Certifiably Optimal Anisotropic Rotation Averaging

Carl Olsson[1] Yaroslava Lochman[2] Johan Malmport[1] Christopher Zach[2]
[1]Lund University [2]Chalmers University of Technology

Abstract

Rotation averaging is a key subproblem in applications of computer vision and robotics. Many methods for solving this problem exist, and there are also several theoretical results analyzing difficulty and optimality. However, one aspect that most of these have in common is a focus on the isotropic setting, where the intrinsic uncertainties in the measurements are not fully incorporated into the resulting optimization task. Recent empirical results suggest that moving to an anisotropic framework, where these uncertainties are explicitly included, can result in an improvement of solution quality. However, global optimization for rotation averaging has remained a challenge in this scenario. In this work we show how anisotropic costs can be incorporated in certifiably optimal rotation averaging. We also demonstrate how existing solvers, designed for isotropic situations, fail in the anisotropic setting. Finally, we propose a stronger relaxation and empirically show that it recovers global optima in all tested datasets and leads to more accurate reconstructions in almost all scenes.[1]

1. Introduction

Rotation averaging has been a topic of interest now for some time, dating back more than two decades, with early work such as [20]. Govindu's groundbreaking work has since been followed up by many others and the study of rotation averaging is now a very important area.

In large part this is because of the great importance it has for many problems within computer vision and robotics. Traditionally, it has been a core component of non-sequential structure from motion (SfM), see e.g. [29, 30, 35]. The key strength here is that rotation averaging optimizes the orientations of all the cameras at the same time, allowing recovery methods to avoid incrementally increasing errors (drift) [14] that easily occur in sequential SfM methods. More recently, [7, 9, 10] point out the importance of rotation averaging to improve the accuracy and speed of simultaneous localization and mapping (SLAM).

The main difficulty in solving the rotation averaging problem is that rotations reside in a nonlinear manifold $SO(3)$. Some approaches [11, 15, 21, 22, 24, 43] address this challenge by using local optimization over $SO(3)^n$ via explicit parametrizations. These, and other methods of local optimization all work by moving along descent directions, which may lead to a global optima in practical and sufficiently "nice" situations. Wilson et al. [47, 48] study when this occurs, and when local optimization is "hard" due to the existence of "bad" local minima. They conclude that this depends on both the measurement noise and on the algebraic connectivity of the graph representing the problem.

More recent methods for solving rotation averaging problems formulate a non-convex quadratically constrained quadratic problem (QCQP) and relax it to a semidefinite program (SDP). The quadratic constraints force the desired matrices to be orthogonal. The non-quadratic determinant constraint is often dropped, effectively replacing optimization over $SO(3)$ with one over $O(3)$ [5, 12, 16, 28, 36, 39]. The SDP formulation *is* a relaxation of the QCQP and results in a strictly less or equal optimal value than the QCQP. There are, however, several theoretical results [18, 22, 39] that point to this "gap" often being non-existent, giving a tight lower bound allowing the recovery of a "certifiably correct" solution. In [18] an explicit bound depending on the measurement noise and algebraic connectivity of the graph is given, in analogy to the local optimization setting.

While the above methods and algorithms are often certifiably correct, they minimize the isotropic chordal distance and, therefore, ignore the relative rotation uncertainties. The anisotropic version of the chordal distance allows to explicitly include uncertainties obtained by local two-view optimization into the rotation averaging framework. Recent empirical results such as [51] suggest that accounting for uncertainties in the optimization could "largely improve reconstruction quality". Therefore the isotropic assumption discussed above can be seen as a disadvantage.

In this paper, we develop a certifiably optimal SDP-formulation able to optimize the anisotropic chordal distance. In light of the above, this is an improvement of many

[1]This work was supported by the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP), funded by the Knut and Alice Wallenberg Foundation, and the Swedish Research Council (#2023-05341).

earlier results. However, global optimization will become more challenging in this case, as—based on our analysis—a direct modification of the objective function results in an SDP that rarely provides a tight lower bound. The main reason for this is, that optimization of an anisotropic cost over $O(3)$ typically results in a solution outside $SO(3)$, implying that the determinant constraint on rotation matrices cannot be dropped in order to expect tight semidefinite relaxations. As mitigation we present a new relaxation that is capable of further constraining the solution to the convex hull of $SO(3)$, $\mathrm{conv}(SO(3))$.

Further, we empirically verify—on a number of synthetic and real datasets—that this new relaxation is sufficient to recover the globally optimal solution, and that our anisotropic model generally achieves solutions that are more accurate than the standard isotropic chordal penalty. In summary, our main contributions are:

- We show how anisotropic costs can be incorporated in certifiably correct rotation averaging.
- We provide an analysis of the objective function that explains why regular solvers only enforcing $O(3)$ membership usually fail in the anisotropic setting.
- We present a stronger convex relaxation able to enforce $\mathrm{conv}(SO(3))$ membership and verify empirically, that the proposed formulation is able to recover a global optimum in all tested instances. To our knowledge, this is the first formulation to yield tight relaxations.

2. Related Work

There are many ways of formulating the rotation averaging problem and representing rotations; [6, 33] represent the rotations using vectors and optimize over them, while [19] uses quaternions. Quaternions have some advantages, but the fact that unit quaternions form a double cover of $SO(3)$ means that quaternion distances are less straightforward and a second optimization step is needed to determine their appropriate signs. In this paper we will only consider matrix representations.

The works [3, 23] also consider certifiable optimization of the anisotropic chordal distance. However, their SDP relaxation is based on the Cayley mapping and the relation to our formulation is somewhat unclear. In addition, [3, 23] have to introduce redundant constraints in order to get a "reasonably tight" SDP relaxation. It is likely, in the light of our analysis in this paper, and the discussion in [6] that this is a result of the authors only enforcing $O(3)$ membership for their optimization.

It is also worth mentioning that—while this paper looks at the anisotropic L_2 chordal distance—it is possible to consider other distances between rotations as well as their robust counterparts. E.g., [47, 48] consider the geodesic distance. Further, in [11, 43, 46] the authors analyze and look at different (robustified) distances in order to obtain a "robust" formulation (resulting in the solution being less susceptible to outliers). Outliers are handled via the introduction of auxiliary variables in the SDP formulation in [8, 38].

From an algorithmic point of view, the SDP relaxation can be solved by general-purpose solvers, e.g. [1]. There are also several dedicated solvers, constructed specifically for the rotation averaging problem. [18, 36] use coordinate descent (fixing all but one row/column and minimizing), [28] use a primal-dual update rule, and in [12, 49] the authors combine a global coordinate descent approach with a local step. An early approach [44] looks at eigenvectors of a specific matrix. A particularly powerful method used in [5, 16] is that of the Riemannian staircase, where one optimizes using block matrices of a certain dimension and increases ("climb the staircase") the dimension until a global optimum is found. The authors show that such a method can be made highly effective.

Like most work on rotation averaging, we focus on point estimates and assume unimodal Gaussian noise in the observations. If the input data has multiple modes (e.g. due to perceptual aliasing), the particle-based algorithm proposed in [4] provides more expressive marginals. Perceptual aliasing can also be explicitly addressed by leveraging cycle consistency constraints (e.g. [25, 42, 50]).

3. Certifiably Optimal Rotation Averaging

In the context of structure from motion and SLAM, the goal of rotation averaging is to estimate absolute camera orientations from relative rotation measurements between pairs of calibrated cameras [19, 20, 22, 44]. If R_i and R_j are rotation matrices, encoding the orientation of cameras $P_i = \begin{bmatrix} R_i & t_i \end{bmatrix}$ and $P_j = \begin{bmatrix} R_j & t_j \end{bmatrix}$, the relative rotation R_{ij} between the two cameras, that is, P_i's orientation in the camera coordinate system of P_j is $R_{ij} = R_i R_j^\top$. Given estimates $\tilde{R}_{ij}$ of R_{ij}, typically obtained by solving two view relative pose [31], the traditional rotation averaging formulation [16, 18, 36] accounts for estimation error using the chordal distance $\|\tilde{R}_{ij} - R_i R_j^\top\|_F$ and aims to solve

$$\min_{R_i \in SO(3)} \sum_{i \neq j} \|\tilde{R}_{ij} - R_i R_j^\top\|_F^2. \tag{1}$$

The derivation of a strong convex relaxation relies on two observations: Firstly, since the set of rotation matrices have constant (Frobenius) norm, the objective function can be replaced by $-\sum_{i\neq j}\langle \tilde{R}_{ij}, R_j R_i^\top \rangle$. If we let $\mathbf{R} = \begin{bmatrix} R_1^\top & R_2^\top & \dots \end{bmatrix}^\top$ and $\tilde{\mathbf{R}}$ be the matrix containing the blocks $\tilde{R}_{ij}$ (and zeros where no relative rotation estimate is available) we can simplify the objective to $-\langle \tilde{\mathbf{R}}, \mathbf{R}\mathbf{R}^\top \rangle$, which is linear in the relative rotations. Secondly, by orthogonality we have $R_i R_i^\top = \mathtt{I}$, resulting in the Quadratically Constrained Quadratic Program (QCQP)

$$p^* := \min_{\mathbf{R}} -\langle \tilde{\mathbf{R}}, \mathbf{R}\mathbf{R}^\top \rangle \quad \text{s.t. } \forall i : R_i R_i^\top = \mathtt{I}. \qquad (2)$$

QCQP are typically non-convex but can be relaxed to convex linear semidefinite programs (SDP) by replacing $\mathbf{R}\mathbf{R}^\top$ with a positive semidefinite matrix $\mathbf{X}$ (technically by taking the Lagrange dual twice), yielding the relaxation

$$d_2^* := \min_{\mathbf{X} \succeq 0} -\langle \tilde{\mathbf{R}}, \mathbf{X} \rangle \quad \text{s.t. } \forall i : X_{ii} = \mathtt{I}. \qquad \text{(SDP-O(3)-ISO)}$$

The above program is convex and can therefore be solved reliably using general purpose [1, 26] or specialized solvers [5, 16, 18]. The difference $p^* - d_2^*$ between objective values is referred to as the duality gap. It is clear that we can construct $\mathbf{X} = \mathbf{R}\mathbf{R}^\top$ from a solution $\mathbf{R}$ of (2), which is feasible in (SDP-O(3)-ISO) and therefore $p^* \geq d_2^*$. Moreover, if the solution $\mathbf{X}$ to (SDP-O(3)-ISO) has $\text{rank}(\mathbf{X}) = 3$ then it can be factorized into $\mathbf{X} = \mathbf{R}\mathbf{R}^\top$, with $\mathbf{R}$ being feasible in (2), which yields $p^* = d_2^*$ and certifies that $\mathbf{R}$ is globally optimal in (2). Thus, a solution of the original non-convex QCQP is obtained via a convex SDP program.

We remark that the use of a linear objective is important to obtain a strong relaxation. Linear objectives admit solutions at extreme points (on the boundary) of the feasible set which are often of low rank. It has been shown, both theoretically and empirically, [10, 18, 38, 39, 44] that the duality gap of synchronization problems is often zero under various bounded noise regimes. In [18], explicit bounds depending on the algebraic connectivity of the camera graph are given. For example, for fully connected graphs the duality gap can be shown to be zero if there is a solution with an angular error of at most $42.9°$—which is a rather generous measurement error.

4. Anisotropic Rotation Averaging

While the use of the chordal distance often results in tight SDP relaxations, a downside is that it does not take into account the uncertainty of the relative rotation estimates $\tilde{R}_{ij}$. For two-view relative pose problems, it is frequently the case that the reprojection errors are much less affected by rotations in certain directions than others. In Figure 1 we illustrate this on a real two-view problem. The left image shows the ground truth solution (obtained after bundle adjustment). To generate the graphs to the right, we sampled rotations of the second camera around its x-,y-, and z-axes between -5 and 5 degrees. For each sampled rotation, we found the best 3D point locations and position of the second camera, and computed the sum-of-squared (calibrated) reprojection errors (solid curves). The result shows that rotations around the black/upward axis affects the reprojection errors significantly less than the others, resulting in a direction of larger uncertainty. By performing sensitivity analysis (i.e. by looking at the singular values of the Jacobian) one can arrive at the similar conclusion. It is therefore desirable to propagate this information to the rotation averaging stage to increase the accuracy of the model. The dashed curves (that are almost indistinguishable from the solid ones) show the approximations, presented below, that can be used for this purpose. In the supplementary material, we also show how isotropic rotation averaging gets negatively affected by the single noisy relative rotation (while the proposed method presented further remains unaffected).

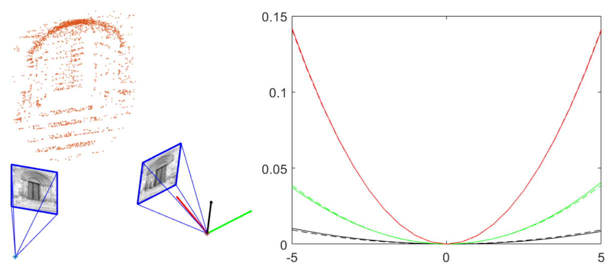

Figure 1. Ground truth solution to two view problem (left). Sum-of-squared reprojection errors (right, solid curves) obtained when rotating the second camera around the black, green and red axes (from -5 to 5 degrees) and anisotropic quadratic approximations (dashed curves) that can be used in rotation averaging.

4.1. Incorporating uncertainties

Consider a relative rotation $\tilde{R}_{ij}$ between cameras i and j, that has been fully optimized using e.g. a Gauss-Newton method for two-view optimization. Using the exponential map we write $R_{ij} = e^{[\Delta\boldsymbol{\omega}_{ij}]_\times}\tilde{R}_{ij}$[2], where $\Delta\boldsymbol{\omega}_{ij}$ is an axis-angle representation of the deviation from $\tilde{R}_{ij}$, i.e. the axis-angle of $R_{ij}\tilde{R}_{ij}^\top$. Since $\tilde{R}_{ij}$ is a minimum, gradient terms vanish and the two-view objective can locally be approximated by a quadratic function (up to an irrelevant constant)

$$\tfrac{1}{2}\Delta\boldsymbol{\omega}_{ij}^\top H_{ij} \Delta\boldsymbol{\omega}_{ij}. \qquad (3)$$

We assume that H_{ij} is positive semidefinite—usually given by a Gauss-Newton approximation $H_{ij} = J_{ij}^\top J_{ij}$, where J_{ij} is the Jacobian of the underlying residual function. We remark that the two-view objective may depend on additional parameters such as camera positions and 3D point locations. Here we assume that these "nuisance" parameters have been marginalized out (using the Schur complement), and therefore we work with a reduced objective solely in terms of the rotation parameters $\Delta\boldsymbol{\omega}_{ij}$.

We assume Laplace's approximation to be valid at least locally in order to quantify the uncertainty of $\Delta\boldsymbol{\omega}_{ij}$, i.e. identify the precision matrix with the Hessian H_{ij}. While in

[2]Other parameterizations are discussed in the supplementary material.

general the uncertainty of a 3×3 matrix R_{ij} is represented by a 9×9 covariance/precision matrix, we aim for a linear cost (in terms of rotation matrices) in order to achieve a strong relaxation, which means that the objective is restricted to an inner product between R_{ij} and a constant but input-dependent matrix. We first determine a matrix M_{ij} such that

$$\begin{aligned} \Delta\boldsymbol{\omega}_{ij}^\top H_{ij} \Delta\boldsymbol{\omega}_{ij} &= \operatorname{tr}\left([\Delta\boldsymbol{\omega}_{ij}]_\times^\top M_{ij} [\Delta\boldsymbol{\omega}_{ij}]_\times\right) \\ &= -\operatorname{tr}\left(M_{ij}[\Delta\boldsymbol{\omega}_{ij}]_\times^2\right). \end{aligned} \tag{4}$$

Using the identity $-[\mathbf{v}]_\times^2 = \mathbf{v}^\top\mathbf{v}\mathtt{I} - \mathbf{v}\mathbf{v}^\top$, we obtain that $H_{ij} = \operatorname{tr}(M_{ij})\mathtt{I} - M_{ij}$. Taking the trace on both sides further shows that $\operatorname{tr}(H_{ij}) = 2\operatorname{tr}(M_{ij})$, resulting in

$$M_{ij} = \tfrac{\operatorname{tr}(H_{ij})}{2}\mathtt{I} - H_{ij}. \tag{5}$$

Next, we deduce via the Taylor expansion of the exponential map that

$$[\Delta\boldsymbol{\omega}_{ij}]_\times \approx R_{ij}\tilde{R}_{ij}^\top - \mathtt{I}, \tag{6}$$

with equality up to first order terms. This yields the objective corresponding to (3) based on (4),

$$\tfrac{1}{2}\operatorname{tr}\left((R_{ij}\tilde{R}_{ij}^\top - \mathtt{I})^\top M_{ij}(R_{ij}\tilde{R}_{ij}^\top - \mathtt{I})\right). \tag{7}$$

Similarly to the isotropic case, by leveraging the properties of rotation matrices and by omitting constants, this objective can be reduced to a linear one[3],

$$-\operatorname{tr}\left(M_{ij}\tilde{R}_{ij}R_{ij}^\top\right) = -\langle M_{ij}\tilde{R}_{ij}, R_{ij}\rangle. \tag{8}$$

We observe that this objective can be viewed as a negative log-likelihood of a Langevin distribution $\mathcal{L}(R_{ij}, M_{ij})$, e.g. [13], with a density function of the type

$$p(\tilde{R}_{ij}) \propto e^{\langle M_{ij}\tilde{R}_{ij}, R_{ij}\rangle}, \tag{9}$$

and that generating rotations $\tilde{R}_{ij}$ according to (9) is approximately (up to second order terms) equivalent to generating axis-angle vectors $\Delta\boldsymbol{\omega}_{ij}$ from a normal distribution $\mathcal{N}(0, H_{ij}^{-1})$ with density

$$p(\Delta\boldsymbol{\omega}_{ij}) \propto e^{-\frac{1}{2}\Delta\boldsymbol{\omega}_{ij}^\top H_{ij}\Delta\boldsymbol{\omega}_{ij}}. \tag{10}$$

Hence, the above derivations suggest to generalize the rotation averaging approach by replacing the relative estimates $\tilde{R}_{ij}$ with a weighted version $M_{ij}\tilde{R}_{ij}$, where M_{ij} is based on the Hessian H_{ij} of the relative pose problem as outlined above. Therefore, our proposed objective to minimize is

$$-\sum\nolimits_{i \neq j}\langle M_{ij}\tilde{R}_{ij}, R_j R_i^\top\rangle = -\langle \mathbf{N}, \mathbf{R}\mathbf{R}^\top\rangle, \tag{11}$$

where $\mathbf{N}$ is a symmetric block matrix containing the blocks $(M_{ij}\tilde{R}_{ij})^\top$ for $i < j$, 0 when $i = j$, and $M_{ji}\tilde{R}_{ji}$ for $i > j$.

We conclude this section by remarking that—in contrast to H_{ij}—the matrix M_{ij} will generally not be positive semidefinite as the following lemma shows.

Lemma 1. *If $H_{ij} \succeq 0$ then $M_{ij} = \frac{\operatorname{tr}(H_{ij})}{2}\mathtt{I} - H_{ij}$ has eigenvalues $\lambda_1 \geq \lambda_2 \geq |\lambda_3|$.*

Proof. If $\eta_1 \geq \eta_2 \geq \eta_3 \geq 0$ are eigenvalues of H_{ij} then $\frac{\operatorname{tr}(H_{ij})}{2} - \eta_i$, $i = 1, 2, 3$ are the eigenvalues of M_{ij}. Sorted in decreasing order we get $\lambda_1 = \frac{1}{2}(\eta_1 + \eta_2 - \eta_3)$, $\lambda_2 = \frac{1}{2}(\eta_1 - \eta_2 + \eta_3)$ and $\lambda_3 = \frac{1}{2}(-\eta_1 + \eta_2 + \eta_3)$. Only λ_3 can be negative and clearly $\lambda_2 \pm \lambda_3 \geq 0$. □

The case $\lambda_3 < 0$ occurs when the leading eigenvalue of H_{ij} is significantly larger than the other two, i.e. the estimate $\boldsymbol{\omega}_{ij}$ is more certain in a specific direction. This is actually a rather common scenario in image-based estimation as the in-plane part of the rotation is usually better constrained than the out-of-plane rotation angles. Table 1 shows the percentage of indefinite matrices M_{ij} for the real datasets tested in Section 5.2.

4.2. Global solutions: $O(3)$ vs. $SO(3)$

Incorporating the anisotropic objective (11) in the standard relaxation (SDP-O(3)-ISO) may seem like a straightforward extension. We refer to this approach as SDP-O(3)-ANISO. The source of the problem with this approach is that it ignores the determinant constraint and optimizes over $O(3)$. To illustrate this problem, we ran 1000 synthetic anisotropic problem instances (the data generation protocol follows Section 5.1 with the eigenvalues of inverse Hessians sampled from $[0.1, 1]$). Fig. 2 shows that the standard rotation averaging relaxation is never able to recover a rank-3 solution[4] for any of the problem instances. In contrast, our relaxation SDP-cSO(3), presented further in Section 4.3, returns rank-3 solutions in all cases.

To better understand the problem, we consider terms of type $-\langle M\tilde{R}, R\rangle$[5] that constitute the objective function. As we have seen in Section 4.1, these are locally accurate and allow uncertainties of $\tilde{R}$ to be propagated to the rotation averaging stage. However, for successful global optimization one has to guarantee that there is no other matrix in the feasible set that gives a lower cost than $\tilde{R}$. The following theorem shows, that when viewed as a function over $O(3)$, $\tilde{R}$ does typically not result in the smallest objective value:

Theorem 1. *Let the eigenvalues of M be such that $\lambda_1 \geq \lambda_2 \geq |\lambda_3|$ and $\lambda_3 < 0$. Then the minimizer of*

$$f(R) = -\langle M\tilde{R}, R\rangle + \langle M\tilde{R}, \tilde{R}\rangle, \tag{12}$$

[3] In the supplementary material we show that we arrive at the same objective using classical first-order uncertainty propagation.

[4] To determine rank, we used the smallest number N of singular values (i.e., the first N largest values) that sum up to $> 99.9\%$ of the total sum.

[5] For the remainder of this section we focus on a single relative rotation and therefore omit the subscripts.

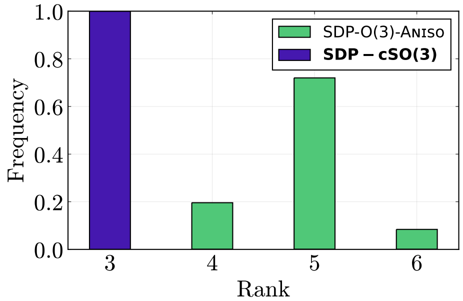

Figure 2. The results of running 1000 synthetic instances of rotation averaging with an anisotropic objective. In contrast to the proposed approach the standard convex relaxations give no solutions of rank 3.

over $SO(3)$ *is given by* $R = \tilde{R}$ *(with* $f(\tilde{R}) = 0$*), but the minimizer* R' *over* $O(3)$ *satisfies* $f(R') = -2|\lambda_3| < 0$.

Proof. We minimize $-\langle M\tilde{R}, R\rangle$ by computing KKT-points over $O(3)$. Introducing a (symmetric) Lagrange multiplier Λ for the constraint $RR^\top - \mathtt{I} = 0$ and differentiating

$$L(R, \Lambda) = -\langle M\tilde{R}, R\rangle + \tfrac{1}{2}\langle \Lambda, RR^\top - \mathtt{I}\rangle \tag{13}$$

with respect to R, gives $\frac{\partial L}{\partial R} = -M\tilde{R} + \Lambda R \stackrel{!}{=} 0$. Hence the task is to find a symmetric multiplier matrix Λ and a rotation matrix $Q = \tilde{R}R^\top$ such that $MQ = \Lambda$. Since M is symmetric, it has an eigen-decomposition $M = UDU^\top$, and we must have $(MQ)(MQ)^\top = MM^\top = UD^2U^\top = \Lambda\Lambda^\top$. Therefore, Λ has the same eigenvalues as M up to their signs, and consequently all solutions for Λ are given by $\Lambda = UDSU^\top$, where $S = \mathrm{diag}(\pm1, \pm1, \pm1)$. Further

$$R = \Lambda^{-1}M\tilde{R} = UD^{-1}S^{-1}U^\top UDU^\top\tilde{R} = USU^\top\tilde{R} \tag{14}$$

with the corresponding objective value

$$-\langle UDU^\top\tilde{R}, USU^\top\tilde{R}\rangle = \pm\lambda_1 \pm \lambda_2 \pm \lambda_3, \tag{15}$$

where $\lambda_1, \lambda_2, \lambda_3$ are the eigenvalues of M. According to Lemma 1 the smallest eigenvalue λ_3 of M is negative if M is indefinite. In this case the objective function is minimized over $O(3)$ by choosing $S = \mathrm{diag}(1, 1, -1)$ which yields the solution $R = USU^\top\tilde{R}$ and $-\langle M\tilde{R}, R\rangle = -\lambda_1 - \lambda_2 - |\lambda_3|$. In contrast $-\langle M\tilde{R}, \tilde{R}\rangle = -\lambda_1 - \lambda_2 + |\lambda_3|$, showing that the smallest value of f is $f(R) = -2|\lambda_3|$.

Since $SO(3) \subset O(3)$ the KKT points of $SO(3)$ are those which satisfy $\det(R) = 1$, which implies that the matrix S has to have either zero or two elements that are -1. Therefore the above choice of S is infeasible, and the optimal S is instead given by $S = I$, which shows that the minimum value of f is $f(\tilde{R}) = 0$. □

The above result shows that loss-terms of the type $-\langle M\tilde{R}, R\rangle$ will favor incorrect solutions (that can be far from the estimation $\tilde{R}$) when optimizing over $O(3)$, which explains the results in Figure 2. While strictly enforcing $R \in SO(3)$ is undesirable, as it results in non-convex constraints, the following result shows that it is enough to use $R \in \mathrm{conv}(SO(3))$ to resolve this issue.

Corollary 1. *If the eigenvalues of* M *fulfill* $\lambda_1 \geq \lambda_2 \geq |\lambda_3|$ *then the function* (12) *is non-negative on* $\mathrm{conv}(SO(3))$.

Proof. For a linear objective function it is well known that the optimal value is attained in an extreme point. For any set C, the extreme points of $\mathrm{conv}(C)$ are always contained in C [37]. Further, since for every point $Q \in SO(3)$ the linear function $-\langle Q, R\rangle$ has its unique minimum in $R = Q$, it can be deduced that the set of extreme points of $\mathrm{conv}(SO(3))$ is exactly $SO(3)$. Thus, the minimum value of f is the same over $SO(3)$ as over $\mathrm{conv}(SO(3))$. □

4.3. A stronger convex relaxation

The results of the previous section show that we need to incorporate constraints allowing us to find solutions in $SO(3)$. The direct incorporation of determinant constraints does not lead to a QCQP, and it is unclear how to compute duals. However, similarly to the previous section we can "solve" single term problems over $SO(3)$ (or at least in the respective convex hull). We therefore introduce auxiliary variables $Q_{ij} = R_iR_j^\top$ and consider the problem

$$\min_{\mathbf{R},\mathbf{Q}} \quad -\langle \mathbf{N}, \mathbf{R}\mathbf{R}^\top\rangle \tag{16}$$

$$\text{s.t} \quad R_iR_i^\top = \mathtt{I}, \tag{17}$$

$$Q_{ij} = R_iR_j^\top, \tag{18}$$

$$Q_{ij} \in SO(3). \tag{19}$$

Next we introduce dual variables Υ_{ii} and Υ_{ij} for the constraints (17) and (18) respectively, but we retain (19) as an explicit constraint. The matrices Υ_{ii} are symmetric, $\Upsilon_{ij} = \Upsilon_{ji}^\top$, and we view them as blocks of a symmetric matrix $\boldsymbol{\Upsilon}$. This yields the Lagrangian

$$\begin{aligned} L(\mathbf{R}, \mathbf{Q}, \boldsymbol{\Upsilon}) &= -\langle \mathbf{N}, \mathbf{R}\mathbf{R}^\top\rangle + \sum\nolimits_i \langle \Upsilon_{ii}, R_iR_i^\top - \mathtt{I}\rangle \\ &\quad + \sum\nolimits_{i\neq j} \langle \Upsilon_{ij}, R_iR_j^\top - Q_{ij}\rangle \\ &= \langle \boldsymbol{\Upsilon} - \mathbf{N}, \mathbf{R}\mathbf{R}^\top\rangle - \mathrm{tr}\,(\boldsymbol{\Upsilon}) - \sum\nolimits_{i\neq j} \langle \Upsilon_{ij}, Q_{ij}\rangle. \end{aligned} \tag{20}$$

The dual variables decouple the constraints and makes the Lagrangian quadratic with respect to $\mathbf{R}$ and linear (and separable) with respect to Q_{ij}. Therefore the dual function

$$d(\boldsymbol{\Upsilon}) = \min_{\mathbf{Q}\in SO(3)^{n\times n}, \mathbf{R}} L(\mathbf{R}, \mathbf{Q}, \boldsymbol{\Upsilon}), \tag{21}$$

can be computed while explicitly forcing $Q_{ij} \in SO(3)$. The minimum value of $\langle \boldsymbol{\Upsilon} - \mathbf{N}, \mathbf{R}\mathbf{R}^\top\rangle$ is 0 if $\boldsymbol{\Upsilon} - \mathbf{N} \succeq 0$

and $-\infty$ otherwise. Let $\mathcal{I}_{SO(3)}$ be the indicator function,

$$\mathcal{I}_{SO(3)}(Q_{ij}) = \begin{cases} 0 & Q_{ij} \in SO(3) \\ \infty & Q_{ij} \notin SO(3) \end{cases}, \tag{22}$$

then we are able to write

$$\begin{aligned} &\min_{Q_{ij} \in SO(3)} -\langle \Upsilon_{ij}, Q_{ij} \rangle \\ &\quad = -\big(\max_{Q_{ij}} \langle \Upsilon_{ij}, Q_{ij} \rangle - \mathcal{I}_{SO(3)}(Q_{ij})\big). \end{aligned} \tag{23}$$

The last term can be identified as $-\mathcal{I}^*_{SO(3)}(\Upsilon_{ij})$ where $\mathcal{I}^*_{SO(3)}$ is the convex conjugate [37] of $\mathcal{I}_{SO(3)}$. We thus obtain the dual problem

$$\begin{aligned} &\max_{\boldsymbol{\Upsilon}} -\mathrm{tr}(\boldsymbol{\Upsilon}) - \textstyle\sum_{i \neq j} \mathcal{I}^*_{SO(3)}(\Upsilon_{ij}) \\ &\text{s.t. } \boldsymbol{\Upsilon} - \mathbf{N} \succeq 0. \end{aligned} \tag{24}$$

Taking the dual once more we get

$$\min_{\mathbf{X} \succeq 0} \max_{\boldsymbol{\Upsilon}} -\mathrm{tr}(\boldsymbol{\Upsilon}) + \langle \mathbf{X}, \boldsymbol{\Upsilon} - \mathbf{N} \rangle - \sum_{i \neq j} \mathcal{I}^*_{SO(3)}(\Upsilon_{ij}), \tag{25}$$

which is the same as

$$\begin{aligned} &\min_{\mathbf{X} \succeq 0} \max_{\boldsymbol{\Upsilon}} -\mathrm{tr}(\mathbf{N}\mathbf{X}) + \textstyle\sum_i \max_{\Upsilon_{ii}} \langle (X_{ii} - \mathtt{I}), \Upsilon_{ii} \rangle \\ &\quad + \textstyle\sum_{i \neq j} \max_{\Upsilon_{ij}} \Big(\langle X_{ij}, \Upsilon_{ij} \rangle - \mathcal{I}^*_{SO(3)}(\Upsilon_{ij}) \Big). \end{aligned} \tag{26}$$

If $X_{ii} \neq \mathtt{I}$ it is clear that the maximum over Υ_{ii} will be unbounded. The second term is $\mathcal{I}^{**}_{SO(3)}(X_{ij})$, which is the convex envelope of the indicator function $\mathcal{I}_{SO(3)}$. This is also the indicator function of $\mathrm{conv}(SO(3))$ [37]. Therefore the bidual program—and consequently our proposed relaxation—is given by

$$\begin{aligned} &\min_{\mathbf{X} \succeq 0} -\mathrm{tr}(\mathbf{N}\mathbf{X}) \\ &\text{s.t. } X_{ii} = \mathtt{I},\ X_{ij} \in \mathrm{conv}(SO(3)). \end{aligned} \tag{SDP-cSO(3)}$$

The constraint $X_{ij} \in \mathrm{conv}(SO(3))$ has been shown to be equivalent to a semidefinite constraint [40, 41]. A 3×3 matrix Y is in $\mathrm{conv}(SO(3))$ if and only if $\mathcal{A}(Y) + \mathtt{I} \succeq 0$, where $\mathcal{A}(Y) =$

$$\begin{pmatrix} -Y_{11} - Y_{22} + Y_{33} & Y_{13} + Y_{31} & Y_{12} - Y_{21} & Y_{23} + Y_{32} \\ Y_{13} + Y_{31} & Y_{11} - Y_{22} - Y_{33} & Y_{23} - Y_{32} & Y_{12} + Y_{21} \\ Y_{12} - Y_{21} & Y_{23} - Y_{32} & Y_{11} + Y_{22} + Y_{33} & Y_{31} - Y_{13} \\ Y_{23} + Y_{32} & Y_{12} + Y_{21} & Y_{31} - Y_{13} & -Y_{11} + Y_{22} - Y_{33} \end{pmatrix}.$$

In contrast to the relaxation in (SDP-O(3)-ISO), the inclusion of the convex hull constraints can rule out solutions with incorrect determinants, as we have shown previously. This is crucial when having indefinite cost matrices. While a theoretical guarantee of a tight relaxation is beyond the scope of this paper, it is clear by construction that our new relaxation will be stronger than (SDP-O(3)-ISO). In addition, our empirical results show that there is a significant difference between the two approaches. Figure 2 shows the result of applying our new convex relaxation to the synthetic problem described in Section 4.2. In all problem instances the new relaxation returns a solution that is of rank 3 and therefore optimal in the primal problem.

5. Experiments

We implemented the SDP program using the conic splitting solver [34] with the JuMP [27] wrapper in Julia[6]. The cost matrix in (SDP-cSO(3)) is down-scaled by the average (across the observed relative poses in the scene) of the largest eigenvalue of the computed Hessian matrix, which proved beneficial empirically. The absolute and relative feasibilities are set to 10^{-5} and 10^{-6}, respectively, and infeasibility tolerance is 10^{-8}. The number of iterations is limited to 500 000. In practice, the methods converge in far fewer iterations to meet the stopping criterion. While it is likely that dedicated solvers are much more efficient, we remark that generalization of efficient methods such as [16] may not be straightforward since they rely heavily on properties $O(d)$, which are not sufficient to ensure $SO(3)$ membership.

We compare[7] three approaches: 1) SDP-O(3)-ISO is the standard rotation averaging approach that uses the regular chordal distance and ignores the $\mathrm{conv}(SO(3))$ constraints, 2) SDP-O(3)-ANISO uses the anisotropic objective function but ignores the $\mathrm{conv}(SO(3))$ constraints, and 3) SDP-cSO(3) is the proposed method that uses both the anisotropic objective and the $\mathrm{conv}(SO(3))$ constraints. For completeness of the synthetic study, we also include SDP-cSO(3)-ISO, which uses the isotropic objective and the $\mathrm{conv}(SO(3))$ constraints. The reported rotation error wrt. ground truth is the difference $\sqrt{\sum_i \|R_i - R_i^*\|_F^2}$, where R_i are the ground truth absolute rotations and R_i^* are the estimated ones. Note that to remove gauge freedom, we align the two sets of rotations $\{R_i\}_i$ and $\{R_i^*\}_i$, by applying a global rotation V to all rotations in $\{R_i\}_i$ such that the rotation error achieves its minimum wrt. V.

5.1. Synthetic experiments

To create synthetic data, we randomly generate absolute rotations R_i and (invertible) covariance matrices H_{ij}^{-1} with eigenvalues in the range $[0.01, 0.1]$. For each instance of the problem, we then draw $\Delta\boldsymbol{\omega}_{ij}$ from $\mathcal{N}(0, H_{ij}^{-1})$ and compute noisy relative rotations $\tilde{R}_{ij} = e^{[\Delta\boldsymbol{\omega}_{ij}]_\times} R_i R_j^\top$.

In Figure 3 we evaluate the effect of using the proposed anisotropic error measurements versus regular chordal distances together with the proposed relaxation versus the standard one. We vary the number of cameras from 2 to 100, and the proportion of missing data from 0% to

[6] The code is available at: https://github.com/ylochman/anisotropic-ra

[7] Complementing results on synthetic and real data can be found in the supplementary material. These also include comparison to the anisotropic extension of the spectral approach [2].

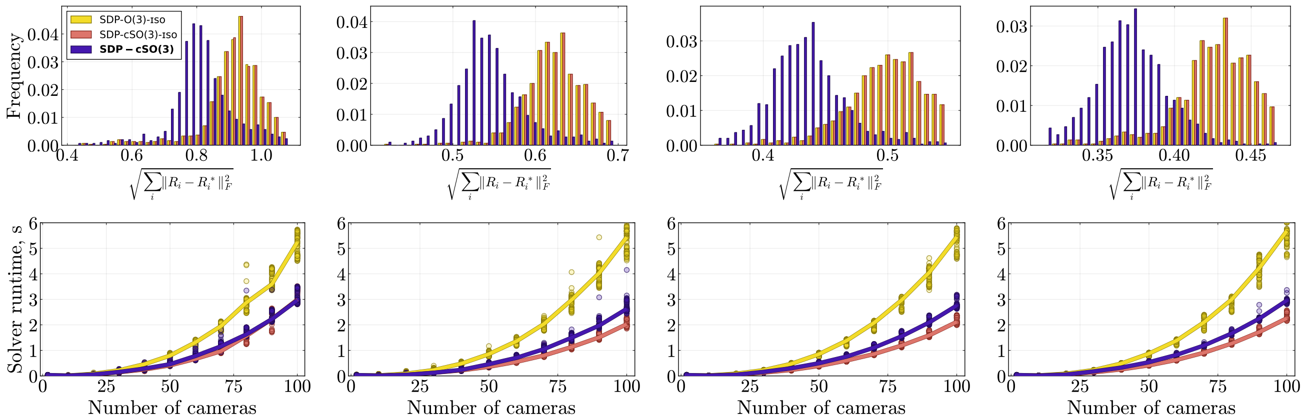

(a) $p = 0.2$ (b) $p = 0.4$ (c) $p = 0.6$ (d) $p = 0.8$

Figure 3. Histograms of rotation errors wrt. ground truth (top). Corresponding solver runtimes (s) wrt. number of cameras, for different fractions p of observed relative rotations (bottom). Scatter plots represent all instances, and the solid lines are the respective medians.

90%. For each configuration pair, we run 100 instances of the problem. In the top row of Figure 3, we report the distance $\sqrt{\sum_i \|R_i - R_i^*\|_F^2}$. As expected, using the proposed method (i.e., both new objective and new relaxation) consistently leads to lower errors. We record the time-to-solution of the solver, shown in the bottom row of Figure 3. In this synthetic setting, the proposed relaxation is more efficient. Our intuition is that the introduction of local constraints—$X_{ij} \in \text{conv}(SO(3))$—reduces the oscillations in the employed first-order SDP solver and therefore leads to faster time-to-solution for optimizing both isotropic and anisotropic costs with the proposed constraints.

5.2. Real experiments

We present the results of applying the proposed framework to a number of public structure from motion datasets [8] from [17, 32]. For every pair of images (with 5 or more point matches), we use RANSAC and the 5 point solver [31] to generate an initial solution to the pairwise geometry which is refined, using bundle adjustment [45], into a local minimum of the sum-of-squared reprojection errors. Here we fix gauge freedom by assuming that the first camera is fixed, the second camera is at distance 1 from the first, and represent 3D points with homogeneous coordinates with norm 1. We then compute the second order Taylor approximation of the objective function and marginalize over the position of the second camera and 3D points. This gives us a quadratic function, in the axis-angle variable that represents the relative rotation, as described in Section 4.1 and Figure 1, that we use for the anisotropic objective.

In Table 1 we report the obtained rank for each method and the distance $\sqrt{\sum_i \|R_i - R_i^*\|_F^2}$, where the ground truth $\{R_i\}_i$ is from the final reconstruction obtained using the pipeline of [32]. When SDP-O(3)-ANISO does not give a solution X^* of rank 3, we compute the closest rank 3 approximation to X^* by setting all but the three largest eigenvalues to 0. We then compute a factorization $X^* = VV^\top$, where V has 3 columns. From each consecutive 3×3 block V_i of V we then compute the closest rotation matrix which we take to be the approximate solution R_i^*.

The proposed method SDP-cSO(3) gives more accurate estimations in all but one of the datasets. We remark that there is no guarantee that anisotropy (or a more correct statistical model in general) will result in better estimates for every noise realization, and all we can expect from MLE is that—on average—it will yield better results. We only have access to a single noise realization for real datasets, and the "ground-truth" solution generated by SfM software may be biased. Our approach is also as fast as the standard method SDP-O(3)-ISO in half of the datasets. The results also confirm that the standard semidefinite relaxation fails on real data when anisotropic costs are used (SDP-O(3)-ANISO).

6. Conclusions

In this work, we propose to incorporate anisotropic costs in a certifiably optimal rotation averaging framework. We demonstrate how existing solvers designed for chordal distances fail to provide sensible solutions for the new objective as these ignore the determinant constraint. Our new SDP formulation enforces $\text{conv}(SO(3))$, and our empirical evaluation shows that this relaxation is able to recover global minima. We evaluated the new approach on both synthetic and real datasets and obtained more accurate reconstructions than isotropic methods in most cases, confirming recent observations based on local optimization. We believe that our work serves as a good entry point for further investigating the anisotropic rotation averaging setting. This work focuses mainly on modeling aspects and under-

[8] Available at: https://www.maths.lth.se/matematiklth/personal/calle/dataset/dataset.html

Dataset		Method	rank(X^*)	$\sqrt{\sum_i \|R_i - R_i^*\|_F^2}$	Runtime, s
LU Sphinx	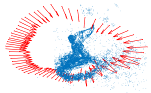	SDP-O(3)-ISO	3	0.0944	2
70 cameras		SDP-O(3)-ANISO	7	18.6037	460
85% indef.		**SDP-cSO(3)**	3	**0.0740**	5
Round Church	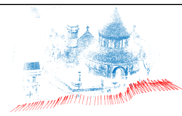	SDP-O(3)-ISO	3	0.1399	6
92 cameras		SDP-O(3)-ANISO	6	26.3808	632
98% indef.		**SDP-cSO(3)**	3	**0.1267**	55
UWO	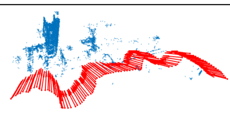	SDP-O(3)-ISO	3	0.3142	14
114 cameras		SDP-O(3)-ANISO	6	22.6873	1929
77% indef.		**SDP-cSO(3)**	3	**0.2274**	7
Tsar Nikolai I	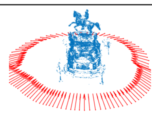	SDP-O(3)-ISO	3	0.1170	7
89 cameras		SDP-O(3)-ANISO	6	26.8944	1245
87% indef.		**SDP-cSO(3)**	3	**0.0534**	5
Vercingetorix	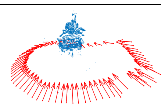	SDP-O(3)-ISO	3	0.3146	2
69 cameras		SDP-O(3)-ANISO	6	14.8244	242
77% indef.		**SDP-cSO(3)**	3	**0.2910**	4
Eglise Du Dome	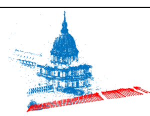	SDP-O(3)-ISO	3	0.0546	4
69 cameras		SDP-O(3)-ANISO	6	20.1954	1840
98% indef.		**SDP-cSO(3)**	3	**0.0487**	5
King's College	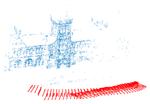	SDP-O(3)-ISO	3	0.1656	4
77 cameras		SDP-O(3)-ANISO	6	17.6508	681
100% indef.		**SDP-cSO(3)**	3	**0.0796**	83
Kronan	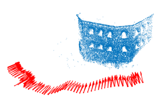	SDP-O(3)-ISO	3	**0.2149**	18
131 cameras		SDP-O(3)-ANISO	6	22.6035	2997
98% indef.		**SDP-cSO(3)**	3	0.3892	201
Alcatraz	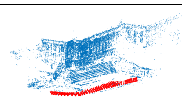	SDP-O(3)-ISO	3	0.1763	19
133 cameras		SDP-O(3)-ANISO	6	22.7623	1931
97% indef.		**SDP-cSO(3)**	3	**0.1268**	107
Mus. Barcelona		SDP-O(3)-ISO	3	0.2255	22
133 cameras		SDP-O(3)-ANISO	7	30.3428	669
97% indef.		**SDP-cSO(3)**	3	**0.1310**	27
Temple Singapore	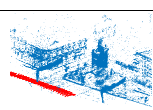	SDP-O(3)-ISO	3	0.2646	31
157 cameras		SDP-O(3)-ANISO	6	26.0824	2313
97% indef.		**SDP-cSO(3)**	3	**0.1696**	628

Table 1. Results on real data from [17, 32]. SDP-O(3)-ISO uses the regular chordal distance and ignores the $\text{conv}(SO(3))$ constraints, SDP-O(3)-ANISO uses the proposed anisotropic objective but ignores the $\text{conv}(SO(3))$ constraints, **SDP-cSO(3)** is the proposed approach that uses both the anisotropic objective and the $\text{conv}(SO(3))$ constraints. The first column provides with the information about the dataset: the number of cameras and the percentage of indefinite matrices M_{ij}.

standing the shortcomings of existing relaxations. We aim to leverage standard SDP solvers; however, the increased strength of our SDP formulation often comes at the cost of more expensive optimization. Designing a dedicated efficient algorithm is an important future direction.

References

[1] MOSEK ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019. 2, 3

[2] Mica Arie-Nachimson, Shahar Z Kovalsky, Ira Kemelmacher-Shlizerman, Amit Singer, and Ronen Basri. Global motion estimation from point matches. *2012 Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*, pages 81–88, 2012. 6

[3] Timothy D. Barfoot, Connor Holmes, and Frederike Dümbgen. Certifiably optimal rotation and pose estimation based on the cayley map. *The International Journal of Robotics Research*, 44(3):366–387, 2025. 2

[4] Tolga Birdal, Michael Arbel, Umut Simsekli, and Leonidas J Guibas. Synchronizing probability measures on rotations via optimal transport. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1569–1579, 2020. 2

[5] Jesus Briales and Javier Gonzalez-Jimenez. Cartan-sync: Fast and global se(d)-synchronization. *IEEE Robotics and Automation Letters*, 2(4):2127–2134, 2017. 1, 2, 3

[6] Jesus Briales and Javier Gonzalez-Jimenez. Convex global 3d registration with lagrangian duality. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5612–5621, 2017. 2

[7] Álvaro Parra Bustos, Tat-Jun Chin, Anders Eriksson, and Ian Reid. Visual slam: Why bundle adjust? In *2019 International Conference on Robotics and Automation (ICRA)*, page 2385–2391. IEEE Press, 2019. 1

[8] Luca Carlone. Estimation contracts for outlier-robust geometric perception. *Found. Trends Robot*, 11(2–3):90–224, 2023. 2

[9] Luca Carlone, David M. Rosen, Giuseppe Carlo Calafiore, John J. Leonard, and Frank Dellaert. Lagrangian duality in 3d slam: Verification techniques and optimal solutions. *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 125–132, 2015. 1

[10] Luca Carlone, Roberto Tron, Kostas Daniilidis, and Frank Dellaert. Initialization techniques for 3d slam: A survey on rotation estimation and its use in pose graph optimization. *Proceedings - IEEE International Conference on Robotics and Automation*, 2015:4597–4604, 2015. 1, 3

[11] Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):958–972, 2017. 1, 2

[12] Yu Chen, Ji Zhao, and Laurent Kneip. Hybrid rotation averaging: A fast and robust rotation averaging approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10358–10367, 2021. 1, 2

[13] Alessandro Chiuso, Giorgio Picci, and Stefano Soatto. Wide-sense estimation on the special orthogonal group. *Commun. Inf. Syst.*, 8(3):185–200, 2008. 4

[14] Kurt Cornelis, Frank Verbiest, and Luc Van Gool. Drift detection and removal for sequential structure from motion algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1249–1259, 2004. 1

[15] Yuchao Dai, Jochen Trumpf, Hongdong Li, Nick Barnes, and Richard Hartley. Rotation averaging with application to camera-rig calibration. In *Computer Vision – ACCV 2009*, 2010. 1

[16] Frank Dellaert, David M. Rosen, Jing Wu, Robert Mahony, and Luca Carlone. Shonan rotation averaging: Global optimality by surfing so(p)$^{\mathrm{n}}$. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI*, page 292–308, Berlin, Heidelberg, 2020. Springer-Verlag. 1, 2, 3, 6

[17] Olof Enqvist, Fredrik Kahl, and Carl Olsson. Non-sequential structure from motion. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 264–271, 2011. 7, 8

[18] Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging with the chordal distance: Global minimizers and strong duality. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):256–268, 2021. 1, 2, 3

[19] Johan Fredriksson and Carl Olsson. Simultaneous multiple rotation averaging using lagrangian duality. In *Computer Vision – ACCV 2012*, pages 245–258, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. 2

[20] Venu Madhav Govindu. Lie-algebraic averaging for globally consistent motion estimation. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 684–691, Los Alamitos, CA, USA, 2004. IEEE Computer Society. 1, 2

[21] Richard Hartley, Khurrum Aftab, and Jochen Trumpf. L1 rotation averaging using the weiszfeld algorithm. In *CVPR 2011*, pages 3041–3048. IEEE, 2011. 1

[22] Richard Hartley, Jochen Trumpf, Yuchao Dai, and Hongdong Li. Rotation averaging. *International Journal of Computer Vision*, 103(3):267 – 305, 2013. 1, 2

[23] Connor Holmes, Frederike Dümbgen, and Timothy Barfoot. On semidefinite relaxations for matrix-weighted state-estimation problems in robotics. *IEEE Transactions on Robotics*, 40:4805–4824, 2024. 2

[24] Matanya B Horowitz, Nikolai Matni, and Joel W Burdick. Convex relaxations of se (2) and se (3) for visual pose estimation. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1148–1154. IEEE, 2014. 1

[25] Gilad Lerman and Yunpeng Shi. Robust group synchronization via cycle-edge message passing. *Foundations of Computational Mathematics*, 22(6):1665–1741, 2022. 2

[26] Johan Lofberg. Yalmip: A toolbox for modeling and optimization in matlab. In *2004 IEEE international conference on robotics and automation (IEEE Cat. No. 04CH37508)*, pages 284–289. IEEE, 2004. 3

[27] Miles Lubin, Oscar Dowson, Joaquim Dias Garcia, Joey Huchette, Benoît Legat, and Juan Pablo Vielma. JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*, 2023. 6

[28] Gabriel Moreira, Manuel Marques, and João Paulo Costeira. Rotation averaging in a split second: A primal-dual method and a closed-form for cycle graphs. In *Proceedings of the*

IEEE/CVF International Conference on Computer Vision (ICCV), pages 5452–5460, 2021. 1, 2

[29] Pierre Moulon and Pascal Monasse. Unordered feature tracking made fast and easy. In *The 9th European Conference on Visual Media Production*, London, United Kingdom, 2012. 1

[30] Pierre Moulon, Pascal Monasse, Romuald Perrot, and Renaud Marlet. Openmvg: Open multiple view geometry. In *Reproducible Research in Pattern Recognition*, pages 60–74, Cham, 2017. Springer International Publishing. 1

[31] David Nistér. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:756–770, 2004. 2, 7

[32] Carl Olsson and Olof Enqvist. Stable structure from motion for unordered image collections. In *Image Analysis*, pages 524–535, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. 7, 8

[33] Carl Olsson and Anders Eriksson. Solving quadratically constrained geometrical problems using lagrangian duality. In *2008 19th International Conference on Pattern Recognition*, pages 1–5, 2008. 2

[34] Brendan O'donoghue, Eric Chu, Neal Parikh, and Stephen Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 169:1042–1068, 2016. 6

[35] Linfei Pan, Dániel Baráth, Marc Pollefeys, and Johannes L. Schönberger. Global structure-from-motion revisited. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XL*, page 58–77, Berlin, Heidelberg, 2024. Springer-Verlag. 1

[36] Alvaro Parra, Shin-Fang Chng, Tat-Jun Chin, Anders Eriksson, and Ian Reid. Rotation coordinate descent for fast globally optimal rotation averaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4298–4307, 2021. 1, 2

[37] Ralph Tyrrell Rockafellar. *Convex analysis*. Princeton university press, 1997. 5, 6

[38] David Rosen, Kevin Doherty, Antonio Espinoza, and John Leonard. Advances in inference and representation for simultaneous localization and mapping. *Annual Review of Control, Robotics, and Autonomous Systems*, 4, 2021. 2, 3

[39] David M Rosen, Luca Carlone, Afonso S Bandeira, and John J Leonard. Se-sync: A certifiably correct algorithm for synchronization over the special euclidean group. *The International Journal of Robotics Research*, 38(2-3):95–125, 2019. 1, 3

[40] Raman Sanyal, Frank Sottile, and Bernd Sturmfels. Orbitopes. *Mathematika*, 57(2):275–314, 2011. 6

[41] James Saunderson, Pablo A Parrilo, and Alan S Willsky. Semidefinite descriptions of the convex hull of rotation matrices. *SIAM Journal on Optimization*, 25(3):1314–1343, 2015. 6

[42] Yunpeng Shi, Cole M Wyeth, and Gilad Lerman. Robust group synchronization via quadratic programming. In *International Conference on Machine Learning*, pages 20095–20105. PMLR, 2022. 2

[43] Chitturi Sidhartha and Venu Madhav Govindu. It is all in the weights: Robust rotation averaging revisited. In *2021 International Conference on 3D Vision (3DV)*, pages 1134–1143, 2021. 1, 2

[44] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and computational harmonic analysis*, 30(1):20–36, 2011. 2, 3

[45] Bill Triggs, Philip F. McLauchlan, Richard I. Hartley, and Andrew W. Fitzgibbon. Bundle adjustment - a modern synthesis. In *Proceedings of the International Workshop on Vision Algorithms: Theory and Practice*, page 298–372, Berlin, Heidelberg, 1999. Springer-Verlag. 7

[46] Lanhui Wang and Amit Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA*, 2(2):145–193, 2013. 2

[47] Kyle Wilson and David Bindel. On the distribution of minima in intrinsic-metric rotation averaging. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6030–6038, 2020. 1, 2

[48] Kyle Wilson, David Bindel, and Noah Snavely. When is rotations averaging hard? In *Proceedings of ECCV 2016*, 2016. 1, 2

[49] Heng Yang and Luca Carlone. Certifiably optimal outlier-robust geometric perception: Semidefinite relaxations and scalable global optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):2816–2834, 2023. 2

[50] C. Zach, M. Klopschitz, and M. Pollefeys. Disambiguating visual relations using loop constraints. In *CVPR*, pages 1426–1433, 2010. 2

[51] Ganlin Zhang, Viktor Larsson, and Daniel Barath. Revisiting rotation averaging: Uncertainties and robust losses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17215–17224, 2023. 1