbf{x}]\boldsymbol{\lambda}=\mathbf{C}\mathbf{x}-\rho\mathbf{L}\mathbf{x}$. The system allows to solve for the multipliers provided the left-hand matrix $[\mathbf{A}_1\mathbf{x}\dots\mathbf{A}_K\mathbf{x}]$, known as the Jacobian, has: (1) not more columns than rows (i.e.,never becomes an underdetermined system); and (2) it is full-rank. Once the candidate to Lagrange multipliers have been computed, it is required to check that they are indeed feasible for the dual problem, that is, the Hessian evaluated at them is PSD. If the Hessian is PSD, then the vector $\boldsymbol{\lambda}$ is dual feasible. By weak duality in (2), we know that $d(\boldsymbol{\lambda},\rho)\leq f(\mathbf{x})$ for all feasible points $(\boldsymbol{\lambda},\rho)$ and $\mathbf{x}$. Since $d(\boldsymbol{\lambda},\rho)=\rho$ and $\rho=\mathbf{x}^T\mathbf{C}\mathbf{x}=f(\mathbf{x})$, then both $(\boldsymbol{\lambda},\rho)$ and $\mathbf{x}$ are the optimal solutions for their respective problems and strong duality holds.

Appendix E
Set of constraints for 3D rotation matrices

In this section we provide the set of constraints employed in this work. We remind the reader that these sets are not new, and have appeared previously in the literature. We include them here for completeness. Since we will work with rows and columns of the same matrix, we opt for employing the notation in [13]. Thus, for the matrix $\mathbf{R}$, the i-th row is denoted by $\mathbf{r}_i\in\mathbb{R}^3$ and the i-th column by $\text{⌙}_i\in\mathbb{R}^3$ (index starts from 0).

$$\text{ROWS}\equiv\begin{Bmatrix}\mathbf{r}_0^T\mathbf{r}_0=1,\ \mathbf{r}_1^T\mathbf{r}_1=1, & \mathbf{r}_2^T\mathbf{r}_2=1\\ \mathbf{r}_0^T\mathbf{r}_1=1,\ \mathbf{r}_0^T\mathbf{r}_2=1, & \mathbf{r}_1^T\mathbf{r}_2=1\end{Bmatrix}. \qquad (32)$$

$$\text{COLS}\equiv\begin{Bmatrix}\text{⌙}_0^T\text{⌙}_0=1,\ \text{⌙}_1^T\text{⌙}_1=1, & \text{⌙}_2^T\text{⌙}_2=1\\ \text{⌙}_0^T\text{⌙}_1=1,\ \text{⌙}_0^T\text{⌙}_2=1, & \text{⌙}_1^T\text{⌙}_2=1\end{Bmatrix}. \qquad (33)$$

As it was shown in [33], the positive determinant constraints of the rotation matrices can be expressed as a quadratic function of its elements. As stated in the document, this set of constraints is also obtained from the relation between the adjugate of the matrix $\text{Adj}(\mathbf{R})$ since $\text{Adj}(\mathbf{R})\mathbf{R}=\text{Det}(\mathbf{R})\mathbf{I}_3$. For rotation matrices, $\text{Det}(\mathbf{R})=+1$ and so $\text{Adj}(\mathbf{R})=+\mathbf{R}^T$, which provides 9 linear independent constraints. Here, r_i for $i=1,\dots,9$ denotes the i-th entry of the rotation matrix,

$$\mathbf{R}\doteq\begin{pmatrix}r_1 & r_2 & r_3\\ r_4 & r_5 & r_6\\ r_7 & r_8 & r_9\end{pmatrix}. \qquad (34)$$

The set DET is given by

$$\text{DET}\equiv\begin{Bmatrix}r_4r_8-r_5r_7=r_3 & r_2r_7-r_1r_8=r_6 & r_1r_5-r_2r_4=r_9\\ r_5r_9-r_6r_8=r_1 & r_3r_8-r_2r_9=r_4 & r_2r_6-r_3r_5=r_7\\ r_6r_7-r_4r_9=r_2 & r_1r_9-r_3r_7=r_5 & r_3r_4-r_1r_6=r_8\end{Bmatrix}. \qquad (35)$$