

A Quaternion-based Certifiably Optimal Solution to the Wahba Problem with Outliers

Heng Yang and Luca Carlone

Laboratory for Information & Decision Systems (LIDS)
Massachusetts Institute of Technology

{hankyang, lcarlone}@mit.edu

Abstract

The Wahba problem, also known as rotation search, seeks to find the best rotation to align two sets of vector observations given putative correspondences, and is a fundamental routine in many computer vision and robotics applications. This work proposes the first polynomial-time certifiably optimal approach for solving the Wahba problem when a large number of vector observations are outliers. Our first contribution is to formulate the Wahba problem using a Truncated Least Squares (TLS) cost that is insensitive to a large fraction of spurious correspondences. The second contribution is to rewrite the problem using unit quaternions and show that the TLS cost can be framed as a Quadratically-Constrained Quadratic Program (QCQP). Since the resulting optimization is still highly non-convex and hard to solve globally, our third contribution is to develop a convex Semidefinite Programming (SDP) relaxation. We show that while a naive relaxation performs poorly in general, our relaxation is tight even in the presence of large noise and outliers. We validate the proposed algorithm, named QUASAR (QUaternion-based Semidefinite relAXation for Robust alignment), in both synthetic and real datasets showing that the algorithm outperforms RANSAC, robust local optimization techniques, global outlier-removal procedures, and Branch-and-Bound methods. QUASAR is able to compute certifiably optimal solutions (i.e. the relaxation is exact) even in the case when 95% of the correspondences are outliers.

1. Introduction

The Wahba problem [52, 18], also known as rotation search [42, 27, 6], is a fundamental problem in computer vision, robotics, and aerospace engineering and consists in finding the rotation between two coordinate frames given vector observations taken in the two frames. The problem finds extensive applications in point cloud registration [8, 54], image stitching [6], motion estimation and 3D

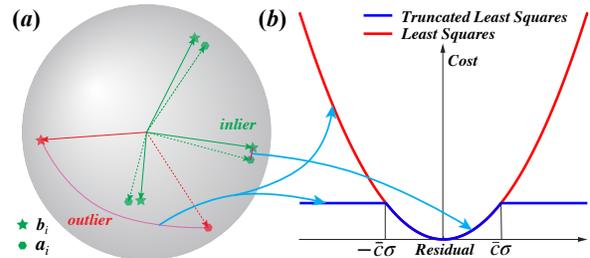


Figure 1. We propose QUASAR (QUaternion-based Semidefinite relaxation for Robust alignment), a certifiably optimal solution to the Wahba problem with outliers. (a) Wahba problem with four vector observations (three inliers and a single outlier). (b) Contrary to standard least squares formulations, QUASAR uses a truncated least squares cost that assigns a constant cost to measurements with large residuals, hence being insensitive to outliers.

reconstruction [10, 19, 35], and satellite attitude determination [52, 18, 15], to name a few.

Given two sets of vectors $\mathbf{a}_i, \mathbf{b}_i \in \mathbb{R}^3, i = 1, \dots, N$, the Wahba problem is formulated as a least squares problem

$$\min_{\mathbf{R} \in \text{SO}(3)} \sum_{i=1}^N w_i^2 \|\mathbf{b}_i - \mathbf{R}\mathbf{a}_i\|^2 \quad (1)$$

which computes the best rotation \mathbf{R} that aligns vectors \mathbf{a}_i and \mathbf{b}_i , and where $\{w_i^2\}_{i=1}^N$ are (known) weights associated to each pair of measurements. Here $\text{SO}(3) \doteq \{\mathbf{R} \in \mathbb{R}^{3 \times 3} : \mathbf{R}^T \mathbf{R} = \mathbf{R}\mathbf{R}^T = \mathbf{I}_3, \det(\mathbf{R}) = 1\}$ is the 3D *Special Orthogonal Group* containing proper 3D rotation matrices and \mathbf{I}_d denotes the identity matrix of size d . Problem (1) is known to be a maximum likelihood estimator for the unknown rotation when the ground-truth correspondences $(\mathbf{a}_i, \mathbf{b}_i)$ are known and the observations are corrupted with zero-mean isotropic Gaussian noise [48]. In other words, (1) computes an accurate estimate for \mathbf{R} when the observations can be written as $\mathbf{b}_i = \mathbf{R}\mathbf{a}_i + \epsilon_i$ ($i = 1, \dots, N$), where ϵ_i is isotropic Gaussian noise. Moreover, problem (1) can be solved in closed form [37, 28, 36, 3, 47, 29].

Unfortunately, in practical application, many of the vector observations may be *outliers*, typically due to incorrect vector-to-vector correspondences, cf. Fig. 1(a). In com-

puter vision, correspondences are established through 2D (e.g., SIFT [33], ORB [44]) or 3D (e.g., FPFH [45]) feature matching techniques, which are prone to produce many outlier correspondences. For instance, it is not uncommon to observe 95% outliers when using FPFH for point cloud registration [13]. In the presence of outliers, the standard Wahba problem (1) is no longer a maximum likelihood estimator and the resulting estimates are affected by large errors [13]. A common approach to gain robustness against outliers is to incorporate solvers for problem (1) in a RANSAC scheme [22]. However, RANSAC’s runtime grows exponentially with the outlier ratio [13] and its performance, as we will see in Section 6, quickly degrades in the presence of noise and high outlier ratios.

This paper is motivated by the goal of designing an approach that (i) can solve the Wahba problem globally (without an initial guess), (ii) can tolerate large noise (e.g., the magnitude of the noise is 10% of the magnitude of the measurement vector) and extreme amount of outliers (e.g., over 95% of the observations are outliers), (iii) runs in polynomial time, and (iv) provides certifiably optimal solutions. The related literature, reviewed in Section 2, fails to simultaneously meet these goals, and only includes algorithms that are robust to moderate amounts of outlier, or are robust to extreme amounts of outliers (e.g., 90%) but only provide sub-optimal solutions (e.g., GORE [42]), or that are globally optimal but run in exponential time in the worst case, such as *branch-and-bound* (BnB) methods (e.g., [27, 7]).

Contribution. Our first contribution, presented in Section 3, is to reformulate the Wahba problem (1) using a *Truncated Least Squares* (TLS) cost that is robust against a large fraction of outliers. We name the resulting optimization problem the (outlier-)Robust Wahba problem.

The second contribution (Section 4) is to depart from the rotation matrix representation and rewrite the Robust Wahba problem using unit quaternions. In addition, we show how to rewrite the TLS cost function by using additional binary variables that decide whether a measurement is an inlier or outlier. Finally, we prove that the mixed-integer program (including a quaternion and N binary variables) can be rewritten as an optimization involving $N + 1$ unit quaternions. This sequence of re-parametrizations, that we call *binary cloning*, leads to a non-convex Quadratically-Constrained Quadratic Program (QCQP).

The third contribution (Section 5) is to provide a polynomial-time certifiably optimal solver for the QCQP. Since the QCQP is highly non-convex and hard to solve globally, we propose a *Semidefinite Programming* (SDP) *relaxation* that is empirically tight. We show that while a naive SDP relaxation is not tight and performs poorly in practice, the proposed method remains tight even when observing 95% outliers. Our approach is *certifiably optimal* [4] in the sense that it provides a way to check optimal-

ity of the resulting solution, and computes optimal solutions in practical problems. To the best of our knowledge, this is the first polynomial-time method that solves the robust Wahba problem with certifiable optimality guarantees.

We validate the proposed algorithm, named QUASAR (QUaternion-based Semidefinite relAXation for Robust alignment), in both synthetic and real datasets for point cloud registration and image stitching, showing that the algorithm outperforms RANSAC, robust local optimization techniques, outlier-removal procedures, and BnB methods.

2. Related Work

2.1. Wahba problem without outliers

The Wahba problem was first proposed by Grace Wahba in 1965 with the goal of estimating satellite attitude given vector observations [52]. In aerospace, the vector observations are typically the directions to visible stars observed by sensors onboard the satellite. The Wahba problem (1) is related to the well-known *Orthogonal Procrustes* problem [24] where one searches for orthogonal matrices (rather than rotations) and all the weights w_i are set to be equal to one. Schonemann [47] provides a closed-form solution for the Orthogonal Procrustes problem using singular value decomposition. Subsequent research effort across multiple communities led to the derivation of closed-form solutions for the Wahba problem (1) using both quaternion [28, 37] and rotation matrix [29, 36, 3, 23, 46, 31] representations.

The computer vision community has investigated the Wahba problem in the context of point cloud registration [8, 54], image stitching [6], motion estimation and 3D reconstruction [10, 19]. In particular, the closed-form solutions from Horn [28] and Arun *et al.* [3] are now commonly used in point cloud registration techniques [8]. While the formulation (1) implicitly assumes zero-mean isotropic Gaussian noise, several authors investigate a generalized version of the Wahba problem (1), where the noise follows an anisotropic Gaussian [12, 15, 30]. Cheng and Crasidis [15] develop a local iterative optimization algorithm. Briaies and Gonzalez-Jimenez [12] propose a convex relaxation to compute global solutions for the anisotropic case. Ahmed *et al.* [1] develop an SDP relaxation for the case with bounded noise and no outliers. All these methods are known to perform poorly in the presence of outliers.

2.2. Wahba problem with outliers

In computer vision applications, the vector-to-vector correspondences are typically established through descriptor matching, which may lead to spurious correspondences and outliers [42, 13]. This observation triggered research into outlier-robust variants of the Wahba problem.

Local Methods. The most widely used method for handling outliers is RANSAC, which is efficient and accurate in

the low-noise and low-outlier regime [22, 38]. However, RANSAC is non-deterministic (different runs give different solutions), sub-optimal (the solution may not capture all the inlier measurements), and its performance quickly deteriorates in the large-noise and high-outlier regime. Other approaches resort to *M-estimators*, which replace the least squares cost in eq. (1) with robust cost functions that are less sensitive to outliers [9, 34]. Zhou *et al.* [56] propose *Fast Global Registration* (FGR) that employs the Geman-McClure robust cost function and solves the resulting optimization iteratively using a continuation method. Since the problem becomes more and more non-convex at each iteration, FGR does not guarantee global optimality in general. In fact, as we show in Section 6, FGR tends to fail when the outlier ratio is above 70% of the observations.

Global Methods. The most popular class of global methods for robust rotation search is based on *Consensus Maximization* [17] and *branch-and-bound* (BnB) [14]. Hartley and Kahl [27] first proposed using BnB for rotation search, and Bazin *et al.* [7] adopted consensus maximization to extend their BnB algorithm with a robust formulation. BnB is guaranteed to return the globally optimal solution, but it runs in exponential time in the worst case. Another class of global methods for consensus maximization enumerates all possible subsets of measurements with size no larger than the problem dimension (3 for rotation search) to analytically compute candidate solutions, and then verify global optimality using computational geometry [41, 21]. These methods still require exhaustive enumeration.

Outlier-removal Methods. Recently, Para and Chin [42] proposed a *guaranteed outlier removal* (GORE) algorithm that removes gross outliers while ensuring that all inliers are preserved. Using GORE as a preprocessing step for BnB is shown to boost the speed of BnB, while using GORE alone still provides a reasonable sub-optimal solution. Besides rotation search, GORE has also been successfully applied to other geometric vision problems such as triangulation [16], and registration [42, 13]. Nevertheless, the existing literature is still missing a polynomial-time algorithm that can simultaneously tolerate extreme amounts of outliers and return globally optimal solutions.

3. Problem Formulation: Robust Wahba

Let $\mathcal{A} = \{\mathbf{a}_i\}_{i=1}^N$ and $\mathcal{B} = \{\mathbf{b}_i\}_{i=1}^N$ be two sets of 3D vectors ($\mathbf{a}_i, \mathbf{b}_i \in \mathbb{R}^3$), such that, given \mathcal{A} , the vectors in \mathcal{B} are described by the following generative model:

$$\begin{cases} \mathbf{b}_i = \mathbf{R}\mathbf{a}_i + \boldsymbol{\epsilon}_i & \text{if } \mathbf{b}_i \text{ is an inlier, or} \\ \mathbf{b}_i = \mathbf{o}_i & \text{if } \mathbf{b}_i \text{ is an outlier} \end{cases} \quad (2)$$

where $\mathbf{R} \in \text{SO}(3)$ is an (unknown, to-be-estimated) rotation matrix, $\boldsymbol{\epsilon}_i \in \mathbb{R}^3$ models the inlier measurement noise, and $\mathbf{o}_i \in \mathbb{R}^3$ is an arbitrary vector. In other words, if \mathbf{b}_i

is an inlier, then it must be a rotated version of \mathbf{a}_i plus noise, while if \mathbf{b}_i is an outlier, then \mathbf{b}_i is arbitrary. In the special case where all \mathbf{b}_i 's are inliers and the noise obeys a zero-mean isotropic Gaussian distribution, *i.e.* $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}_3, \sigma_i^2 \mathbf{I}_3)$, then the Maximum Likelihood estimator of \mathbf{R} takes the form of eq. (1), with the weights chosen as the inverse of the measurement variances, $w_i^2 = 1/\sigma_i^2$. In this paper, we are interested in the case where measurements include outliers.

3.1. The Robust Wahba Problem

We introduce a novel formulation for the (outlier-)Robust Wahba problem, that uses a *truncated least squares* (TLS) cost function:

$$\min_{\mathbf{R} \in \text{SO}(3)} \sum_{i=1}^N \min \left(\frac{1}{\sigma_i^2} \|\mathbf{b}_i - \mathbf{R}\mathbf{a}_i\|^2, \bar{c}^2 \right) \quad (3)$$

where the inner “ $\min(\cdot, \cdot)$ ” returns the minimum between two scalars. The TLS cost has been recently shown to be robust against high outlier rates in pose graph optimization [32]. Problem (3) computes a least squares solution for measurements with small residuals, *i.e.* when $\frac{1}{\sigma_i^2} \|\mathbf{b}_i - \mathbf{R}\mathbf{a}_i\|^2 \leq \bar{c}^2$ (or, equivalently, $\|\mathbf{b}_i - \mathbf{R}\mathbf{a}_i\|^2 \leq \sigma_i^2 \bar{c}^2$), while discarding measurements with large residuals (when $\frac{1}{\sigma_i^2} \|\mathbf{b}_i - \mathbf{R}\mathbf{a}_i\|^2 > \bar{c}^2$ the i -th term becomes a constant \bar{c}^2 and has no effect on the optimization). Problem (3) implements the TLS cost function illustrated in Fig. 1(b). The parameters σ_i and \bar{c}^2 are fairly easy to set in practice, as discussed in the following remarks.

Remark 1 (Probabilistic choice of σ_i and \bar{c}^2). *Let us assume that the inliers follow the generative model (2) with $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}_3, \sigma_i^2 \mathbf{I}_3)$; we will not make assumptions on the generative model for the outliers, which is unknown in practice. Since the noise on the inliers is Gaussian, it holds:*

$$\frac{1}{\sigma_i^2} \|\mathbf{b}_i - \mathbf{R}\mathbf{a}_i\|^2 = \frac{1}{\sigma_i^2} \|\boldsymbol{\epsilon}_i\|^2 \sim \chi^2(3) \quad (4)$$

where $\chi^2(3)$ is the Chi-squared distribution with three degrees of freedom. Therefore, with desired probability p , the weighted error $\frac{1}{\sigma_i^2} \|\boldsymbol{\epsilon}_i\|^2$ for the inliers satisfies:

$$\mathbb{P} \left(\frac{\|\boldsymbol{\epsilon}_i\|^2}{\sigma_i^2} \leq \bar{c}^2 \right) = p, \quad (5)$$

where \bar{c}^2 is the quantile of the χ^2 distribution with three degrees of freedom and lower tail probability equal to p . Therefore, one can simply set the σ_i in Problem (3) to be the standard deviation of the inlier noise, and compute \bar{c}^2 from the $\chi^2(3)$ distribution for a desired probability p (e.g., $p = 0.99$). The constant \bar{c}^2 monotonically increases with p ; therefore, setting p close to 1 makes the formulation (3) more prone to accept measurements with large residuals, while a small p makes (3) more selective.

Remark 2 (Set membership choice of σ_i and \bar{c}^2). Let us assume that the inliers follow the generative model (2) with $\|\epsilon_i\| \leq \beta_i$, where β_i is a given noise bound; this is the typical setup assumed in set membership estimation [40]. In this case, it is easy to see that the inliers satisfy:

$$\|\epsilon_i\| \leq \beta_i \iff \|\mathbf{b}_i - \mathbf{R}\mathbf{a}_i\|^2 \leq \beta_i^2 \quad (6)$$

hence one can simply choose $\sigma_i^2 \bar{c}^2 = \beta_i^2$. Intuitively, the constant $\sigma_i^2 \bar{c}^2$ is the largest (squared) residual error we are willing to tolerate on the i -th measurement.

3.2. Quaternion formulation

We now adopt a quaternion formulation for (3). Quaternions are an alternative representation for 3D rotations [49, 11] and their use will simplify the derivation of our convex relaxation in Section 5. We start by reviewing basic facts about quaternions and then state the quaternion-based Robust Wahba formulation in Problem 1 below.

Preliminaries on Unit Quaternions. We denote a unit quaternion as a unit-norm column vector $\mathbf{q} = [\mathbf{v}^\top \ s]^\top$, where $\mathbf{v} \in \mathbb{R}^3$ is the *vector part* of the quaternion and the last element s is the *scalar part*. We also use $\mathbf{q} = [q_1 \ q_2 \ q_3 \ q_4]^\top$ to denote the four entries of the quaternion.

Each quaternion represents a 3D rotation and the composition of two rotations \mathbf{q}_a and \mathbf{q}_b can be computed using the *quaternion product* $\mathbf{q}_c = \mathbf{q}_a \otimes \mathbf{q}_b$:

$$\mathbf{q}_c = \mathbf{q}_a \otimes \mathbf{q}_b = \Omega_1(\mathbf{q}_a)\mathbf{q}_b = \Omega_2(\mathbf{q}_b)\mathbf{q}_a, \quad (7)$$

where $\Omega_1(\mathbf{q})$ and $\Omega_2(\mathbf{q})$ are defined as follows:

$$\Omega_1(\mathbf{q}) = \begin{bmatrix} q_4 & -q_3 & q_2 & q_1 \\ q_3 & q_4 & -q_1 & q_2 \\ -q_2 & q_1 & q_4 & q_3 \\ -q_1 & -q_2 & -q_3 & q_4 \end{bmatrix}, \quad \Omega_2(\mathbf{q}) = \begin{bmatrix} q_4 & q_3 & -q_2 & q_1 \\ -q_3 & q_4 & q_1 & q_2 \\ q_2 & -q_1 & q_4 & q_3 \\ -q_1 & -q_2 & -q_3 & q_4 \end{bmatrix}. \quad (8)$$

The inverse of a quaternion $\mathbf{q} = [\mathbf{v}^\top \ s]^\top$ is defined as:

$$\mathbf{q}^{-1} = \begin{bmatrix} -\mathbf{v} \\ s \end{bmatrix}, \quad (9)$$

where one simply reverses the sign of the vector part.

The rotation of a vector $\mathbf{a} \in \mathbb{R}^3$ can be expressed in terms of quaternion product. Formally, if \mathbf{R} is the (unique) rotation matrix corresponding to a unit quaternion \mathbf{q} , then:

$$\begin{bmatrix} \mathbf{R}\mathbf{a} \\ 0 \end{bmatrix} = \mathbf{q} \otimes \hat{\mathbf{a}} \otimes \mathbf{q}^{-1}, \quad (10)$$

where $\hat{\mathbf{a}} = [\mathbf{a}^\top \ 0]^\top$ is the homogenization of \mathbf{a} , obtained by augmenting \mathbf{a} with an extra entry equal to zero.

The set of unit quaternions, denoted as $\mathcal{S}^3 = \{\mathbf{q} \in \mathbb{R}^4 : \|\mathbf{q}\| = 1\}$, is the *3-Sphere* manifold. \mathcal{S}^3 is a *double cover* of $\text{SO}(3)$ since \mathbf{q} and $-\mathbf{q}$ represent the same rotation (intuitively, eq. (10) implements the same rotation if we replace \mathbf{q} with $-\mathbf{q}$, since the matrices in (8) are linear in \mathbf{q}).

Quaternion-based Robust Wahba Problem. Equipped with the relations reviewed above, it is now easy to rewrite (3) using unit quaternions.

Problem 1 (Quaternion-based Robust Wahba). *The robust Wahba problem (3) can be equivalently written as:*

$$\min_{\mathbf{q} \in \mathcal{S}^3} \sum_{i=1}^N \min \left(\frac{1}{\sigma_i^2} \|\hat{\mathbf{b}}_i - \mathbf{q} \otimes \hat{\mathbf{a}}_i \otimes \mathbf{q}^{-1}\|^2, \bar{c}^2 \right), \quad (11)$$

where we defined $\hat{\mathbf{a}}_i \doteq [\mathbf{a}_i^\top \ 0]^\top$ and $\hat{\mathbf{b}}_i \doteq [\mathbf{b}_i^\top \ 0]^\top$, and \otimes denotes the quaternion product.

The equivalence between (11) and (3) can be easily understood from eq. (10). The main advantage of using (11) is that we replaced the set $\text{SO}(3)$ with a simpler set, the set of unit-norm vectors \mathcal{S}^3 .

4. Binary Cloning and QCQP

The goal of this section is to rewrite (11) as a Quadratically-Constrained Quadratic Program (QCQP). We do so in three steps: (i) we show that (11) can be written using binary variables, (ii) we show that the problem with binary variables can be written using $N + 1$ quaternions, and (iii) we manipulate the resulting problem to expose its quadratic nature. The derivation of the QCQP will pave the way to our convex relaxation (Section 5).

From Truncated Least Squares to Mixed-Integer Programming. Problem (11) is hard to solve globally, due to the non-convexity of both the cost function and the constraint (\mathcal{S}^3 is a non-convex set). As a first reparametrization, we expose the non-convexity of the cost by rewriting the TLS cost using binary variables. Towards this goal, we rewrite the inner “min” in (11) using the following property, that holds for any pair of scalars x and y :

$$\min(x, y) = \min_{\theta \in \{+1, -1\}} \frac{1 + \theta}{2} x + \frac{1 - \theta}{2} y. \quad (12)$$

Eq. (12) can be verified to be true by inspection: the right-hand-side returns x (with minimizer $\theta = +1$) if $x < y$, and y (with minimizer $\theta = -1$) if $x > y$. This enables us to rewrite problem (11) as a mixed-integer program including the quaternion \mathbf{q} and binary variables θ_i , $i = 1, \dots, N$:

$$\min_{\substack{\mathbf{q} \in \mathcal{S}^3 \\ \theta_i \in \{\pm 1\}}} \sum_{i=1}^N \frac{1 + \theta_i}{2} \frac{\|\hat{\mathbf{b}}_i - \mathbf{q} \otimes \hat{\mathbf{a}}_i \otimes \mathbf{q}^{-1}\|^2}{\sigma_i^2} + \frac{1 - \theta_i}{2} \bar{c}^2. \quad (13)$$

The reformulation is related to the Black-Rangarajan duality between robust estimation and line processes [9]: the TLS cost is an extreme case of robust function that results in a binary line process. Intuitively, the binary variables $\{\theta_i\}_{i=1}^N$ in problem (13) decide whether a given measurement i is an inlier ($\theta_i = +1$) or an outlier ($\theta_i = -1$).

From Mixed-Integer to Quaternions. Now we convert the mixed-integer program (13) to an optimization over $N + 1$ quaternions. The intuition is that, if we define extra quaternions $\mathbf{q}_i \doteq \theta_i \mathbf{q}$, we can rewrite (13) as a function of

\mathbf{q} and \mathbf{q}_i ($i = 1, \dots, N$). This is a key step towards getting a quadratic cost (Proposition 4). The re-parametrization is formalized in the following proposition.

Proposition 3 (Binary cloning). *The mixed-integer program (13) is equivalent (in the sense that they admit the same optimal solution \mathbf{q}) to the following optimization*

$$\min_{\substack{\mathbf{q} \in S^3 \\ \mathbf{q}_i \in \{\pm \mathbf{q}\}}} \sum_{i=1}^N \frac{\|\hat{\mathbf{b}}_i - \mathbf{q} \otimes \hat{\mathbf{a}}_i \otimes \mathbf{q}^{-1} + \mathbf{q}^\top \mathbf{q}_i \hat{\mathbf{b}}_i - \mathbf{q} \otimes \hat{\mathbf{a}}_i \otimes \mathbf{q}_i^{-1}\|^2}{4\sigma_i^2} + \frac{1 - \mathbf{q}^\top \mathbf{q}_i}{2} \bar{c}^2. \quad (14)$$

which involves $N + 1$ quaternions (\mathbf{q} and \mathbf{q}_i , $i = 1, \dots, N$).

While a formal proof is given in the Supplementary Material, it is fairly easy to see that if $\mathbf{q}_i = \{\pm \mathbf{q}\}$, or equivalently, $\mathbf{q}_i = \theta_i \mathbf{q}$ with $\theta_i \in \{\pm 1\}$, then $\mathbf{q}_i^\top \mathbf{q} = \theta_i$, and $\mathbf{q} \otimes \hat{\mathbf{a}}_i \otimes \mathbf{q}_i^{-1} = \theta_i (\mathbf{q} \otimes \hat{\mathbf{a}}_i \otimes \mathbf{q}^{-1})$ which exposes the relation between (13) and (14). We dubbed the re-parametrization (14) *binary cloning* since now we created a ‘‘clone’’ \mathbf{q}_i for each measurement, such that $\mathbf{q}_i = \mathbf{q}$ for inliers (recall that $\mathbf{q}_i = \theta_i \mathbf{q}$) and $\mathbf{q}_i = -\mathbf{q}$ for outliers.

From Quaternions to QCQP. We conclude this section by showing that (14) can be actually written as a QCQP. This observation is non-trivial since (14) has a *quartic* cost and $\mathbf{q}_i = \{\pm \mathbf{q}\}$ is not in the form of a quadratic constraint. The re-formulation as a QCQP is given in the following.

Proposition 4 (Binary Cloning as a QCQP). *Define a single column vector $\mathbf{x} = [\mathbf{q}^\top \mathbf{q}_1^\top \dots \mathbf{q}_N^\top]^\top$ stacking all variables in Problem (14). Then, Problem (14) is equivalent (in the sense that they admit the same optimal solution \mathbf{q}) to the following Quadratically-Constrained Quadratic Program:*

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^{4(N+1)}} & \sum_{i=1}^N \mathbf{x}^\top \mathbf{Q}_i \mathbf{x} \\ \text{subject to} & \mathbf{x}_q^\top \mathbf{x}_q = 1 \\ & \mathbf{x}_{q_i} \mathbf{x}_{q_i}^\top = \mathbf{x}_q \mathbf{x}_q^\top, \forall i = 1, \dots, N \end{aligned} \quad (15)$$

where $\mathbf{Q}_i \in \mathbb{R}^{4(N+1) \times 4(N+1)}$ ($i = 1, \dots, N$) are known symmetric matrices that depend on the 3D vectors \mathbf{a}_i and \mathbf{b}_i (the explicit expression is given in the Supplementary Material), and the notation \mathbf{x}_q (resp. \mathbf{x}_{q_i}) denotes the 4D subvector of \mathbf{x} corresponding to \mathbf{q} (resp. \mathbf{q}_i).

A complete proof of Proposition 4 is given in the Supplementary Material. Intuitively, (i) we developed the squares in the cost function (14), (ii) we used the properties of unit quaternions (Section 3) to simplify the expression to a quadratic cost, and (iii) we adopted the more compact notation afforded by the vector \mathbf{x} to obtain (15).

5. Semidefinite Relaxation

Problem (15) writes the Robust Wahba problem as a QCQP. Problem (15) is still a non-convex problem

(quadratic equality constraints are non-convex). Here we develop a convex semidefinite programming (SDP) relaxation for problem (15).

The crux of the relaxation consists in rewriting problem (15) as a function of the following matrix:

$$\mathbf{Z} = \mathbf{x} \mathbf{x}^\top = \begin{bmatrix} \mathbf{q} \mathbf{q}^\top & \mathbf{q} \mathbf{q}_1^\top & \dots & \mathbf{q} \mathbf{q}_N^\top \\ \star & \mathbf{q}_1 \mathbf{q}_1^\top & \dots & \mathbf{q}_1 \mathbf{q}_N^\top \\ \vdots & \vdots & \ddots & \vdots \\ \star & \star & \dots & \mathbf{q}_N \mathbf{q}_N^\top \end{bmatrix}. \quad (16)$$

For this purpose we note that if we define $\mathbf{Q} \doteq \sum_{i=1}^N \mathbf{Q}_i$:

$$\sum_{i=1}^N \mathbf{x}^\top \mathbf{Q}_i \mathbf{x} = \mathbf{x}^\top \mathbf{Q} \mathbf{x} = \text{tr}(\mathbf{Q} \mathbf{x} \mathbf{x}^\top) = \text{tr}(\mathbf{Q} \mathbf{Z}) \quad (17)$$

and that $\mathbf{x}_{q_i} \mathbf{x}_{q_i}^\top = [\mathbf{Z}]_{q_i q_i}$, where $[\mathbf{Z}]_{q_i q_i}$ denotes the 4×4 diagonal block of \mathbf{Z} with row and column indices corresponding to \mathbf{q}_i . Since any matrix in the form $\mathbf{Z} = \mathbf{x} \mathbf{x}^\top$ is a positive-semidefinite rank-1 matrix, we obtain:

Proposition 5 (Matrix Formulation of Binary Cloning). *Problem (15) is equivalent (in the sense that optimal solutions of a problem can be mapped to optimal solutions of the other) to the following non-convex program:*

$$\begin{aligned} \min_{\mathbf{Z} \succeq 0} & \text{tr}(\mathbf{Q} \mathbf{Z}) \\ \text{subject to} & \text{tr}([\mathbf{Z}]_{qq}) = 1 \\ & [\mathbf{Z}]_{q_i q_i} = [\mathbf{Z}]_{qq}, \forall i = 1, \dots, N \\ & \text{rank}(\mathbf{Z}) = 1 \end{aligned} \quad (18)$$

At this point it is straightforward to develop a (standard) SDP relaxation by dropping the rank-1 constraint, which is the only source of non-convexity in (18).

Proposition 6 (Naive SDP Relaxation). *The following SDP is a convex relaxation of Problem (18):*

$$\begin{aligned} \min_{\mathbf{Z} \succeq 0} & \text{tr}(\mathbf{Q} \mathbf{Z}) \\ \text{subject to} & \text{tr}([\mathbf{Z}]_{qq}) = 1 \\ & [\mathbf{Z}]_{q_i q_i} = [\mathbf{Z}]_{qq}, \forall i = 1, \dots, N \end{aligned} \quad (19)$$

The following theorem proves that the naive SDP relaxation (6) is tight in the absence of noise and outliers.

Theorem 7 (Tightness in Noiseless and Outlier-free Wahba). *When there is no noise and no outliers in the measurements, and there are at least two vector measurements (\mathbf{a}_i 's) that are not parallel to each other, the SDP relaxation (19) is always tight, i.e.:*

1. the optimal cost of (19) matches the optimal cost of the QCQP (15),
2. the optimal solution \mathbf{Z}^* of (19) has rank 1, and

3. \mathbf{Z}^* can be written as $\mathbf{Z}^* = (\mathbf{x}^*)(\mathbf{x}^*)^\top$ where $\mathbf{x}^* \doteq [(\mathbf{q}^*)^\top (\mathbf{q}_1^*)^\top \dots (\mathbf{q}_N^*)^\top]$ is a global minimizer of the original non-convex problem (15).

A formal proof, based on Lagrangian duality theory, is given in the Supplementary Material. While Theorem 7 ensures that the naive SDP relaxation computes optimal solutions in noiseless and outlier-free problems, our original motivation was to solve problems with many outliers. One can still empirically assess tightness by solving the SDP and verifying if a rank-1 solution is obtained. Unfortunately, the naive relaxation produces solutions with rank larger than 1 in the presence of outliers, cf. Fig. 2(a). Even when the rank is larger than 1, one can *round* the solution by computing a rank-1 approximation of \mathbf{Z}^* ; however, we empirically observe that, whenever the rank is larger than 1, the rounded estimates exhibit large errors, as shown in Fig. 2(b).

To address these issues, we propose to add *redundant constraints* to *tighten* (improve) the SDP relaxation, inspired by [12, 51]. The following proposed relaxation is tight (empirically satisfies the three claims of Theorem 7) even in the presence of noise and 95% of outliers.

Proposition 8 (SDP Relaxation with Redundant Constraints). *The following SDP is a convex relaxation of (18):*

$$\begin{aligned} \min_{\mathbf{Z} \succeq 0} \quad & \text{tr}(\mathbf{Q}\mathbf{Z}) & (20) \\ \text{subject to} \quad & \text{tr}([\mathbf{Z}]_{qq}) = 1 \\ & [\mathbf{Z}]_{q_i q_i} = [\mathbf{Z}]_{qq}, \forall i = 1, \dots, N \\ & [\mathbf{Z}]_{q_i q_i} = [\mathbf{Z}]_{qq_i}^\top, \forall i = 1, \dots, N \\ & [\mathbf{Z}]_{q_i q_j} = [\mathbf{Z}]_{q_i q_j}^\top, \forall 1 \leq i < j \leq N \end{aligned}$$

and it is always tighter, i.e. the optimal objective of (20) is always closer to the optimal objective of (15), when compared to the naive relaxation (19).

We name the improved relaxation (20) QUASAR (QUATernion-based Semidefinite relAXation for Robust alignment). While our current theoretical results only guarantee tightness in the noiseless and outlier-free case (Theorem 7 also holds for QUASAR, since it is always tighter than the naive relaxation), in the next section we empirically demonstrate the tightness of (20) in the face of noise and extreme outlier rates, cf. Fig. 2.

6. Experiments

We evaluate QUASAR in both synthetic and real datasets for point cloud registration and image stitching showing that (i) the proposed relaxation is tight even with extreme (95%) outlier rates, and (ii) QUASAR is more accurate and robust than state-of-the-art techniques for rotation search.

We implemented QUASAR in Matlab, using cvx [25] to model the convex programs (19) and (20), and used MOSEK [2] as the SDP solver. The parameters in QUASAR are set according to Remark 1 with $p = 1 - 10^{-4}$.

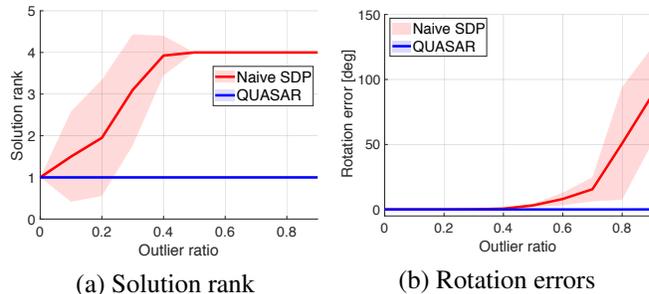


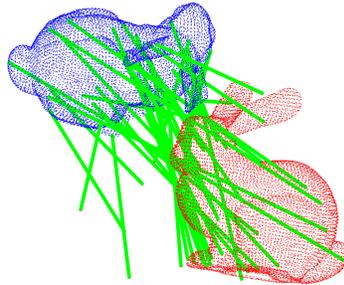
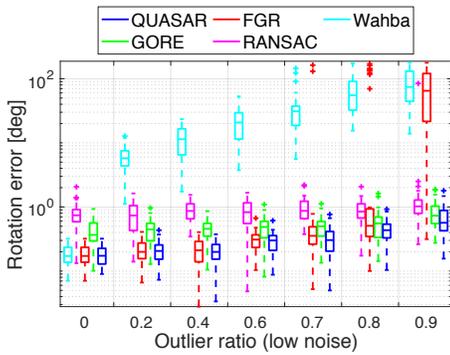
Figure 2. Comparison between the naive relaxation (19) and the proposed relaxation (20) for increasing outlier percentages. (a) Rank of solution (relaxation is tight when rank is 1); (b) rotation errors (solid line: mean; shaded area: 1-sigma standard deviation).

6.1. Evaluation on Synthetic Datasets

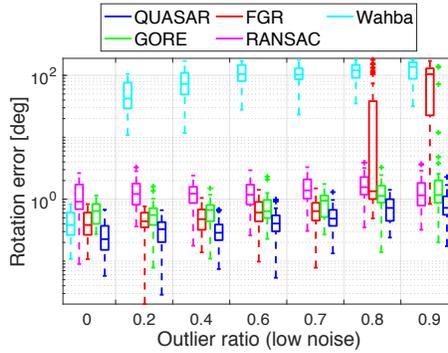
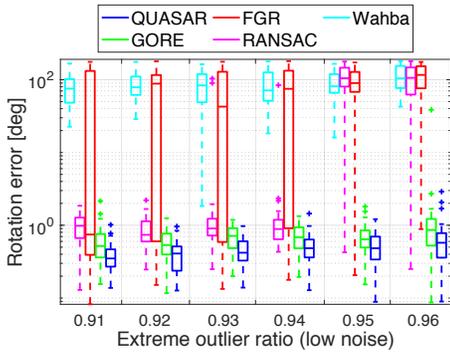
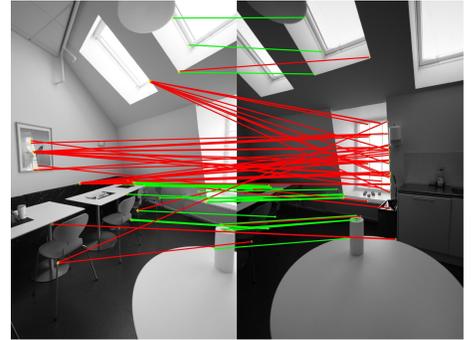
Comparison against the Naive SDP Relaxation. We first test the performance of the naive SDP relaxation (19) and QUASAR (20) under zero noise and increasing outlier rates. In each test, we sample $N = 40$ unit-norm vectors $\mathcal{A} = \{\mathbf{a}_i\}_{i=1}^N$ uniformly at random. Then we apply a random rotation \mathbf{R} to \mathcal{A} according to (2) with no additive noise to get $\mathcal{B} = \{\mathbf{b}_i\}_{i=1}^N$. To generate outliers, we replace a fraction of \mathbf{b}_i 's with random unit-norm vectors. Results are averaged over 40 Monte Carlo runs. Fig. 2(a) shows the rank of the solution produced by the naive SDP relaxation (19) and QUASAR (20); recall that the relaxation is tight when the rank is 1. Both relaxations are tight when there are no outliers, which validates Theorem 7. However, the performance of the naive relaxation quickly degrades as the measurements include more outliers: at 10 – 40%, the naive relaxation starts becoming loose and completely breaks when the outlier ratio is above 40%. On the other hand, QUASAR produces a certifiably optimal rank-1 solution even with 90% outliers. Fig. 2(b) confirms that a tighter relaxation translates into more accurate rotation estimates.

Comparison against the State of the Art. Using the setup described in the previous section, we test the performance of QUASAR in the presence of noise and outliers and benchmark it against (i) the closed-form solution [29] to the standard Wahba problem (1) (label: Wahba); (ii) RANSAC with 1000 maximum iteration (label: RANSAC); (iii) *Fast Global Registration* [56] (label: FGR); (iv) *Guaranteed Outlier Removal* [42] (label: GORE). We used default parameters for all approaches. We also benchmarked QUASAR against BnB methods [7, 42], but we found GORE had similar or better performance than BnB. Therefore, the comparison with BnB is presented in the Supplementary Material.

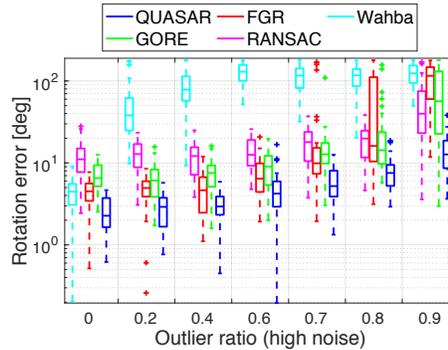
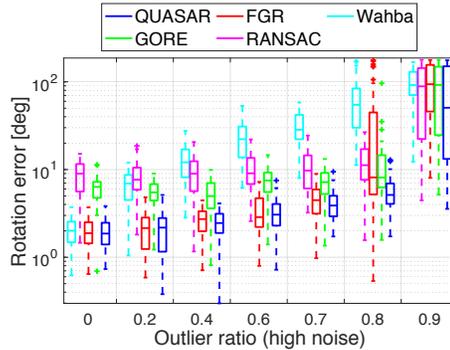
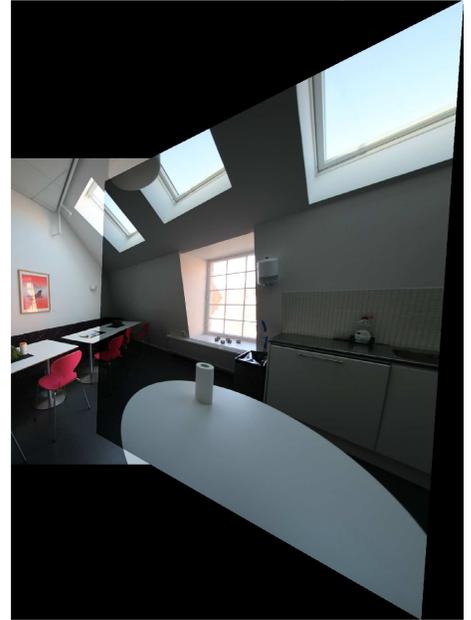
Fig. 3(a, first row) shows the box-plot of the rotation errors produced by the compared techniques for increasing ratios (0 – 90%) of outliers when the inlier noise is low ($\sigma_i = 0.01, i = 1, \dots, N$). As expected, Wahba only produces reasonable results in the outlier-free case. FGR is robust against 70% outliers but breaks at 90% (many tests re-



Input images and SURF correspondences (Green: inliers, Red: outliers)



QUASAR stitched image



(a) Synthetic datasets

(b) Point cloud registration

(c) Image stitching

Figure 3. (a) Rotation errors for increasing levels of outliers in synthetic datasets with low and high inlier noise. (b) Rotation errors on the Bunny dataset for point cloud registration. (c) Sample result of QUASAR on the PASSTA image stitching dataset.

sult in large errors even at 80%, see red “+” in the figure). RANSAC, GORE, and QUASAR are robust to 90% outliers, with QUASAR being slight more accurate than the others. In some of the tests with 90% outliers, RANSAC converged to poor solutions (red “+” in the figure).

Fig. 3(a, second row) shows the rotation errors for low noise ($\sigma_i = 0.01, i = 1, \dots, N$) and extreme outlier ratios (91 – 96%). Here we use $N = 100$ to ensure a sufficient number of inliers. Not surprisingly, Wahba, FGR, and RANSAC break at such extreme levels of outliers. GORE fails only once at 96% outliers (red “+” in the figure), while QUASAR returns highly-accurate solutions in all tests.

Fig. 3(a, third row) shows the rotation errors for a higher noise level ($\sigma_i = 0.1, i = 1, \dots, N$) and increasing outlier

ratios (0 – 90%). Even with large noise, QUASAR is still robust against 80% outliers, an outlier level where all the other algorithms fail to produce an accurate solution.

6.2. Point Cloud Registration

In point cloud registration, [53] showed how to build invariant measurements to decouple the estimation of scale, rotation and translation. Therefore, in this section we test QUASAR to solve the rotation-only subproblem. We use the Bunny dataset from the Stanford 3D Scanning Repository [20] and resize the corresponding point cloud to be within the $[0, 1]^3$ cube. The Bunny is first down-sampled to $N = 40$ points, and then we apply a random rotation with additive noise and random outliers, according to eq. (2). Results are averaged over 40 Monte Carlo runs. Fig. 3(b, first row)

Datasets	Noise level	Outlier ratio	Relative relaxation gap		Solution rank		Solution stable rank	
			Mean	SD	Mean	SD	Mean	SD
Synthetic	Low	0 – 0.9	4.32e−9	3.89e−8	1	0	1 + 1.19e−16	1.20e−15
	Low	0.91 – 0.96	1.47e−8	2.10e−7	1	0	1 + 6.76e−9	1.05e−7
	High	0 – 0.8	2.25e−8	3.49e−7	1	0	1 + 9.08e−8	1.41e−6
	High	0.9	0.03300	0.05561	33.33	36.19	1.1411	0.2609
Bunny	Low	0 – 0.9	1.53e−8	1.64e−7	1	0	1 + 4.04e−16	4.83e−15
	High	0 – 0.9	9.96e−12	4.06e−11	1	0	1 + 7.53e−18	3.85e−16

Table 1. Mean and standard deviation (SD) of the relative relaxation gap, the solution rank, and the solution stable rank of QUASAR on the synthetic and the Bunny datasets. The relaxation is *tight* in all tests except in the synthetic tests with high noise and 90% outliers.

shows an example of the point cloud registration problem with vector-to-vector putative correspondences (including outliers) shown in green. Fig. 3(b, second row) and Fig. 3(b, third row) evaluate the compared techniques for increasing outliers in the low noise ($\sigma_i = 0.01$) and the high noise regime ($\sigma_i = 0.1$), respectively. The results confirm our findings from Fig. 3(a, first row) and Fig. 3(a, third row). QUASAR dominates the other techniques. Interestingly, in this dataset QUASAR performs even better (more accurate results) at 90% outliers and high noise.

Tightness of QUASAR. Table 1 provides a detailed evaluation of the quality of QUASAR in both the synthetic and the Bunny datasets. Tightness is evaluated in terms of (i) the *relative relaxation gap*, defined as $\frac{f_{\text{QCQP}} - f_{\text{SDP}}^*}{f_{\text{QCQP}}}$, where f_{SDP}^* is the optimal cost of the relaxation (20) and f_{QCQP} is the cost attained in (15) by the corresponding (possibly rounded) solution (the relaxation is exact when the gap is zero), (ii) the rank of the optimal solution of (20) (the relaxation is exact when the rank is one). Since the evaluation of the rank requires setting a numeric tolerance (10^{-6} in our tests), we also report the *stable rank*, the squared ratio between the Frobenius norm and the spectral norm, which is less sensitive to the choice of the numeric tolerance.

6.3. Image Stitching

We use the PASSTA dataset to test QUASAR in challenging panorama stitching applications [39]. To merge two images together, we first use SURF [5] feature descriptors to match and establish putative feature correspondences. Fig. 3(c, first row) shows the established correspondences between two input images from the *Lunch Room* dataset. Due to significantly different lighting conditions, small overlapping area, and different objects having similar image features, 46 of the established 70 SURF correspondences (66%) are outliers (shown in red). From the SURF feature points, we apply the inverse of the known camera intrinsic matrix \mathbf{K} to obtain unit-norm bearing vectors ($\{\mathbf{a}_i, \mathbf{b}_i\}_{i=1}^{70}$) observed in each camera frame. Then we use QUASAR (with $\sigma^2 \bar{c}^2 = 0.001$) to find the relative rotation \mathbf{R} between the two camera frames. Using the estimated rotation, we compute the homography matrix as $\mathbf{H} = \mathbf{K} \mathbf{R} \mathbf{K}^{-1}$ to stitch the pair of images together. Fig. 3(c) shows an example of the

image stitching results. QUASAR performs accurate stitching in very challenging instances with many outliers and small image overlap. On the other hand, applying the *M-estimator Sample Consensus* (MSAC) [50, 26] algorithm, as implemented by the Matlab “estimateGeometricTransform” function, results in an incorrect stitching (see the Supplementary Material for extra results and statistics).

7. Conclusions

We propose the first polynomial-time certifiably optimal solution to the Wahba problem with outliers. The crux of the approach is the use of a TLS cost function that makes the estimation insensitive to a large number of outliers. The second key ingredient is to write the TLS problem as a QCQP using a quaternion representation for the unknown rotation. Despite the simplicity of the QCQP formulation, the problem remains non-convex and hard to solve globally. Therefore, we develop a convex SDP relaxation. While a naive relaxation of the QCQP is loose in the presence of outliers, we propose an improved relaxation with redundant constraints, named QUASAR. We provide a theoretical proof that QUASAR is tight (computes an exact solution to the QCQP) in the noiseless and outlier-free case. More importantly, we empirically show that the relaxation remains tight in the face of high noise and extreme outliers (95%). Experiments on synthetic and real datasets for point cloud registration and image stitching show that QUASAR outperforms RANSAC, as well as state-of-the-art robust local optimization techniques, global outlier-removal procedures, and BnB optimization methods. While running in polynomial time, the general-purpose SDP solver used in our current implementation scales poorly in the problem size (about 1200 seconds with MOSEK [2] and 500 seconds with SDPNAL+ [55] for 100 correspondences). Current research effort is devoted to developing fast specialized SDP solvers along the line of [43].

Acknowledgements

This work was partially funded by ARL DCIST CRA W911NF-17-2-0181, ONR RAIDER N00014-18-1-2828, and Lincoln Laboratory’s Resilient Perception in Degraded Environments program. The authors want to thank Álvaro Parra Bustos and Tat-Jun Chin for kindly sharing BnB implementations used in [7, 42].

References

- [1] Shakil Ahmed, Eric C. Kerrigan, and Imad M. Jaimoukha. A semidefinite relaxation-based algorithm for robust attitude estimation. *IEEE Transactions on Signal Processing*, 60(8):3942–3952, 2012. 2
- [2] MOSEK ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 8.1.*, 2017. 6, 8
- [3] K. Somani Arun, Thomas S. Huang, and Steven D. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Trans. Pattern Anal. Machine Intell.*, 9(5):698–700, 1987. 1, 2
- [4] Afonso S. Bandeira. A note on probably certifiably correct algorithms. *Comptes Rendus Mathématique*, 354(3):329–333, 2016. 2
- [5] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: speeded up robust features. In *European Conf. on Computer Vision (ECCV)*, 2006. 8
- [6] Jean-Charles Bazin, Yongduek Seo, Richard Hartley, and Marc Pollefeys. Globally optimal inlier set maximization with unknown rotation and focal length. In *European Conf. on Computer Vision (ECCV)*, pages 803–817, 2014. 1, 2
- [7] Jean-Charles Bazin, Yongduek Seo, and Marc Pollefeys. Globally optimal consensus set maximization through rotation search. In *Asian Conference on Computer Vision*, pages 539–551. Springer, 2012. 2, 3, 6, 8
- [8] Paul J. Besl and Neil D. McKay. A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(2), 1992. 1, 2
- [9] Michael J. Black and Anand Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *Intl. J. of Computer Vision*, 19(1):57–91, 1996. 3, 4
- [10] Gérard Blais and Martin D. Levine. Registering multiview range data to create 3d computer objects. *IEEE Trans. Pattern Anal. Machine Intell.*, 17(8):820–824, 1995. 1, 2
- [11] William G. Breckenridge. Quaternions - proposed standard conventions. In *JPL, Tech. Rep. INTEROFFICE MEMORANDUM IOM 343-79-1199*, 1999. 4
- [12] Jesus Briales and Javier Gonzalez-Jimenez. Convex Global 3D Registration with Lagrangian Duality. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 6
- [13] Álvaro Parra Bustos and Tat-Jun Chin. Guaranteed outlier removal for point cloud registration with correspondences. *IEEE Trans. Pattern Anal. Machine Intell.*, 40(12):2868–2882, 2018. 2, 3
- [14] Dylan Campbell, Lars Petersson, Laurent Kneip, and Hongdong Li. Globally-optimal inlier set maximisation for simultaneous camera pose and feature correspondence. In *Intl. Conf. on Computer Vision (ICCV)*, pages 1–10, 2017. 3
- [15] Yang Cheng and John L. Crassidis. A total least-squares estimate for attitude determination. In *AIAA Scitech 2019 Forum*, page 1176, 2019. 1, 2
- [16] Tat-Jun Chin, Yang Heng Kee, Anders Eriksson, and Frank Neumann. Guaranteed outlier removal with mixed integer linear programs. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5858–5866, 2016. 3
- [17] Tat-Jun Chin and David Suter. The maximum consensus problem: recent algorithmic advances. *Synthesis Lectures on Computer Vision*, 7(2):1–194, 2017. 3
- [18] Tat-Jun Chin and David Suter. Star tracking using an event camera. *ArXiv Preprint: 1812.02895*, 12 2018. 1
- [19] Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun. Robust reconstruction of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5556–5565, 2015. 1, 2
- [20] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH*, pages 303–312, 1996. 7
- [21] Olof Enqvist, Erik Ask, Fredrik Kahl, and Kalle Åström. Robust fitting for multiple view geometry. In *European Conf. on Computer Vision (ECCV)*, pages 738–751. Springer, 2012. 3
- [22] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with application to image analysis and automated cartography. *Commun. ACM*, 24:381–395, 1981. 2, 3
- [23] James Richard Forbes and Anton H. J. de Ruiter. Linear-Matrix-Inequality-Based solution to Wahba’s problem. *Journal of Guidance, Control, and Dynamics*, 38(1):147–151, 2015. 2
- [24] John Gower and Garnt B. Dijkstra. Procrustes problems. *Procrustes Problems, Oxford Statistical Science Series*, 30, 01 2005. 2
- [25] Michael Grant, Stephen Boyd, and Yinyu Ye. CVX: Matlab software for disciplined convex programming, 2014. 6
- [26] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition, 2004. 8
- [27] Richard I. Hartley and Fredrik Kahl. Global optimization through rotation space search. *Intl. J. of Computer Vision*, 82(1):64–79, 2009. 1, 2, 3
- [28] Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *J. Opt. Soc. Amer.*, 4(4):629–642, 1987. 1, 2
- [29] Berthold K. P. Horn, Hugh M. Hilden, and Shahriar Negahdaripour. Closed-form solution of absolute orientation using orthonormal matrices. *J. Opt. Soc. Amer.*, 5(7):1127–1135, 1988. 1, 2, 6
- [30] F. Landis Markley John L. Crassidis and Yang Cheng. Survey of nonlinear attitude estimation methods. *Journal of Guidance, Control, and Dynamics*, 30(1):12–28, 2007. 2
- [31] Kourosh Khoshelham. Closed-form solutions for estimating a rigid motion from plane correspondences extracted from point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 114:78 – 91, 2016. 2
- [32] Pierre-Yves Lajoie, Siyi Hu, Giovanni Beltrame, and Luca Carlone. Modeling perceptual aliasing in SLAM via Discrete–Continuous Graphical Models. *IEEE Robotics and Automation Letters*, 4(2):1232–1239, 2019. 3
- [33] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Intl. J. of Computer Vision*, 60(2):91–110, 2004. 2
- [34] Kirk Mac Tavish and Timothy D. Barfoot. At all costs: A comparison of robust cost functions for camera correspon-

- dence outliers. In *Computer and Robot Vision (CRV), 2015 12th Conference on*, pages 62–69. IEEE, 2015. 3
- [35] Ameesh Makadia, Alexander Patterson, and Kostas Daniilidis. Fully automatic registration of 3d point clouds. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1297–1304, 2006. 1
- [36] F. Landis Markley. Attitude determination using vector observations and the singular value decomposition. *The Journal of the Astronautical Sciences*, 36(3):245–258, 1988. 1, 2
- [37] F. Landis Markley and John L. Crassidis. *Fundamentals of spacecraft attitude determination and control*, volume 33. Springer, 2014. 1, 2
- [38] Peter Meer, Doron Mintz, Azriel Rosenfeld, and Dong Yoon Kim. Robust regression methods for computer vision: A review. *Intl. J. of Computer Vision*, 6(1):59–70, 1991. 3
- [39] Giulia Meneghetti, Martin Danelljan, Michael Felsberg, and Klas Nordberg. Image alignment for panorama stitching in sparsely structured environments. In *Scandinavian Conference on Image Analysis*, pages 428–439. Springer, 2015. 8
- [40] Mario Milanese. Estimation and prediction in the presence of unknown but bounded uncertainty: A survey. In *Robustness in Identification and Control*, pages 3–24. Springer US, Boston, MA, 1989. 4
- [41] Carl Olsson, Olof Enqvist, and Fredrik Kahl. A polynomial-time bound for matching and registration with outliers. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008. 3
- [42] Álvaro Parra Bustos and Tat-Jun Chin. Guaranteed outlier removal for rotation search. In *Intl. Conf. on Computer Vision (ICCV)*, pages 2165–2173, 2015. 1, 2, 3, 6, 8
- [43] David M. Rosen, Luca Carlone, Afonso S. Bandeira, and John J. Leonard. SE-Sync: a certifiably correct algorithm for synchronization over the Special Euclidean group. *Intl. J. of Robotics Research*, 2018. 8
- [44] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *Intl. Conf. on Computer Vision (ICCV)*, pages 2564–2571, 2011. 2
- [45] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (FPFH) for 3d registration. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3212–3217, 2009. 2
- [46] James Saunderson, Pablo A. Parrilo, and Alan S. Willsky. Semidefinite descriptions of the convex hull of rotation matrices. *SIAM J. OPTIM.*, 25(3):1314–1343, 2015. 2
- [47] Peter H. Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31:1–10, 1966. 1, 2
- [48] Malcolm D. Shuster. Maximum likelihood estimation of spacecraft attitude. *J. Astronautical Sci.*, 37(1):79–88, 01 1989. 1
- [49] Malcolm D. Shuster. A survey of attitude representations. *Journal of the Astronautical Sciences*, 41(4):439–517, 1993. 4
- [50] Philip H. S. Torr and Andrew Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Comput. Vis. Image Underst.*, 78(1):138–156, 2000. 8
- [51] Roberto Tron, David M. Rosen, and Luca 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”*, 2015. 6
- [52] Grace Wahba. A least squares estimate of satellite attitude. *SIAM review*, 7(3):409–409, 1965. 1, 2
- [53] Heng Yang and Luca Carlone. A polynomial-time solution for robust registration with extreme outlier rates. In *Robotics: Science and Systems (RSS)*, 2019. 7
- [54] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-ICP: A globally optimal solution to 3D ICP point-set registration. *IEEE Trans. Pattern Anal. Machine Intell.*, 38(11):2241–2254, Nov. 2016. 1, 2
- [55] Liuqin Yang, Defeng Sun, and Kim-Chuan Toh. SDPNAL+: a majorized semismooth Newton-CG augmented lagrangian method for semidefinite programming with nonnegative constraints. *Math. Program. Comput.*, 7(3):331–366, 2015. 8
- [56] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In *European Conf. on Computer Vision (ECCV)*, pages 766–782. Springer, 2016. 3, 6