A Sufficient Condition of Optimality for the Relative Pose Problem Between Cameras

Mercedes Garcia-Salguero* and Javier Gonzalez-Jimenez†
Machine Perception and Intelligent Robotics (MAPIR) Group, System Engineering and Automation Department,
University of Malaga, Campus de Teatinos, 29071 Malaga, Spain
Email: *mercedesgarsal@uma.es, †javiergonzalez@uma.es

Abstract

The Relative Pose problem (RPp) seeks for the relative rotation and translation between two central, calibrated cameras given a set of pair-wise feature correspondences. The RPp is a fundamental block for many 3D computer vision tasks, and hence, the quality of the estimated relative pose is of key importance for the correct performance of these applications. Nonetheless, the RPp is a non-convex problem that presents multiple local minima. Recent non-minimal solvers provide relatively fast certifiable solutions, usually relying on a convex relaxation of the problem; however, there is no guarantee a priori that these relaxations return the optimal solution, *i.e*. are tight.

This work presents a sufficient condition to guarantee that a given solution of the Relative Pose problem (RPp) is the global optimum, in a faster way than evaluating a certifiable algorithm (up to 4 times faster). We state RPp as an optimization problem that minimizes the squared normalized epipolar error over the set of normalized essential matrices. The proposed condition is derived through spectral analysis and builds up on the recently proposed certifiable algorithm in [1]. The results of extensive experiments, with both synthetic and real data, support that by using the proposed conditions we can detect a large number of optimal solutions for most common problem instances.

Index Terms

Relative Pose problem; sufficient optimality condition; optimality certification; spectral analysis

I. Introduction

The Relative Pose problem (RPp) aims to estimate the relative rotation $\boldsymbol{R}$ and translation $\boldsymbol{t}$ between two central, calibrated cameras given a (noisy) set of N pair-wise feature correspondences $(\boldsymbol{f}_i, \boldsymbol{f}'_i)$, as shown in Figure (1). The estimation of the relative pose is the cornerstone of visual odometry [2], [3] and more complex 3D computer vision applications, including Simultaneous Localization and Mapping (SLAM) [4]–[6] and Structure-from-Motion (SfM) [7], [8]. The gold standard approach for RPp formulates it as a 2-view Bundle-adjustment that minimizes the reprojection error [9]. This problem is highly non-convex [8] and a good initialization, usually derived from a simplification of the problem, is crucial for a successful estimation [10], [9].

Common approaches for solving the RPp state it in terms of the estimation of the *essential matrix* $\boldsymbol{E}$ [9, Sec. 9], a 3×3 matrix that encapsulates all the information about the relative

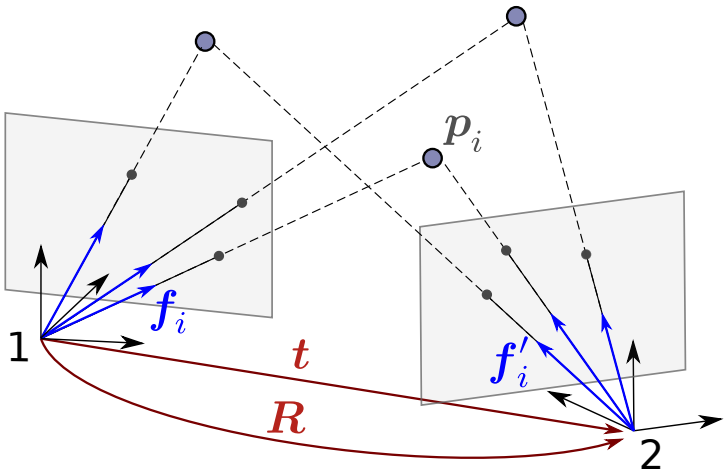

Fig. 1: The RPp seeks the relative rotation $\boldsymbol{R}$ and translation $\boldsymbol{t}$ *up-to-scale* between two central, calibrated cameras $(1-2)$ from a set of N pair-wise feature correspondences $(\boldsymbol{f}_i, \boldsymbol{f}'_i)$ $i = 1, \ldots, N$. Figure taken from [1].

pose. The essential matrix relates each pair of correspondence points $(\boldsymbol{f}_i, \boldsymbol{f}'_i)$ through the so-called *epipolar constraint* $\boldsymbol{f}_i'^T \boldsymbol{E} \boldsymbol{f}_i = 0$. For noisy correspondences, though, this relation does not hold exactly, and so $\boldsymbol{f}_i'^T \boldsymbol{E} \boldsymbol{f}_i = \epsilon_i \neq 0$. Thus, a common method to compute the matrix $\boldsymbol{E}$ is through the minimization of the (squared) epipolar error $\epsilon_i^2 = (\boldsymbol{f}_i'^T \boldsymbol{E} \boldsymbol{f}_i)^2$ [9, Sec. 9]. This basic cost function and more elaborated constructions upon the same concept of keeping minimum the epipolar error, like in [11], [12], lead to non-convex optimization problems, presenting multiple local minima, see *e.g.* [1], [13].

The relative pose, and thus the essential matrix, has only five degrees of freedom: three for the 3D rotation and two for the 3D translation direction, since the scale cannot be recovered for central cameras [9]. Since the epipolar constraint is linear in the entries of the $\boldsymbol{E}$, only five points are necessary to estimate it. This is known as the minimal solver (five-point algorithm, 5-PT) [14], [15], that can be embedded into robust frameworks, such as RANSAC [13], [16]. A simpler method that also estimates the essential matrix is the well-known eight-point algorithm (8-PT), which, although originally devised for the fundamental matrix (the equivalent entity for uncalibrated cameras), can be adapted to the calibrated case, *i.e.* for $\boldsymbol{E}$ [9], [17]. Minimal solvers are highly sensible to noise in the data, hence iterative algorithms that consider all the correspondences (non-minimal solvers) are usually employed, since they achieve better results in the presence of noise [11], [12], [18], [19]. Usually, non-minimal solvers take an initial guess from a minimal solver [11], [14]–[16], [20]–[22] and iteratively refine the solution on the manifold of essential matrices. The main motivation behind this approach is the fact that, although the RPp is non-convex, iterative solvers perform well in practice and repeatedly attain the global solution [11], [12], [19]. Nevertheless, most of the times these methods are sensitive to the given initialization, both for accuracy, optimality and convergence rate [1], [13], [23]. The quality of this initial guess, though, cannot be measured a priori. In summary, given the non-convexity of RPp, there is no guarantee that the solution found by non-minimal solvers is the global optimum.

Though certifiable solvers have been recently proposed for the RPp [13], [23]–[25], they

usually come with the drawback of being inefficient because off-the-shelf tools employed for their resolution have polynomial [26], [27] (or even exponential) time complexity. Special attention has been given recently to the semidefinite positive relaxations (convex) of RPp [23], [24], since they can be actually solved in polynomial time and tend to recover the solution to the original problem in most problem instances, even with high noisy data and low number of correspondences. Also in this line, it was recently proposed in [1] a fast optimality certification algorithm that leverages another convex relaxation of the RPp, and is able to certify if a given solution is optimal or is inconclusive about its optimality. Since the certification does not require convex optimization tools and the relative pose can be estimated by any solver, in practice the computational cost of the algorithm is lower than the above-mentioned methods.

Although for most problem instances the certifiable solvers that leverage convex relaxations have shown to recover and certify the solution to the original problem, there is no guarantee to do so. An approach that usually helps to shed some light on why these solvers perform well in practice is through analyzing *sufficient conditions* for optimality. These conditions are derived from an explicit parameterization of the problem, and *assure* that, if fulfilled, *strong duality* holds for a given solution and problem instance, *i.e*. the convex relaxation is able to estimate the solution to the original, non-convex problem. Thus, the conditions are a mechanism to detect optimality. They tend to be also simpler (and hence, faster) to compute than others certifiable approaches for the same problem including fast certifiers, if they exist; however, they are just sufficient and not necessary for optimality. If the condition is *tight*, though, it can certify most of the optimal solutions for a wide variety of problem instances. Such sufficient conditions do appear in the literature for other problems and are proved to be tight and useful in practice. In [28] Eriksson *et al*. proposed a sufficient condition for the problem of rotation averaging, while Iglesias *et al*. in [29] presented a similar condition for the point cloud registration with missing data. Interestingly, Hartley and Seo in [30] showed that it is possible to bound an optimality zone for the triangulation problem with known rotation or known plane, allowing to certify if a given solution is optimal.

A. Contributions

In this line, and exploiting the certifiable algorithm in [1], we propose a fast and simple sufficient optimality condition for the RPp.

We formulate the problem as an optimization program that minimizes the squared, normalized epipolar error over the set of normalized essential matrices. This condition is defined through meaningful terms that help to understand the hardness of the problem instance. The certifier from which we derive our condition relies on a specific relaxation of the set of essential matrices; however, five other relaxations can be applied to obtain similar yet different certifiers. Any of them is suitable for certification and hence, also as starting point for our condition. We begin with the derivation of one of them and then illustrate how the final form of our condition is extended to any other relaxation of the set of essential matrices, *i.e*. to any certifier. Further, we are able to detect *a priori* which relaxation provides the best (wider) region of optimality for a given problem instance. Additionally, and given the simplicity of our condition, we state under which conditions it can be guaranteed that a given problem instance and a given relaxation attain *strong duality*, *i.e*. the certifier associated can detect the optimal solution. For that, we analyze the form of the eigenvector associated with the smallest eigenvalue of the data matrix (formed by the pair-wise correspondences), without explicitly providing any solution. This statement is actually related to the output of the 8-PT algorithm and the Direct Linear Transform (DLT)

method, the natural extension of the 8-PT to all the correspondences [9]. Through extensive experiments on both synthetic and real data, our condition is proved to be tight in practice for common scenarios, while being faster to compute than any other certifiable approach, including the certifier.

Before continue, we want to remark that our work builds upon the assumption that a set of relaxations is tight, which cannot be assured a priori. Nevertheless, empirically we show that it is indeed tight and thus our proposal works on real problem instances. Our main results are stated in Theorem (3.1) and Corollary (3.1.2).

B. Notation

We denote (column) vectors by bold, lower-case letters, *e.g.* $\boldsymbol{e}$, and matrices by bold, upper-case letters, *e.g.* $\boldsymbol{E}$. The operator $\mathrm{vec}(\boldsymbol{E})$ vectorizes the matrix $\boldsymbol{E}$ column-wise. We denote the kronecker product by $\otimes$ and the direct sum by $\oplus$. The identity matrix of size n is denoted by $\boldsymbol{I}_n$ and the zero matrix of size $n \times m$ by $\mathbf{0}_{n\times m}$. The set of symmetric matrices of size n is denoted by $\mathbb{S}^n$ and the cone of positive semidefinite (PSD) matrices by $\mathbb{S}^n_+$. A matrix $\boldsymbol{Q}$ that is PSD is also indicated as $\boldsymbol{Q} \succeq 0$. We reserve λ for the *Lagrange multipliers*, μ for eigenvalues and σ for singular values. Last by $[\boldsymbol{t}]_\times$ we denote the 3×3 skew-symmetric matrix that forms the cross-product with the 3D vector $\boldsymbol{t} = [t_1, t_2, t_3]^T$, *i.e.*, $\boldsymbol{t} \times (\bullet) = [\boldsymbol{t}]_\times (\bullet)$ with

$$[\boldsymbol{t}]_\times = \begin{bmatrix} 0 & -t_3 & t_2 \\ t_3 & 0 & -t_1 \\ -t_2 & t_1 & 0 \end{bmatrix}.$$

II. Relative Pose problem and Optimality Certification

A. Relative Pose problem

This works tackles the Relative Pose problem (RPp) between two central calibrated cameras, given a set of N pair-wise correspondences $(\boldsymbol{f}_i, \boldsymbol{f}'_i)$, which are the images in the two different views of N (unknown) world points, see Figure (1). We pose the RPp as an optimization problem over the set of normalized essential matrices $\mathbb{E}$, defined as [9], [31]

$$\mathbb{E} = \{\boldsymbol{E} \mid \boldsymbol{E} = [\boldsymbol{t}]_\times \boldsymbol{R},\ \boldsymbol{t} \in \mathcal{S}^2, \boldsymbol{R} \in \mathrm{SO}(3)\}, \tag{1}$$

where $\mathcal{S}^2$ is the 2-sphere

$$\mathcal{S}^2 \doteq \{\boldsymbol{t} \in \mathbb{R}^3 \mid \boldsymbol{t}^T\boldsymbol{t} = 1\} \tag{2}$$

and $\mathrm{SO}(3)$ is the set of 3D rotations as 3×3 matrices,

$$\mathrm{SO}(3) = \{\boldsymbol{R} \in \mathbb{R}^{3\times 3} \mid \boldsymbol{R}\boldsymbol{R}^T = \boldsymbol{I}_3, \det(\boldsymbol{R}) = +1\}. \tag{3}$$

The cost function is defined as the sum of the squared, normalized epipolar error [9], [32] for each pair-wise observation $\epsilon_i^2 = (\boldsymbol{f}_i'^T \boldsymbol{E} \boldsymbol{f}_i)^2$, where we have identified these observations with points in the 2-sphere. This has been the approach followed by previous works [23], [24], and although more complex cost functions have been proposed, this is already a non-convex problem that presents many challenges.

For the purpose of this work, the problem has to be written as a Quadratically Constrained Quadratic Program (QCQP) [33]. Hence, the set of essential matrices has to be defined explicitly and globally; among all the definitions available for the essential matrices [17], [24], a minimal representation is preferred. We employ the minimal parameterization proposed by Faugeras *et*

al. in [17] and recently leveraged in [1], [24]. This definition provides with a set of seven independent constraints given by:

$$h_0 \equiv \boldsymbol{t}^T\boldsymbol{t} = 1 \tag{4}$$
$$h_1 \equiv \boldsymbol{e}_1^T\boldsymbol{e}_1 = t_2^2 + t_3^2 \quad , \quad h_6 \equiv \boldsymbol{e}_1^T\boldsymbol{e}_2 = -t_1t_2 \tag{5}$$
$$h_2 \equiv \boldsymbol{e}_2^T\boldsymbol{e}_2 = t_1^2 + t_3^2 \quad , \quad h_5 \equiv \boldsymbol{e}_1^T\boldsymbol{e}_3 = -t_1t_3 \tag{6}$$
$$h_3 \equiv \boldsymbol{e}_3^T\boldsymbol{e}_3 = t_1^2 + t_2^2 \quad , \quad h_4 \equiv \boldsymbol{e}_2^T\boldsymbol{e}_3 = -t_2t_3 \tag{7}$$

where $\boldsymbol{e}_i \in \mathbb{R}^3$ stands for the i-th row of the matrix $\boldsymbol{E}$ and $\boldsymbol{t}$ is the translation vector associated to $\boldsymbol{E}$.

We write the RPp as a standard QCQP by introducing the positive semidefinite (PSD) data matrix $\boldsymbol{Q}_9 = \frac{1}{N}\sum_{i=1}^{N}(\boldsymbol{f}_i' \otimes \boldsymbol{f}_i)(\boldsymbol{f}_i' \otimes \boldsymbol{f}_i)^T \in \mathbb{S}_+^9$ [1] such that the cost function $f(\boldsymbol{E}) = \sum_{i=1}^{N}\frac{1}{N}\epsilon_i^2$ is written as $f(\boldsymbol{E}) = \mathrm{vec}(\boldsymbol{E})^T\boldsymbol{Q}_9\,\mathrm{vec}(\boldsymbol{E})$. The constraints are also expressed in their quadratic forms as $h_i \equiv \boldsymbol{x}^T\boldsymbol{A}_i\boldsymbol{x} = c_i$ for $i = 0, \ldots, 6$, where $\boldsymbol{A}_i \in \mathbb{S}^{12}, c_i \in \mathbb{R}$ and the vector $\boldsymbol{x}$ contains all the variables of the problem, *i.e.* $\boldsymbol{x} \doteq [\mathrm{vec}(\boldsymbol{E})^T, \boldsymbol{t}^T]^T \in \mathbb{R}^{12}$. By defining the matrix $\boldsymbol{Q}_{12} \doteq \boldsymbol{Q}_9 \oplus \boldsymbol{0}_{3\times 3} \in \mathbb{S}_+^{12}$ the RPp has the standard QCQP form:

$$f^\star = \min_{\boldsymbol{x}\in\mathbb{R}^{12}} \boldsymbol{x}^T\boldsymbol{Q}_{12}\boldsymbol{x} \quad \text{subject to } \boldsymbol{x}^T\boldsymbol{A}_i\boldsymbol{x} = c_i,\ i = 0\ldots, 6 \tag{O}$$

B. Optimality Certification

This problem is, however, non-convex and thus, presents multiple local minima. We proposed recently in [1] a fast algorithm that allows to certify the optimality of a given solution $\hat{\boldsymbol{x}}$ for this same problem, leveraging the minimal parameterization of the set of essential matrices in (6). It was shown that this parameterization cannot be leveraged in such form and a relaxation becomes necessary in order to derive a suitable fast certifier: one of the six constraints associated with $\boldsymbol{E}$ (Equations (5)-(7)) must be dropped.

Since the condition derived in this paper relies on the specific form of the certifier and the associated relaxation of the set (6), let us concretize them. Among all these six relaxations we choose the one given by the feasible set

$$h_0 \equiv \boldsymbol{t}^T\boldsymbol{t} = 1 \quad , \quad h_5 \equiv \boldsymbol{e}_1^T\boldsymbol{e}_3 = -t_1t_3 \tag{8}$$
$$h_1 \equiv \boldsymbol{e}_1^T\boldsymbol{e}_1 = t_2^2 + t_3^2 \quad , \quad h_4 \equiv \boldsymbol{e}_2^T\boldsymbol{e}_3 = -t_2t_3 \tag{9}$$
$$h_2 \equiv \boldsymbol{e}_2^T\boldsymbol{e}_2 = t_1^2 + t_3^2 \quad , \quad h_3 \equiv \boldsymbol{e}_3^T\boldsymbol{e}_3 = t_1^2 + t_2^2 \tag{10}$$

For this relaxation and assuming *strong duality* holds (see [1, Th. 5.1]), the closed-form expression for dual candidate $\check{\boldsymbol{\lambda}}$ given a potential solution $\hat{\boldsymbol{x}}$ to the original problem (O) is expressed as

$$[\boldsymbol{A}_0 \mid \boldsymbol{A}_1 \mid \boldsymbol{A}_2 \mid \boldsymbol{A}_3 \mid \boldsymbol{A}_4 \mid \boldsymbol{A}_5](\boldsymbol{I}_6 \otimes \hat{\boldsymbol{x}})\check{\boldsymbol{\lambda}} = \boldsymbol{Q}_{12}\hat{\boldsymbol{x}}, \tag{11}$$

where $\check{\boldsymbol{\lambda}} \in \mathbb{R}^6$ are the *Lagrange multipliers* and recall that $\boldsymbol{I}_6$ is the identity matrix of size 6. We say that the solution $\hat{\boldsymbol{x}}$ is optimal *iff* the *Hessian of the Lagrangian* $\check{\boldsymbol{H}}(\check{\boldsymbol{\lambda}}) \doteq \boldsymbol{Q}_{12} - \sum_{i=0}^{5}\check{\lambda}_i\boldsymbol{A}_i$ is positive semidefinite, *i.e.* $\check{\boldsymbol{H}}(\check{\boldsymbol{\lambda}}) \succeq 0$.

This certification requires (1) the estimation of the candidates to Lagrange multipliers $\check{\boldsymbol{\lambda}}$ in (11) and (2) the computation of the minimum eigenvalue of $\check{\boldsymbol{H}}(\check{\boldsymbol{\lambda}})$. Since there exist 6 different

[1]The scale $\frac{1}{N}$ does not affect the solution of the problem since does not depend on $\boldsymbol{E}$, but it is necessary for our condition to be independent on the number of correspondences. The reason behind this is provided in the *Supplementary material* (C).

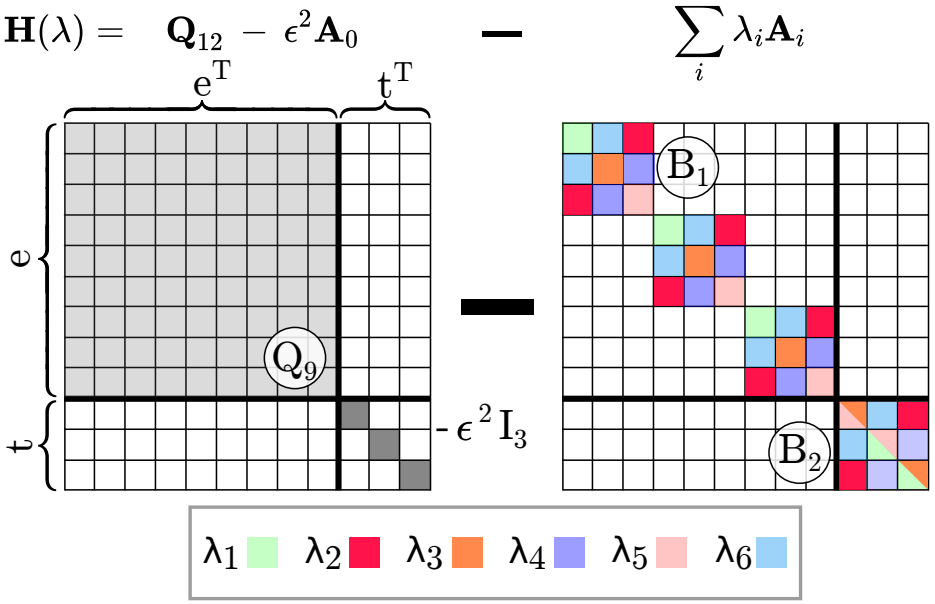

Fig. 2: We can express the Hessian $\boldsymbol{H}(\boldsymbol{\lambda})$ as the sum of block-diagonal matrices $\hat{\boldsymbol{Q}}_{12}, \boldsymbol{A}_0, \ldots, \boldsymbol{A}_6$ with the same pattern: two blocks of size 9×9 and 3×3. We show here all the constraints for completeness; for the relaxation employed through the paper, the multiplier λ_6 is not included. See text for more details.

relaxations, we can derive also 6 closed-forms for dual candidates and their associated Hessian matrices. It is not known *a priori* which relaxation performs better (in terms of certification) and hence, in the worst case, 6 different certifications must be carried out.

III. A Sufficient Condition for Optimality

In this paper we aim to provide a simpler sufficient condition that guarantees that the solution $\hat{\boldsymbol{x}}$ is optimal. The intuitive idea is to define a region of solutions for which the positive semi-definiteness of the Hessian $\check{\boldsymbol{H}}(\check{\boldsymbol{\lambda}})$ is guaranteed without computing explicitly its eigenvalues nor the vector $\check{\boldsymbol{\lambda}}$. Despite the soundness of this concept, such a condition has not been proposed for the RPp to the best of our knowledge.

Our main results are stated below:

Theorem 3.1 (Sufficient Optimality Condition): Let a problem instance has as data matrix $\boldsymbol{Q}_9 \in \mathbb{S}^9_+$. Let the 9D eigenvector associated with the smallest eigenvalue be $\boldsymbol{e}_0$ and its 3×3 form $\boldsymbol{E}_0$. The 3D left eigenvector associated with the smallest eigenvalue of $\boldsymbol{E}_0$ is $\boldsymbol{t}_0$, and the 12D vector $\boldsymbol{x}_0 = [\boldsymbol{e}_0^T, \boldsymbol{t}_0^T]^T$. Let i be the index associated with the best relaxation, such that the minimum singular value $\sigma^i_{\text{null}}(\boldsymbol{x}_0)$ of $\boldsymbol{A}^i_{\text{P}}(\boldsymbol{x})$, obtained by dropping the corresponding rows and columns of $\boldsymbol{A}_{\text{g}}(\boldsymbol{x})$ and $\boldsymbol{P}_{\text{g}}$, is the maximum over all the relaxations (see (D)). A given solution $\boldsymbol{x} = [\boldsymbol{e}^T, \boldsymbol{t}^T]^T$ to the problem (O) with cost value $\epsilon^2 \leq \epsilon_0$, being $\epsilon_0 \geq 0$ a threshold, is *guaranteed* to be optimal (with an optimality certificate) and strong duality holds for the (relaxed) dual problem and the original problem (O) if

$$\eta(\boldsymbol{Q}_9, \boldsymbol{x}) \leq \epsilon_0 \sigma^i_{\text{null}}(\boldsymbol{x}_0), \tag{12}$$

holds where

$$\eta(\boldsymbol{Q}_9, \boldsymbol{x}) \doteq \sqrt{\boldsymbol{e}^T \boldsymbol{Q}_9^2 \boldsymbol{e} + (\epsilon^2)^2} + (\epsilon_0 - \epsilon^2) \left\| \boldsymbol{A}_{\text{P}}^i (\boldsymbol{x} - \boldsymbol{x}_0) \right\|_{\text{F}} + \epsilon^2 \sigma_{\text{null}}^i(\boldsymbol{x}_0). \tag{13}$$

Corollary 3.1.1: In the noiseless case, the sufficient condition derived from any relaxation is able to certify the solution as optimal.

Proof 3.1: In this case, all the correspondences fulfill the epipolar constraint *exactly*, that is, $\epsilon_i = 0$ for all i, and hence $\epsilon^2 = \frac{1}{N}\sum_{i=1}^{N} \epsilon_i^2 = 0$. Further, since $\epsilon^2 = \boldsymbol{e}^T \boldsymbol{Q}_9 \boldsymbol{e} = 0$ and $\boldsymbol{e} \neq \boldsymbol{0}_{9\times 1}$, it follows that $\boldsymbol{Q}_9$ has an eigenvalue equal to zero. We defined $\boldsymbol{e}_0$ as the eigenvector associated with the smaller eigenvalue, but since the solution $\boldsymbol{x}$ lays in the nullspace of the data matrix $\boldsymbol{Q}_9$, we have that $\boldsymbol{x} = \boldsymbol{x}_0$. Then, $\boldsymbol{x} - \boldsymbol{x}_0 = 0 \implies \boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0) = 0$ and $\boldsymbol{e}^T \boldsymbol{Q}_9^2 \boldsymbol{e} = \boldsymbol{e}_0^T \boldsymbol{Q}_9 \boldsymbol{Q}_9 \boldsymbol{e}_0 = 0$. Therefore, $\eta(\boldsymbol{Q}_9, \boldsymbol{x}) = 0 \leq \epsilon_0 \sigma_{\text{null}}^i(\boldsymbol{x}_0)$ and since the right-hand side (RHS) is always non-negative, the solution $\boldsymbol{x}$ is certified as optimal independently of the term $\sigma_{\text{null}}^i(\boldsymbol{x}_0)$, *i.e.* for any relaxation.

Further, given the form of the condition in Theorem (3.1), we can derive an even simpler condition that *guarantees* that strong duality holds for a given problem without providing explicitly any solution

Corollary 3.1.2: Let $k_i \doteq k_i(\boldsymbol{Q}_9, \boldsymbol{e}_0)$, $i = 0, 1, 2$ be three scalar functions of $\boldsymbol{Q}_9$ and $\boldsymbol{e}_0$, as defined in Theorem (3.1), whose explicit forms are provided in the *Supplementary material* (E). If the expression

$$\sqrt{k_1 + k_0^2} + \sqrt{k_2}\epsilon_0 - \sqrt{k_2} k_0 + k_0 \sigma_{\text{null}}^i(\boldsymbol{x}_0) \leq \epsilon_0 \sigma_{\text{null}}^i(\boldsymbol{x}_0), \tag{14}$$

holds, then (1) strong duality holds for this problem (with certificate); (2) the output from the 8-PT algorithm (DLT) lies sufficiently close to the global optimum to be considered as optimal; and hence (3) any refinement of this output leads to the global optimum.

Proof 3.2: Although this may seem a strong claim, the intuition behind it is pretty simple: The expression in (14) is the explicit form of our condition in Theorem (3.1) applied to the output of the 8-PT algorithm (DLT) in terms of the SVD decomposition of $\boldsymbol{E}_0$. Due to space limits, we provide the proof in the *Supplementary material* (E).

The remainder of this Section is devoted to obtain these results. We evaluate the performance of our optimality condition in Section (V).

We start by simplifying the certification algorithm in Section (II-B). First, recall that the closed-form expression for $\check{\boldsymbol{\lambda}}$ in (11) assumes that *strong duality* holds, *i.e.* the optimal cost of the original problem $f^\star$ and the *dual problem* $d^\star$ [33] are equal up to some accuracy. Since the dual problem for this relaxation has the form [1, App. C]

$$d^\star = \max_{\check{\boldsymbol{\lambda}}} \check{\lambda}_0 \quad \text{subject to } \check{\boldsymbol{H}}(\check{\boldsymbol{\lambda}}) \succeq 0 \tag{D-R}$$

we have that $f^\star = d^\star = \check{\lambda}_0$. Let us denote the cost value attained by the solution $\hat{\boldsymbol{x}}$ by $\epsilon^2 \doteq f(\hat{\boldsymbol{x}})$. The certifier assumes that $f(\hat{\boldsymbol{x}}) = f^\star$ and so $\check{\lambda}_0 = \epsilon^2$. In what follows, we will drop the *hat* to clear notation. Hence, we can simplify the expression in (11) by "moving" the first lagrange multiplier λ_0 to the right-hand side and substituting it with ϵ^2. This reduced closed-form is formally given by

$$[\boldsymbol{A}_1 \mid \boldsymbol{A}_2 \mid \boldsymbol{A}_3 \mid \boldsymbol{A}_4 \mid \boldsymbol{A}_5](\boldsymbol{I}_5 \otimes \boldsymbol{x})\boldsymbol{\lambda} = \boldsymbol{Q}_{12}\boldsymbol{x} - \epsilon^2 \boldsymbol{A}_0 \boldsymbol{x}, \tag{15}$$

with $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_5]^T \in \mathbb{R}^5$. For short we will denote

$$\boldsymbol{A}(\boldsymbol{x}) \doteq [\boldsymbol{A}_1 \mid \boldsymbol{A}_2 \mid \boldsymbol{A}_3 \mid \boldsymbol{A}_4 \mid \boldsymbol{A}_5](\boldsymbol{I}_5 \otimes \boldsymbol{x}) \tag{16}$$

$$\hat{\boldsymbol{Q}}_{12} \doteq \boldsymbol{Q}_{12} - \epsilon^2 \boldsymbol{A}_0 = \boldsymbol{Q}_{12} - \epsilon^2 \big(\boldsymbol{0}_{9\times 9} \oplus \boldsymbol{I}_3\big) = (\boldsymbol{Q}_9 \oplus -\epsilon^2 \boldsymbol{I}_3). \tag{17}$$

The Hessian for this reduced certifier is given by $\boldsymbol{H}(\boldsymbol{\lambda}) = \boldsymbol{Q}_{12} - \boldsymbol{A}_0\epsilon^2 - \sum_{i=1}^{5}\lambda_i\boldsymbol{A}_i$, which is the sum of seven block-diagonal matrices with two blocks of size 9×9 and 3×3, as we show in Figure (2). Hence, $\boldsymbol{H} \doteq \boldsymbol{H}(\boldsymbol{\lambda})$ is also a block-diagonal matrix. Let us denote the top-left 9×9 block by $\boldsymbol{H}^E$ and the bottom-right 3×3 block by $\boldsymbol{H}^t$. The matrix $\boldsymbol{H}$ is PSD if and only if each block is also PSD and we only need to analyze these two individuals blocks instead of the full matrix. The explicit expressions for these blocks $\boldsymbol{H}^E, \boldsymbol{H}^t$ are given by (see also Figure (2)):

$$\boldsymbol{H}^E \doteq \boldsymbol{Q}_9 - (\boldsymbol{B}_1 \oplus \boldsymbol{B}_1 \oplus \boldsymbol{B}_1), \qquad \boldsymbol{H}^t \doteq -\epsilon^2\boldsymbol{I}_3 - \boldsymbol{B}_2. \tag{18}$$

where $\boldsymbol{B}_1, \boldsymbol{B}_2$ are 3×3 symmetric matrices with the form:

$$\boldsymbol{B}_1 \doteq \begin{pmatrix} \lambda_1 & 0 & \lambda_2/2 \\ 0 & \lambda_3 & \lambda_4/2 \\ \lambda_2/2 & \lambda_4/2 & \lambda_5 \end{pmatrix}, \qquad \boldsymbol{B}_2 \doteq \begin{pmatrix} -\lambda_3-\lambda_5 & 0 & \lambda_2/2 \\ 0 & -\lambda_1-\lambda_5 & \lambda_4/2 \\ \lambda_2/2 & \lambda_4/2 & -\lambda_1-\lambda_3 \end{pmatrix}. \tag{19}$$

See that, by definition, any PSD matrix has as minimum eigenvalue a non-negative number. Hence, we can re-write the two PSD conditions as $\mu^E \geq 0, \mu^t \geq 0$, where μ^E, μ^t are the minimum eigenvalues of the matrices $\boldsymbol{H}^E, \boldsymbol{H}^t$, respectively. In order to accommodate numerical errors, we will ask the minimum eigenvalues to be greater than a negative threshold $-\epsilon_0$, with $\epsilon_0 \geq 0$.

Weyl's inequality for eigenvalues [34, Sec. 4.3] allows us to express this PSD conditions in terms of the eigenvalues of the $\boldsymbol{B}_1, \boldsymbol{B}_2$. Let us define the minimum eigenvalue of $\boldsymbol{Q}_9$ by μ^Q, the minimum eigenvalue of $\boldsymbol{B}_1$ by $\mu^{\text{B1}}_{\text{min}}$ and its maximum eigenvalue by $\mu^{\text{B1}}_{\text{max}}$ (similarly for $\boldsymbol{B}_2$ and $\mu^{\text{B2}}_{\text{min}}, \mu^{\text{B2}}_{\text{max}}$). Therefore, Weyl's inequality assures that:

$$\mu^E \geq \mu^Q - \mu^{\text{B1}}_{\text{max}}, \quad \text{and} \quad \mu^t = -\epsilon^2 - \mu^{\text{B2}}_{\text{max}}, \tag{20}$$

where we use the fact that $\mu_{\text{min}}(-\boldsymbol{X}) = -\mu_{\text{max}}(\boldsymbol{X})$ for any matrix $\boldsymbol{X}$ and that the minimum eigenvalue of a block-diagonal matrix is the minimum eigenvalue among all blocks. Notice that in the second expression equality holds since $\boldsymbol{I}_3$ is the identity matrix of size 3.

By restricting the right-hand sides of each inequality in (20) to be greater than $-\epsilon_0$, we are automatically assuring that μ^E, μ^t are also greater than this threshold.

$$\mu^Q - \mu^{\text{B1}}_{\text{max}} \geq -\epsilon_0 \Leftrightarrow \mu^{\text{B1}}_{\text{max}} \leq \mu^Q + \epsilon_0 \tag{21}$$

$$-\epsilon^2 - \mu^{\text{B2}}_{\text{max}} \geq -\epsilon_0 \Leftrightarrow \mu^{\text{B2}}_{\text{max}} \leq \epsilon_0 - \epsilon^2 \tag{22}$$

Note that the right-hand side of the inequality (21) is always greater than zero since $\boldsymbol{Q}_9$ is PSD by construction and $\epsilon_0 \geq 0$. On the other hand, the right-hand side of (22) may be less than zero if the sum of the residuals ϵ^2 is greater than the threshold ϵ_0. In what follows, we assume that the expression $\epsilon_0 - \epsilon^2 \geq 0$ holds, which in turn imposes a condition on the cost attained by the optimal solution. While the condition $\epsilon_0 - \epsilon^2 \geq 0$ may appear as a strong assumption, it is necessary to recall that the derived optimality condition will be sufficient but not necessary. That is, we expect the condition to fail for some problem instances, specially those with large noise and/or low number of correspondences, as we show empirically in Section (V). For those problem instances in which the final condition can work, generally the cost of the solution ϵ^2 will be smaller than the threshold ϵ_0. While we contemplate to analyze the case in which the assumption $\epsilon_0 - \epsilon^2 \geq 0$ does not hold, problem instances with optimal cost of the magnitude of ϵ_0 present highly noisy data, which may hinder the tightness of the relaxations making the optimality condition to fail eventually.

Instead of working with the (signed) maximum eigenvalues, we will treat their absolute values and introduce the spectral radius $\rho(\boldsymbol{B}_1), \rho(\boldsymbol{B}_2)$ for each matrix. By definition, the spectral radius of a matrix $\boldsymbol{B}_1$ takes the form $\rho(\boldsymbol{B}_1) \doteq \max(\left|\mu_{\text{min}}^{\text{B1}}\right|, \left|\mu_{\text{max}}^{\text{B1}}\right|)$ and hence, for any eigenvalue μ_i of the matrix (including the extremes) $\mu_i \leq |\mu_i| \leq \rho(\boldsymbol{B}_1)$. We can further restrict the inequalities in (21), (22) as:

$$\mu_{\text{max}}^{\text{B1}} \leq \left|\mu_{\text{max}}^{\text{B1}}\right| \leq \rho(\boldsymbol{B}_1) \leq \mu^Q + \epsilon_0 \tag{23}$$

$$\mu_{\text{max}}^{\text{B2}} \leq \left|\mu_{\text{max}}^{\text{B2}}\right| \leq \rho(\boldsymbol{B}_2) \leq \epsilon_0 - \epsilon^2. \tag{24}$$

Our interest in the spectral radius is its relation with any matrix norm $\|\bullet\|$ since the inequality $\rho(\boldsymbol{X}) \leq \|\boldsymbol{X}\|$ holds for any matrix $\boldsymbol{X}$. We bound the spectral radius in (23),(24) with the frobenius norm of each matrix as

$$\rho(\boldsymbol{B}_1) \leq \|\boldsymbol{B}_1\|_{\text{F}} \leq \mu^Q + \epsilon_0 \tag{25}$$

$$\rho(\boldsymbol{B}_2) \leq \|\boldsymbol{B}_2\|_{\text{F}} \leq \epsilon_0 - \epsilon^2 \tag{26}$$

For this relaxation, the Frobenius norms have the form:

$$\|\boldsymbol{B}_1\|_{\text{F}}^2 = \lambda_1^2 + \lambda_3^2 + \lambda_5^2 + (\lambda_2/\sqrt{2})^2 + (\lambda_4/\sqrt{2})^2 \tag{27}$$

$$\begin{aligned}\|\boldsymbol{B}_2\|_{\text{F}}^2 &= (\lambda_3 + \lambda_5)^2 + (\lambda_1 + \lambda_3)^2 + (\lambda_1 + \lambda_5)^2 + \\ &+ (\lambda_2/\sqrt{2})^2 + (\lambda_4/\sqrt{2})^2.\end{aligned} \tag{28}$$

These norms can be reformulated as quadratic forms in the entries of $\boldsymbol{\lambda}$ by introducing two positive definite matrices $\hat{\boldsymbol{S}}, \hat{\boldsymbol{P}} \in \mathbb{S}_+^5$ such $\|\boldsymbol{B}_1\|_{\text{F}}^2 = \boldsymbol{\lambda}^T \hat{\boldsymbol{S}}^{-1} \boldsymbol{\lambda}$ and $\|\boldsymbol{B}_2\|_{\text{F}}^2 = \boldsymbol{\lambda}^T \hat{\boldsymbol{P}}^{-1} \boldsymbol{\lambda}$. These matrices depend on the chosen relaxation and for the one performed in this work (10) they have the form

$$\hat{\boldsymbol{S}}^{-1} \doteq \text{diag}(1, 1/2, 1, 1/2, 1),\ \hat{\boldsymbol{P}}^{-1} \doteq \begin{pmatrix} 2 & 0 & 1 & 0 & 1 \\ 0 & 1/2 & 0 & 0 & 0 \\ 1 & 0 & 2 & 0 & 1 \\ 0 & 0 & 0 & 1/2 & 0 \\ 1 & 0 & 1 & 0 & 2 \end{pmatrix}. \tag{29}$$

Consider now two full-rank matrices $\boldsymbol{S}, \boldsymbol{P} \in \mathbb{R}^{5\times 5}$ such that $\hat{\boldsymbol{S}} = \boldsymbol{S}\boldsymbol{S}^T$ and $\hat{\boldsymbol{P}} = \boldsymbol{P}\boldsymbol{P}^T$. See that $\hat{\boldsymbol{S}}^{-1} = \boldsymbol{S}^{-T}\boldsymbol{S}^{-1}$ and $\hat{\boldsymbol{P}}^{-1} = \boldsymbol{P}^{-T}\boldsymbol{P}^{-1}$ hold since $\boldsymbol{P}^T$ is the transpose of $\boldsymbol{P}$ (similar for $\boldsymbol{S}$). Let us define two vectors $\tilde{\boldsymbol{\lambda}} = [\tilde{\lambda}_1, \ldots, \tilde{\lambda}_5]^T$ and $\hat{\boldsymbol{\lambda}} = [\hat{\lambda}_1, \ldots, \hat{\lambda}_5]^T$ as linear combinations of the original Lagrange multipliers $\boldsymbol{\lambda}$ such that $\boldsymbol{P}\tilde{\boldsymbol{\lambda}} = \boldsymbol{S}\hat{\boldsymbol{\lambda}} = \boldsymbol{\lambda}$. Then, we can write the frobenius norms in (27), (28) in terms of these vectors as $\|\boldsymbol{B}_1\|_{\text{F}} = \left\|\hat{\boldsymbol{\lambda}}\right\|_2 = \left\|\boldsymbol{S}^{-1}\boldsymbol{\lambda}\right\|_2$ and $\|\boldsymbol{B}_2\|_{\text{F}} = \left\|\tilde{\boldsymbol{\lambda}}\right\|_2 = \left\|\boldsymbol{P}^{-1}\boldsymbol{\lambda}\right\|_2$. We are left to bound the norms of $\hat{\boldsymbol{\lambda}}, \tilde{\boldsymbol{\lambda}}$.

Remark 1: Before continuing, we want to point out that the matrices $\boldsymbol{P}, \boldsymbol{S}$ are not uniquely defined. See that $\hat{\boldsymbol{S}}, \hat{\boldsymbol{P}}$ are symmetric by construction, and thus they admit an eigendecomposition of the form $\hat{\boldsymbol{S}} = \boldsymbol{U}_{\text{S}}\boldsymbol{D}_{\text{S}}^2\boldsymbol{U}_{\text{S}}^T$ and $\hat{\boldsymbol{P}} = \boldsymbol{U}_{\text{P}}\boldsymbol{D}_{\text{P}}^2\boldsymbol{U}_{\text{P}}^T$. We can define the matrices $\boldsymbol{P}, \boldsymbol{S}$ as $\boldsymbol{S} = \boldsymbol{U}_{\text{S}}\boldsymbol{D}_{\text{S}}\boldsymbol{O}_{\text{S}}$ and $\boldsymbol{P} = \boldsymbol{U}_{\text{P}}\boldsymbol{D}_{\text{P}}\boldsymbol{O}_{\text{P}}$, where $\boldsymbol{O}_{\text{P}}, \boldsymbol{O}_{\text{S}}$ are orthogonal matrices of size 5. Due to the form of our main result, this ambiguity does not present any problem. For now, we will just refer to these matrices without any explicit form, and clarify later the reasons for this.

Recall that the Lagrange multipliers were obtained from equation (15). Hence, the 5D vectors $\tilde{\boldsymbol{\lambda}}, \hat{\boldsymbol{\lambda}}$ are obtained as

$$\boldsymbol{A}(\boldsymbol{x})\boldsymbol{P}\tilde{\boldsymbol{\lambda}} = \hat{\boldsymbol{Q}}_{12}\boldsymbol{x} \text{ and } \boldsymbol{A}(\boldsymbol{x})\boldsymbol{S}\hat{\boldsymbol{\lambda}} = \hat{\boldsymbol{Q}}_{12}\boldsymbol{x}. \tag{30}$$

Let us define the matrices $\boldsymbol{A}_{\text{P}}(\boldsymbol{x}) \doteq \boldsymbol{A}(\boldsymbol{x})\boldsymbol{P}$ and $\boldsymbol{A}_{\text{S}}(\boldsymbol{x}) \doteq \boldsymbol{A}(\boldsymbol{x})\boldsymbol{S}$. With them, the solutions to the linear overdetermined systems in (30) are obtained in a least squares sense [2] as $\tilde{\boldsymbol{\lambda}} = \boldsymbol{A}_{\text{P}}(\boldsymbol{x})^{\dagger}\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}, \hat{\boldsymbol{\lambda}} = \boldsymbol{A}_{\text{S}}(\boldsymbol{x})^{\dagger}\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}$, where $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})^{\dagger}, \boldsymbol{A}_{\text{S}}(\boldsymbol{x})^{\dagger}$ are the Moore-Penrose pseudo-inverses of $\boldsymbol{A}_{\text{P}}(\boldsymbol{x}), \boldsymbol{A}_{\text{S}}(\boldsymbol{x})$, respectively.

We can relate the ℓ_2 norm of $\tilde{\boldsymbol{\lambda}}, \hat{\boldsymbol{\lambda}}$ with compatible matrix-vector norms of these expressions and introduce them as upper bounds on the frobenius norms in (27) and (28) as

$$\left\|\hat{\boldsymbol{\lambda}}\right\|_2 \leq \left|\left|\left|\boldsymbol{A}_{\text{S}}(\boldsymbol{x})^{\dagger}\right|\right|\right|_2 \left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2 \leq \mu^Q + \epsilon_0, \tag{31}$$

$$\left\|\tilde{\boldsymbol{\lambda}}\right\|_2 \leq \left|\left|\left|\boldsymbol{A}_{\text{P}}(\boldsymbol{x})^{\dagger}\right|\right|\right|_2 \left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2 \leq \epsilon_0 - \epsilon^2, \tag{32}$$

where $\left|\left|\left|\boldsymbol{A}_{\text{P}}(\boldsymbol{x})^{\dagger}\right|\right|\right|_2$ is the spectral norm of the matrix, *i.e.* the maximum singular value of $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})^{\dagger}$, which in turn is the inverse of the minimum non-zero singular value of $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})$ [3] (and similar for $\boldsymbol{A}_{\text{S}}(\boldsymbol{x})$). Let the minimum singular values of $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})$ and $\boldsymbol{A}_{\text{S}}(\boldsymbol{x})$ be $\sigma_{\boldsymbol{A}}(\boldsymbol{x})$ and $\sigma_{\boldsymbol{S}}(\boldsymbol{x})$, respectively. Then, we can express the conditions in (31),(32) as

$$\left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2 \leq (\mu^Q + \epsilon_0)\sigma_{\boldsymbol{S}}(\boldsymbol{x}), \tag{33}$$

$$\left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2 \leq (\epsilon_0 - \epsilon^2)\sigma_{\boldsymbol{A}}(\boldsymbol{x}), \tag{34}$$

See that $\sigma_{\boldsymbol{A}}(\boldsymbol{x}) \leq \sigma_{\boldsymbol{S}}(\boldsymbol{x})$ by construction of $\boldsymbol{P}, \boldsymbol{S}$ (see *Supplementary material* (A)) so that the next chain of inequalities holds

$$\left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2 \leq (\epsilon_0 - \epsilon^2)\sigma_{\boldsymbol{A}}(\boldsymbol{x}) \leq (\epsilon_0 - \epsilon^2)\sigma_{\boldsymbol{S}}(\boldsymbol{x}) \leq (\mu^Q + \epsilon_0)\sigma_{\boldsymbol{S}}(\boldsymbol{x}).$$

Hence, we only need to analyze the inequality in Equation (34).

The matrix $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})$ and thus the minimum singular value $\sigma_{\boldsymbol{A}}(\boldsymbol{x})$, however, depends on the solution $\boldsymbol{x}$. To remove this dependency, let us first define the one-dimensional eigenspace associated with the smallest eigenvalue of the data matrix $\boldsymbol{Q}_9$ by $\boldsymbol{e}_0$. After reshaping it into a 3×3 matrix, the left eigenvector associated with its smallest eigenvalue is $\boldsymbol{t}_0$. Let us also define the vector $\boldsymbol{x}_0 = [\boldsymbol{e}_0^T, \boldsymbol{t}_0^T]^T$ and the matrix $\boldsymbol{A}_{\text{P}}(\boldsymbol{x}_0)$ with the same construction as $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})$. Further, let $\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)$ have this same structure such that $\boldsymbol{A}_{\text{P}}(\boldsymbol{x}) = \boldsymbol{A}_{\text{P}}(\boldsymbol{x}_0) + \boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)$.

We leverage again the Weyl's inequality but for singular values. Let the minimum singular value of $\boldsymbol{A}_{\text{P}}(\boldsymbol{x}_0)$ be $\sigma_{\text{null}}(\boldsymbol{x}_0)$. Then, Weyl's inequality relates these values as

$$|\sigma_{\boldsymbol{A}}(\boldsymbol{x}) - \sigma_{\text{null}}(\boldsymbol{x}_0)| \leq |||\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)|||_2, \tag{35}$$

where again $|||\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)|||_2$ is the spectral norm of $\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)$. We want to replace the singular value $\sigma_{\boldsymbol{A}}(\boldsymbol{x})$ in (34) (which depends on $\boldsymbol{x}$) by the value $\sigma_{\text{null}}(\boldsymbol{x}_0)$ (which only depends on the data $\boldsymbol{Q}_9$). To do this, our substitution must always represent a lower bound on the upper limit $(\epsilon_0 - \epsilon^2)\sigma_{\boldsymbol{A}}(\boldsymbol{x})$. Then, the worst case is given when $\sigma_{\text{null}}(\boldsymbol{x}_0) \geq \sigma_{\boldsymbol{A}}(\boldsymbol{x})$ and so

$$0 \leq \sigma_{\text{null}}(\boldsymbol{x}_0) - \sigma_{\boldsymbol{A}}(\boldsymbol{x}) \leq |||\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)|||_2 \Leftrightarrow \tag{36}$$

$$\sigma_{\text{null}}(\boldsymbol{x}_0) - |||\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)|||_2 \leq \sigma_{\boldsymbol{A}}(\boldsymbol{x}) \tag{37}$$

[2] Due to numerical errors, the right-hand side on the system does not lie in the span of the coefficient matrix, and the "closer" solution in the least squares sense is computed.

[3] Recall that the singular values are, by definition, non-negative.

Since $(\epsilon_0 - \epsilon^2) \geq 0$, we have that $(\epsilon_0 - \epsilon^2)(\sigma_{\text{null}}(\boldsymbol{x}_0) - |||\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)|||_2) \leq (\epsilon_0 - \epsilon^2)\sigma_{\boldsymbol{A}}(\boldsymbol{x})$. See that in the other case where $\sigma_{\boldsymbol{A}}(\boldsymbol{x}) \geq \sigma_{\text{null}}(\boldsymbol{x}_0)$, this inequality also holds since the ℓ_2 norm of any matrix is always non-negative.

Then, condition (34) is guaranteed to hold if the expression

$$\left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2 \leq (\epsilon_0 - \epsilon^2)(\sigma_{\text{null}}(\boldsymbol{x}_0) - |||\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)|||_2) \tag{38}$$

holds. Further, since $|||\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)|||_2 \leq \|\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)\|_{\text{F}}$ for any matrix $\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)$, we can give a simpler condition

$$\left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2 + (\epsilon_0 - \epsilon^2)\|\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)\|_{\text{F}} + \epsilon^2\sigma_{\text{null}}(\boldsymbol{x}_0) \leq \epsilon_0\sigma_{\text{null}}(\boldsymbol{x}_0), \tag{39}$$

where $\|\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)\|_{\text{F}}$ has a closed-form expression on the entries of $\boldsymbol{x} - \boldsymbol{x}_0$.

Notice that the expression in (39) can be re-written in such a form that represents an upper bound on the actual absolute value of the minimum eigenvalue of the Hessian by dividing the left-hand side (LHS) by the term $\sigma_{\text{null}}(\boldsymbol{x}_0)$. Last, we want to point out that the matrix $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})$ appears in the sufficient condition as a frobenius norm and as a singular value. This makes the explicit form of $\boldsymbol{P}$ not important, provided some relation with the matrix $\hat{\boldsymbol{P}}$ holds. We provide the proof in the *Supplementary material* (B).

IV. A Unified Sufficient Condition: Introducing the Other Relaxations

Nevertheless, the sufficient condition proposed in (39) depends on the specific relaxation employed (set (10)). Other relaxations of the same set given in (6) lead to similar, yet different certifiers, and thus, a different sufficient condition can be derived from any of them for the same problem. Notice that we can choose any of them to certify the same solution for the same problem without loss of generality. Our interest here is that if at least one of them is able to certify the solution as optimal, then the solution is optimal, even if the rest of the certifiers cannot certify it as such. Nevertheless, to the best of our knowledge, it is not possible to know a priori which relaxation performs better in terms of certification, *i.e*. we do not know which relaxation is the tightest.

Fortunately, with our sufficient condition it is *possible* to detect *a priori* the best relaxation for a given problem instance. For any other relaxation, a similar sufficient condition can be derived with a minimal effort following the procedure given here. The form of the condition for all the relaxations is (39) and the specific set employed only affects the matrix $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})$. Hence, we only need to change the form of $\boldsymbol{A}_{\text{P}}(\boldsymbol{x}) = \boldsymbol{A}(\boldsymbol{x})\boldsymbol{P}$ to range over all the relaxations. The matrix $\boldsymbol{A}^i_{\text{P}}(\boldsymbol{x}) = \boldsymbol{A}^i(\boldsymbol{x})\boldsymbol{P}^i$ associated with the i-th relaxation (with $i = 1, \ldots, 5$) is obtained by discarding the i-th column of $\boldsymbol{A}_{\text{g}}(\boldsymbol{x})$ and the i-th row and column of $\boldsymbol{P}_{\text{g}}$, whose explicit expressions are given in the *Supplementary material* (D).

Hence, we are left to provide the criteria for this selection. For that, we need to observe the expression in (39). For a given problem instance $\boldsymbol{Q}_9$ and solution $\boldsymbol{x}$, the terms $\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}$ and ϵ^2 are constant. Further, the objective value ϵ^2 tends to be close to zero and hence $\epsilon^2 \ll \epsilon_0$. Two terms are left that depends on the chosen relaxation via the matrix $\boldsymbol{A}_{\text{P}}(\boldsymbol{x})$: (1) the minimum singular value $\sigma_{\text{null}}(\boldsymbol{x}_0)$ and (2) the closed-form expression for $\|\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)\|_{\text{F}}$. For the latter, we observe that, for the same vector $\boldsymbol{x} - \boldsymbol{x}_0$, its value for the different relaxations is similar and the difference between them goes to $1e-04$. The selection of the relaxation can be fully based on the minimum singular value $\sigma_{\text{null}}(\boldsymbol{x}_0)$. Empirically, we observe that the difference between these terms for the same problem instance is *not* negligible and the "best" relaxation should be detected and employed; the index of this relaxation is not constant, though. Thus,

in practice, we will check all the relaxations and choose the one that provides with the largest value $\sigma_{\text{null}}(\boldsymbol{x}_0)$; see that we only need to compute it once per problem instance. This criteria can be also derived by re-writting the inequality in (39) as

$$\left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2 \leq (\epsilon_0 - \epsilon^2)(\sigma_{\text{null}}(\boldsymbol{x}_0) - \|\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)\|_{\text{F}})$$

. Since $\left\|\hat{\boldsymbol{Q}}_{12}\boldsymbol{x}\right\|_2$ and $(\epsilon_0 - \epsilon^2)$ are constant for the same problem and solution, the "best" relaxation is the one with largest $\sigma_{\text{null}}(\boldsymbol{x}_0) - \|\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)\|_{\text{F}}$. Since the solution $\boldsymbol{x}$ tends to be close to the eigenvector $\boldsymbol{x}_0$, the term $\|\boldsymbol{A}_{\text{P}}(\boldsymbol{x} - \boldsymbol{x}_0)\|_{\text{F}}$ is small and therefore the $\sigma_{\text{null}}(\boldsymbol{x}_0)$ is the dominant term. This concludes the proof of Theorem (3.1).

V. Experimental Validation

We evaluate the performance of our final proposal in Theorem (3.1) on both synthetic and real data under a wide variety of configurations. We estimate the potential optimal solution for a given problem instance by refining iteratively on the manifold of essential matrices [19] the initial guess output by the 8-PT algorithm, as it was done in [1].

A. Evaluation on Synthetic Data

We create a set of synthetic scenes with different configurations of parameters and measure the performance of our condition w.r.t. the certifier.

Data generation: We generate the synthetic problem instances as follows: We place the first camera to the origin (identity rotation and zero translation) and place a set of N world points in the square frustum formed by the Field of View (FOV) of the camera and with depth ranging from one to eight meters. Then, we generate the second camera pose such that (1) all the world points lie in front of the second camera and (2) are within its FOV. To obtain the observations, we project the world points by assuming a pin-hole camera model with fixed focal length. We consider all the points since they all lie within the FOV of the cameras. The optical center is defined in the middle of the image for both cameras. We introduce the specified level of noise σ (in $\texttt{pix}$) in both observations by sampling the (isotropic) gaussian distribution centered at the projections with standard deviation σ. The unit-norm vectors $\boldsymbol{f}_i, \boldsymbol{f}'_i$ are finally obtained by applying the inverse of the camera matrix $\boldsymbol{K}$ and normalizing the results.

The second pose of the camera is generated with a bounded rotation of 0.5 degrees and setting the following parameters for both cameras as default: FOV $= 100$ degrees, noise level $\sigma = 0.5$ $\texttt{pix}$, translation magnitude (parallax) of 2 meters (ratio $2:1$ w.r.t. the minimum depth of the points) and focal length $f = 800$ $\texttt{pix}$. For each set of experiments, we create a configuration of cameras and point clouds with number of correspondences $N \in \{8, 15, 20, 40, 70, 100, 150, 200\}$. We vary one parameter each time, and for each combination parameter-number of correspondences, we generate 1000 random problem instances. The threshold ϵ_0 is fixed to $2e-04$ in all the experiments.

Performance measurement: We take the result returned by the certifier as ground-truth [4], and consider as True Positive (TP) a solution that is certified as optimal by both the certifier and our condition. A False Positive (FP) is a solution that is certified as optimal by our condition but not by the certifier. Last, a False Non-Positive (FNP) is a solution that is considered optimal according to the certifier, but not by the condition. Then, we measure the tightness of our

[4]We employ the relaxation for the certifier that provides us with the best (wider) sufficient condition.

sufficient condition (39) by the Recall metric as in [1]. Precision [5] is not employed here since, by construction, our condition can not incur in FP outcomes. Thus, the precision metric will be always one. The Recall metric reads:

$$\text{RECALL} = \frac{\text{TP}}{\text{TP} + \text{FNP}}. \tag{40}$$

A deviation of the Recall from one means that there are some optimal solutions that are not detected by the condition; the recall, though, is shown to tend to one for common problem instances.

Noise level: We vary the noise level with values in $\sigma \in \{0.1, 0.5, 1.0, 1.5, 2.0\}$ $\texttt{pix}$. Figure (3a) depicts the recall metrics. The precision is always 1, as expected, while we observe that the recall increases with the number of correspondences. With low-medium noise $0.1 - 0.5$ $\texttt{pix}$, our condition is tight (100% of detected optimal solutions), while with high noise, we obtain an acceptable ratio with N large. Since in practice the number of correspondences is larger than 70, we perform another set of experiments with $N \in \{70, 100, 15, 200, 300\}$ and noise level from 1.0 $\texttt{pix}$ to 2.5 $\texttt{pix}$, with step size of 0.1 $\texttt{pix}$. For each configuration, 1000 problem instances are created. The precision was always one and Figure (3b) shows the recall for these experiments. Observe how with $N = 200$ our condition detect more than 95% of optimal solutions, while with $N = 300$ the recall is close to 1 even for high noise $2.0 - 2.5$ $\texttt{pix}$. Last, for this same set of experiments, we show in Figure (3c) the mean value of the absolute minimum eigenvalue of the Hessian of the certifier $\boldsymbol{H}$ (dashed line) and the approximation obtained by our condition as $\eta(\boldsymbol{Q}_9, \boldsymbol{x})/\sigma^i_{\text{null}}(\boldsymbol{x}_0)$ (solid line). The mean value of our condition is always under the threshold $\epsilon_0 = 2e - 04$ even for high noise and $N = 70$ correspondences. Additionally, for this same set of experiments we show in the *Supplementary material* (F) the recall metric for those cases in which we could assure strong duality by Corollary (3.1.2).

Translation magnitude: With a noise level of 0.5 $\texttt{pix}$, we vary the translation magnitude of the second camera with values $\|\boldsymbol{t}\|_2 \in \{0.3, 0.8, 3, 4.5, 8, 10, 25\}$. We relate these magnitudes with the depth of the point cloud of the scene in order to reflect the amount of parallax for each translation. Notice that for small $\|\boldsymbol{t}\|_2$ (and hence small ratio translation-depth), the configuration of cameras approaches one of the degenerate situations (zero translation). On the other hand, a large ratio creates more parallax (a better constrained problem instance); however, if this is too large, the scene approaches the coplanar (degenerate) configuration. We define this ratio as t/c in the graphics, where c is the centroid of the frustum where the world points lay and t is the translation magnitude. Figure (4a) shows the recall metrics for this set of experiments. We observe that for more than $N = 40$ correspondences, we are able to detect *all* the optimal solutions for all the different magnitudes t/c; below this number, the worst case is attained with the smaller translation magnitude (the configuration of cameras is almost degenerate [35]), although even in this case our condition detects 80% of optimal solutions with 8 correspondences, and it goes beyond 95% for $N = 15$. We notice that the largest magnitude (25 $\texttt{m}$) has a lower percentage of detected optimal solutions.

Focal length: In this set of experiments, we let the image size fixed to 1900 $\texttt{pix}$ and vary the focal length, hence varying the FOV in practice. We let the focal length in $f \in \{300, 450, 600, 1000\}$ $\texttt{pix}$. Figure (4b) depicts the recall metrics for this configuration of parameters. The smaller focal length $f = 300$ has the worst performance, although our condition is able to detect more than 80% of optimal solutions with $N = 8$ and 96% for $N = 40$. This result

[5] We define precision as $\text{PRECISION} = \frac{\text{TP}}{\text{TP}+\text{FP}}$.

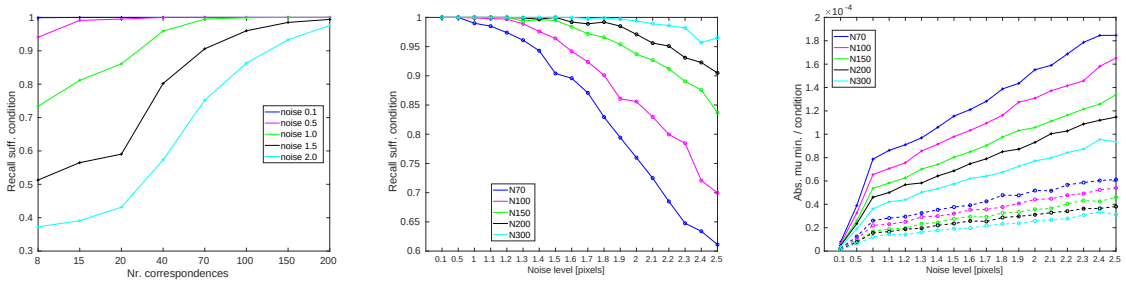

(a) (b) (c)

Fig. 3: (3a) and (3b) Recall metrics for the set of experiments with varying level of noise and number of correspondences. Figure (3a) shows the recall for problem instances with low number of correspondences N (X-axis) and noise up to 2.0 `pix` (see legend). Figure (3b) depicts the recall for N large (see legend) and level of noise from 0.1 to 2.5 `pix` (X-axis). (3c) Mean absolute value of the minimum eigenvalue of the Hessian (dashed lines) and approximation computed by our condition as $\eta(\boldsymbol{Q}_9, \boldsymbol{x})/\sigma^i_{\text{null}}(\boldsymbol{x}_0)$ (solid lines) for the noise and number of correspondences considered in (3b). The mean is always below the threshold $\epsilon_0 = 2e-04$.

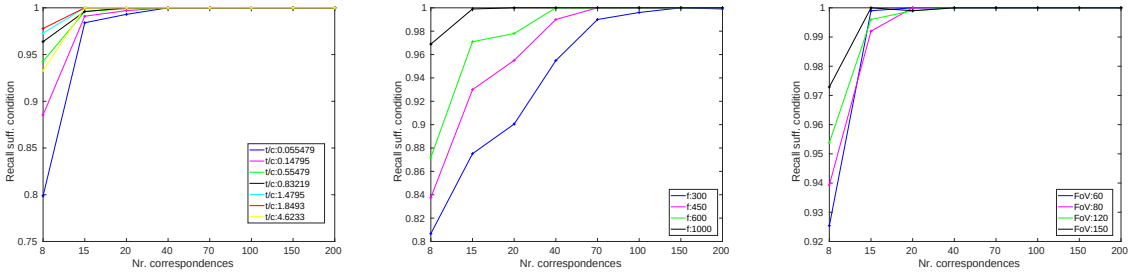

(a) (b) (c)

Fig. 4: Recall metrics for the set of experiments with varying (4a) translation magnitude; (4b) focal length; and (4c) FOV. Noise level is 0.5 `pix`. In (4b), the focal length changes the FOV, but not the image size (1900 `pix`); while in (4c) the image size changes the FOV, but not the focal length (800 `pix`).

agrees on how we are introducing the noise; the angles between the noiseless observation and its noisy version for these parameters are $0.0955, 0.0637, 0.0477, 0.0286$ degrees, respectively. For a focal length of 800 `pix`, these values correspond to noise levels of $1.3, 0.9, 0.67, 0.4$, respectively (*c.f.* Figure (3)). Hence, in practice we are only changing the signal-to-noise ratio.

Field of View: Last, we vary the FOV of both cameras. We let the focal length to 800 `pix` and vary the FOV in FOV $\in \{60, 80, 120, 150\}$ degrees, hence changing the image size. Figure (4c) depicts the precision (one for all the cases) and recall metrics for each combination of parameters. The narrower FOV attains the worst result, although we are still able to detect 92% of optimal solutions even with 8 correspondences; and with 40 correspondences, all the

optimal solutions are detected for all the tested FOV.

Computational cost: One of the greatest advantage of our proposal is the low computational cost incurred by the optimality certification: for each solution to a given problem, we only need to evaluate the polynomial $\eta(\boldsymbol{Q}_9, \boldsymbol{x})$ in Theorem (3.1). The computation of $\sigma_{\text{null}}(\boldsymbol{x}_0)$ is performed only once per problem instance and does not depend on the solution $\boldsymbol{x}$. We compare the computational cost of our proposal against that of the current faster certification algorithm [1]. Both approaches are implemented in MATLAB and run in a standard computer: CPU i7-4702MQ, 2.2GHz and 8 GB RAM. We only measure the cost of the certification stage for both approaches and report the average times for all the synthetic experiments described above. Our proposal has two steps: (1) computation of $\sigma^i_{\text{null}}(\boldsymbol{x}_0)$ (once, but up to six matrices) which takes 0.07996 ms per relaxation, that is, 0.47978 ms if the six relaxations are considered; and (2) evaluation of $\eta(\boldsymbol{Q}_9, \boldsymbol{x})$ for a given solution $\boldsymbol{x}$ (once per solution and problem, independent of relaxation) with 0.040015 ms. In total, 0.1199 ms when only one relaxation is employed and 0.52395 ms if the six relaxations are considered. The certification algorithm in [1] has also two steps: (1) estimation of the dual candidates (once per solution and relaxation) which requires 0.24532 ms; and (2) estimation of the minimum eigenvalue of the Hessian (once per solution and relaxation) with 0.14325 ms. In total 0.3886 ms per relaxation and solution. If the relaxation is not tight, one may need to compute another certifier for the same problem and solution. We observe that this is usually the case, and at least two certifiers must be computed. However, we do not know a priori which certifier will be capable to certify the solution, then, in the worst case, the six relaxations must be considered, which rises the time up to 2.4 ms, per solution. Notice also that if for the same problem we want to certify another solution, the certifiers need to be computed from scratch again, while the condition only requires to evaluate the polynomial, that is, with the condition the certification takes 0.040015 ms, while the certifier requires in the worst case 2.4 ms.

B. Evaluation on Real Data

To conclude this Section, we evaluate our proposal on real data. We sample images from the ETH3D dataset [36], the TUM dataset [37] and [38] (denoted by CVPR08). We list the sequences' names in the *Supplementary material* (G). For each pair of images, we extract and match SURF features [39]. The observations are normalized by the provided camera matrix with a pin-hole camera model. Since our proposal does not contemplate bad matches, *i.e*. outliers, we remove them with the provided ground truth. We discard all the correspondences whose squared epipolar error is larger than $2e-05$ w.r.t. the provided ground-truth relative pose. We employ at most 200 correspondences for each pair of images. For the ETH3D dataset our condition was able to detect 97.13% of the optimal solutions (w.r.t. the detected by the certifier), for TUM the percentage goes to 97.76% and for CVPR08, to 98.89%. Due to space limits, we provide the minimum eigenvalue of the Hessian (certifier) and our approximation as $\eta(\boldsymbol{Q}_9, \boldsymbol{x})/\sigma^i_{\text{null}}(\boldsymbol{x}_0)$ for each dataset and the recall metric for each sequence in the *Supplementary material* (H).

VI. Conclusion

In this work we have proposed a sufficient condition of optimality for the RPp between two central, calibrated cameras. The problem was posed as an optimization problem that minimizes the squared normalized epipolar error over the set of normalized essential matrices. The condition relied on the recently proposed certifiable algorithm in [1] and it was derived through spectral analysis. The form of the condition allows to detect and employ the best

relaxation in terms of certification. Further, we decouple the variables associated with the relaxation, the problem and the solution, which reduces the complexity of certifying different solutions for the same problem. We have shown with extensive experiments that the condition is tight in practice, and can certify optimal solutions in most problem instances even with high noise and low number of correspondences. The computational cost associated with the condition is also lower than the faster certification algorithm under a MATLAB implementation, being the condition between 3 and 4.5 times faster to evaluate. The main advantage, though, relies on its decoupling from the problem and its ability to detect the best relaxation, which makes it faster to certify solutions from the same problem, without requiring to re-compute everything from scratch.

Nonetheless, under some configurations the condition is not able to detect the optimal solutions. Given the tight relation between the parameterization of the set of essential matrices and the derived sufficient condition, we contemplate similar conditions derived from different certifiers associated with distinct minimal representations of said space. These conditions could potentially detect optimality when others do not. A more efficient implementation is also being considered, as well as the joint use of the certifier and the condition for a faster certification.

ACKNOWLEDGMENTS

We would like to thank two anonymous reviewers whose comments and suggestions helped clarify and improve the present manuscript. This work was supported by the grant program FPU18/01526 funded by the Spanish Government and the research projects WISER (DPI2017-84827-R) and ARPEGGIO (PID2020-117057).

REFERENCES

[1] Mercedes Garcia-Salguero, Jesus Briales, and Javier Gonzalez-Jimenez. Certifiable relative pose estimation. *Image and Vision Computing*, 109:104142, 2021.

[2] David Nistér, Oleg Naroditsky, and James Bergen. Visual odometry. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. Ieee, 2004.

[3] Davide Scaramuzza and Friedrich Fraundorfer. Visual odometry [tutorial]. *IEEE robotics & automation magazine*, 18(4):80–92, 2011.

[4] Ruben Gomez-Ojeda, Francisco-Angel Moreno, David Zuñiga-Noël, Davide Scaramuzza, and Javier Gonzalez-Jimenez. Pl-slam: A stereo slam system through the combination of points and line segments. *IEEE Transactions on Robotics*, 35(3):734–746, 2019.

[5] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015.

[6] Rana Azzam, Tarek Taha, Shoudong Huang, and Yahya Zweiri. Feature-based visual simultaneous localization and mapping: a survey. *SN Applied Sciences*, 2(2):1–24, 2020.

[7] Onur Ozyesil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion. *arXiv preprint arXiv:1701.08493*, 2017.

[8] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment—a modern synthesis. In *International workshop on vision algorithms*, pages 298–372. Springer, 1999.

[9] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.

[10] Vadim Indelman, Richard Roberts, and Frank Dellaert. Incremental light bundle adjustment for structure from motion and robotics. *Robotics and Autonomous Systems*, 70:63–82, 2015.

[11] Yi Ma, Jana Košecká, and Shankar Sastry. Optimization criteria and geometric algorithms for motion and structure estimation. *International Journal of Computer Vision*, 44(3):219–249, 2001.

[12] Uwe Helmke, Knut Hüper, Pei Yean Lee, and John Moore. Essential matrix estimation using gauss-newton iterations on a manifold. *International Journal of Computer Vision*, 74(2):117–136, 2007.

[13] Laurent Kneip and Simon Lynen. Direct optimization of frame-to-frame rotation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2352–2359, 2013.

[14] Henrik Stewenius, Christopher Engels, and David Nistér. Recent developments on direct relative orientation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 60(4):284–294, 2006.

[15] David Nistér. An efficient solution to the five-point relative pose problem. *IEEE transactions on pattern analysis and machine intelligence*, 26(6):756–770, 2004.

[16] Tom Botterill, Steven Mills, and Richard Green. Refining essential matrix estimates from ransac. In *Proceedings Image and Vision Computing New Zealand*, pages 1–6, 2011.

[17] Olivier D Faugeras and Steve Maybank. Motion from point matches: multiplicity of solutions. *International Journal of Computer Vision*, 4(3):225–246, 1990.

[18] Vincent Lui and Tom Drummond. An iterative 5-pt algorithm for fast and robust essential matrix estimation. In *BMVC*, 2013.

[19] Roberto Tron and Kostas Daniilidis. The space of essential matrices as a riemannian quotient manifold. *SIAM Journal on Imaging Sciences*, 10(3):1416–1445, 2017.

[20] Zuzana Kukelova and Tomas Pajdla. Two minimal problems for cameras with radial distortion. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.

[21] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Polynomial eigenvalue solutions to the 5-pt and 6-pt relative pose problems. In *BMVC*, volume 2, page 2008, 2008.

[22] H Christopher Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293(5828):133–135, 1981.

[23] Jesus Briales, Laurent Kneip, 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.

[24] Ji Zhao. An efficient solution to non-minimal case essential matrix estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[25] Richard I Hartley and Fredrik Kahl. Global optimization through searching rotation space and optimal estimation of the essential matrix. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.

[26] Kim-Chuan Toh, Michael J Todd, and Reha H Tütüncü. Sdpt3—a matlab software package for semidefinite programming, version 1.3. *Optimization methods and software*, 11(1-4):545–581, 1999.

[27] Jos F Sturm. Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. *Optimization methods and software*, 11(1-4):625–653, 1999.

[28] Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging and strong duality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 127–135, 2018.

[29] José Pedro Iglesias, Carl Olsson, and Fredrik Kahl. Global optimality for point set registration using semidefinite programming. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8287–8295, 2020.

[30] Richard Hartley and Yongduek Seo. Verifying global minima for l 2 minimization problems. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.

[31] Yi Ma, Stefano Soatto, Jana Kosecka, and S Shankar Sastry. *An invitation to 3-d vision: from images to geometric models*, volume 26. Springer Science & Business Media, 2012.

[32] Seong Hun Lee and Javier Civera. Geometric interpretations of the normalized epipolar error. *arXiv preprint arXiv:2008.01254*, 2020.

[33] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[34] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.

[35] Richard I Hartley. In defense of the eight-point algorithm. *IEEE Transactions on pattern analysis and machine intelligence*, 19(6):580–593, 1997.

[36] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3260–3269, 2017.

[37] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012.

[38] Christoph Strecha, Wolfgang Von Hansen, Luc Van Gool, Pascal Fua, and Ulrich Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. Ieee, 2008.

[39] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European conference on computer vision*, pages 404–417. Springer, 2006.

[40] Wei Dai and Yongsheng Ye. A characterization on singular value inequalities of matrices. *Journal of Function Spaces*, 2020, 2020.

[41] Chris Aholt, Sameer Agarwal, and Rekha Thomas. A qcqp approach to triangulation. In *European Conference on Computer Vision*, pages 654–667. Springer, 2012.