

Rotation Averaging in a Split Second: A Primal-Dual Method and a Closed-Form for Cycle Graphs

Gabriel Moreira Manuel Marques João Paulo Costeira
Institute for Systems and Robotics, Instituto Superior Técnico
Av. Rovisco Pais, Lisboa, Portugal
{gmoreira, manuel, jpc}@isr.tecnico.ulisboa.pt

Abstract

A cornerstone of geometric reconstruction, rotation averaging seeks the set of absolute rotations that optimally explains a set of measured relative orientations between them. In spite of being an integral part of bundle adjustment and structure-from-motion, averaging rotations is both a non-convex and high-dimensional optimization problem. In this paper, we address it from a maximum likelihood estimation standpoint and make a twofold contribution. Firstly, we set forth a novel initialization-free primal-dual method which we show empirically to converge to the global optimum. Further, we derive what is to our knowledge, the first optimal closed-form solution for rotation averaging in cycle graphs and contextualize this result within spectral graph theory. Our proposed methods achieve a significant gain both in precision and performance.

1. Introduction

Rotation averaging, also known as group synchronization, is an estimation problem wherein we want to find a set of rotations $\{R_1, \dots, R_n\} \in \text{SO}(p)^n$, where

$$\text{SO}(p) = \{R \in \mathbb{R}^{p \times p} : RR^\top = I, \det(R) = 1\}, \quad (1)$$

that optimally explains a set of m noisy pairwise measurements $\{\tilde{R}_{ij}\}_{i \sim j} \in \text{SO}(p)^m$ of the relative orientations $R_i R_j^\top$. The notation $i \sim j$ refers to the existence of a measurement between rotations i and j . As a sub-problem of several 3D reconstruction tasks, namely bundle adjustment [1, 26], structure-from-motion [23, 17] and camera network calibration [25], rotation averaging is of particular interest in computer vision. Nevertheless, the high-dimensionality of the aforementioned problems and the non-convexity of $\text{SO}(p)$ render this problem difficult.

Under the hypothesis of the Langevin noise model [6, 9] adopted in [13, 12, 3, 20] we formalize rotation averaging




Figure 1: Comparison between our primal-dual method and the solutions produced by Shonan Averaging [12] and SE-Sync [22] for rotation averaging problems adapted from pose graph optimization datasets [8].

as the Maximum Likelihood Estimation (MLE) problem

$$\begin{aligned} & \underset{R_1, \dots, R_n}{\text{minimize}} && \sum_{i \sim j} \|\tilde{R}_{ij} - R_i R_j^\top\|_F^2 \\ & \text{subject to} && R_i \in \text{SO}(3), \quad i = 1, \dots, n. \end{aligned} \quad (2)$$

Contribution Firstly, we present a primal-dual method to solve (2) inspired in optimization algorithms with orthogonality constraints [15]. We show empirically that this algorithm converges to the global optimum when the dual variable is initialized with the graph degree matrix (Fig. 1). Secondly, we put forward the first optimal closed-form solution for rotation averaging problems with a cycle graph topology. This solution allows for the retrieval of machine-precision global optima several orders of magnitude faster than the state-of-the-art and compounds the results in spectral graph theory set forth in [13]. Our code is available at <https://github.com/gabmoreira/maks>.

2. Related work

The literature on rotation averaging spans a large array of optimization schemes, from iterative search methods [2, 12, 13, 14, 21, 22] to closed-form suboptimal solutions [3, 19, 20]. Recently, global optimality has taken the spotlight in rotation averaging papers. Nevertheless, finding the global minimum and thus solving (2) remains a difficult task to accomplish efficiently.

In the domain of iterative algorithms, Gauss-Newton (GN) and Levenberg-Marquardt (LM) methods such as those available in the pose graph optimization frameworks g2o [18] and GTSAM [11] were until recently, the most prominent techniques for solving rotation averaging. The non-convexity of this problem however, makes these techniques initialization dependent. A distributed Riemannian gradient descent algorithm in the manifold of 3D rotations has been proposed by Tron *et al.* [25] but this method is arguably less efficient than GN and LM.

In order to circumvent the retrieval of local optima, which are of no practical interest for the applications we are considering, Eriksson *et al.* [13] derived the dual for problem (2) and set forth a multitude of results pertaining to optimality verification relying on duality theory. Assuming strong duality holds, the solution of the semidefinite program (SDP) corresponding to the dual of problem (2) yields the sought-after set of rotations. To solve this SDP, the authors presented a block coordinate descent method.

Alternative globally optimal strategies have been put forward. The SE-Sync pose graph optimization framework by Rosen *et al.* [22] relies on a Riemannian staircase [5, 7] to solve a SDP relaxation and guarantees globally optimal solutions under a few assumptions on the noise model. Shonan averaging, proposed by Dellaert *et al.* [12] and now part of GTSAM, solves the problem locally on $\text{SO}(3)$ and then increases the dimension of the manifold to start the optimization again. This is carried out iteratively, until a globally optimal solution is attained. This method combines the performance of GN and LM methods with a strategy for guaranteeing global optimality.

As opposed to the methods aforementioned, a number of works have proposed closed-form approximations of the optimal solution. Martinec *et al.* [19] treat the problem as a least-squares and then project the solution to the space of rotations. More recently, Arrigoni *et al.* [4, 3] and Moreira *et al.* [20] have proposed eigenspace-based solutions attending to the fact that for noise-free measurements solving rotation averaging is tantamount to solving an eigenvector equation. Whilst these closed-form approaches may yield satisfactory results for moderate noise levels, it is difficult to ascertain their domain of applicability.

3. Problem statement

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a connected graph and $\{\tilde{R}_{ij}\}_{i \sim j} \in \text{SO}(3)^m$ a set of relative rotation measurements between nodes i and j . Under the assumption of isotropic Langevin noise [6, 9], rotation averaging seeks the set of rotations $\{R_i^*\}_{i \in \mathcal{V}}$, which minimize the chordal distance [16] between each measurement \tilde{R}_{ij} and the respective pairwise estimate $R_i^* R_j^{*\top}$, over all edges of the graph. The MLE estimate [8, 20, 12] is

$$\{R_i^*\}_{i \in \mathcal{V}} = \arg \max_{\{R_i\}_{i \in \mathcal{V}} \in \text{SO}(3)^n} \sum_{i \sim j} \text{tr}(R_i^\top \tilde{R}_{ij} R_j). \quad (3)$$

We now introduce a block-matrix notation that we will use throughout the paper. Let S^p be the set of symmetric $p \times p$ matrices, $I_p \in \mathbb{R}^{p \times p}$ the identity and $0_p \in \mathbb{R}^{p \times p}$ the null matrix. We define the block-vector $R \in \text{SO}(3)^n$ as

$$R := [R_1^\top \quad \dots \quad R_n^\top]^\top \quad (4)$$

and the pairwise block-matrix $\tilde{R} \in S^{3n}$ as

$$\tilde{R} := \begin{cases} I_3 & \text{if } i = j \\ \tilde{R}_{ij} \in \text{SO}(3) & \text{if } i \sim j \\ 0_3 & \text{if } i \not\sim j. \end{cases} \quad (5)$$

The block-entry i, j of \tilde{R} contains the measured rotation of that edge or a null block if $i \not\sim j$. We set rotations from each node to itself as the identity and consider $\tilde{R}_{ji} = \tilde{R}_{ij}^\top$. Defining the cost function

$$f(R) := -\text{tr}(R^\top \tilde{R} R), \quad (6)$$

we can write the optimization problem in (3) as

$$\begin{aligned} & \underset{R}{\text{minimize}} && f(R) \\ & \text{subject to} && R \in \text{SO}(3)^n. \end{aligned} \quad (7)$$

4. Primal-dual method

In this section, we present a novel primal-dual update method to solve problem (7). We will show in Section 6 that in the applications considered, this algorithm succeeds in retrieving the global optimum.

As derived in [13], the Lagrangian for rotation averaging under the orthogonality constraint $R \in \text{O}(3)^n$ is

$$\mathcal{L}(R, \Lambda) = -\text{tr}(R^\top \tilde{R} R) - \text{tr}(\Lambda(I - RR^\top)), \quad (8)$$

where the Lagrange multiplier is the block diagonal matrix $\Lambda = \text{blockdiag}(\Lambda_1, \dots, \Lambda_n)$, with $\Lambda_i \in S^3$. Differentiating (8), we have the stationarity condition

$$(\Lambda - \tilde{R})R = 0 \quad (9)$$

i.e., the optimal rotations are in the kernel of $\Lambda - \tilde{R}$. Conversely, we can obtain Λ from (9) via

$$\Lambda_i = R_i R_i^\top + \sum_{i \sim j} \tilde{R}_{ij} R_j R_i^\top, \quad i = 1, \dots, n. \quad (10)$$

Our primal-dual method consists of combining (9) and (10) with projections to $\text{SO}(3)^n$ and S^{3n} respectively, in order to create primal and dual feasible update rules.

Primal update Given an estimate of the dual variable at the k -th iteration, which we denote by Λ^k , Eq. (9) will in general not have a solution in $\text{SO}(3)^n$. We resort thus to an approximation. In order to avoid the trivial solution $R = 0$, we look for X on the Stiefel manifold

$$\text{St}(3n, 3) := \{X \in \mathbb{R}^{3n \times 3} : X^\top X = I_3\} \quad (11)$$

that minimizes $\text{tr}(X^\top (\Lambda^k - \tilde{R}) X)$. We then project the result to $\text{SO}(3)^n$ by solving n Procrustes problems [24]. Our primal updates consist thus of

$$X^{k+1} = \arg \min_{X \in \text{St}(3n, 3)} \text{tr}(X^\top (\Lambda^k - \tilde{R}) X) \quad (12)$$

$$R^{k+1} = \arg \min_{R \in \text{SO}(3)^n} \|R - X^{k+1}\|_F^2. \quad (13)$$

The solution of (12) is given by the three eigenvectors of $\Lambda^k - \tilde{R}$ associated with the three smallest eigenvalues. These eigenspaces can be computed efficiently by means of sparse symmetric eigensolvers. The optimization problem in (13) can be solved via singular value decompositions of 3×3 matrices. As demonstrated in [13], if a primal-dual pair (R^*, Λ^*) verifies the stationarity condition (9), then $\Lambda^* - \tilde{R} \succeq 0$ is sufficient for strong duality to hold and for (R^*, Λ^*) to be optimal. Thus, (12) allows for an optimality assessment at each iteration.

Dual update Drawing from the work of Gao *et al.* on optimization problems with orthogonality constraints [15], we form the dual update by symmetrizing (10). Let $\Psi(X)$ denote the projection to S^3 i.e., $\Psi(X) := \frac{1}{2}(X + X^\top)$. Given a primal variable estimate R^k , we compute Λ according to

$$\Lambda_i^k = R_i^k R_i^k{}^\top + \Psi\left(\sum_{j \sim i} \tilde{R}_{ij} R_j^k R_i^k{}^\top\right), \quad i = 1, \dots, n \quad (14)$$

Initialization Instead of initializing the aforementioned primal-dual updates with an estimate of the primal variable i.e., a set of rotation estimates, we leverage the fact that for noise-free measurements Λ^* depends only on the graph topology. In this case, the optimal rotations R^* verify

$$((D + I) \otimes I_3 - \tilde{R}) R^* = 0, \quad (15)$$




Figure 2: Cosine of the principal angle between the noisy and noise-free eigenspaces of $\Lambda_{\text{nf}} - \tilde{R}$ for a graph with 50 nodes and different algebraic connectivities ρ (Fiedler value). Averaged over 3×10^4 simulations.

where $D \in \mathbb{R}^{n \times n}$ is the graph degree matrix. The optimal Lagrange multiplier for noise-free measurements is thus

$$\Lambda_{\text{nf}} := (D + I) \otimes I_3, \quad (16)$$

as stated in [3, 20]. In our method, we set $\Lambda^0 = \Lambda_{\text{nf}}$. To understand why this initialization allows the primal-dual iterations to attain optimality, we show empirically that for moderate noise levels, the subspace containing the ground-truth rotations and the subspace which solves (12) at $k = 0$ are close together. Let $\tilde{R}_{\text{nf}} \in \text{S}^{3n}$ denote the ground-truth pairwise block-matrix and $\tilde{R} \in \text{S}^{3n}$ (5) the matrix obtained by perturbing the non-null blocks of \tilde{R}_{nf} with Langevin noise (standard deviation σ_{noise}). We represent in Fig. 2 the cosine of the principal angle between the kernel of $\Lambda_{\text{nf}} - \tilde{R}_{\text{nf}}$, which we denote by U_{nf} and the subspace spanned by the three eigenvectors of $\Lambda_{\text{nf}} - \tilde{R}$ associated with the smallest eigenvalues, which we denote by U . The cosine is computed according to $\cos^2 = \text{tr}(U^\top U_{\text{nf}} U_{\text{nf}}^\top U) / 3$. As expected, these subspaces are closer for well connected graphs (large Fiedler value ρ). However, even for poorly connected ones the cosine is on average close to 1.

In Algorithm 1 we show how our primal-dual updates were implemented. This method will henceforth be referred to as Rotation Averaging in a Split Second (RAveSS). The parameter σ used in the sparse eigensolver corresponds to the eigenvalue target for the eigenvectors we are computing. Since our primal update is achieved by solving (12), we pick $\sigma < 0$ such that the three eigenvectors retrieved correspond to the three smallest eigenvalues. Note that prior to projecting the solution of the primal problem to $\text{SO}(3)^n$ we fix the gauge freedom by anchoring the first rotation.

Algorithm 1: RAveSS (Primal-Dual)

Result: R, Λ $\Lambda \leftarrow (D + I_n) \otimes I_3;$ **for** $t \leftarrow 0$ **to** *maxiter* **do** $X, \lambda_1, \lambda_2, \lambda_3 \leftarrow \text{eigensolver}(\Lambda - \tilde{R}, \sigma = -10^{-6});$ $X \leftarrow X X_1^{-1};$ **for** $i \leftarrow 1$ **to** n **do** $U \Sigma V^T \leftarrow \text{SVD}(X_i);$ $R_i \leftarrow U \text{diag}([1 \ 1 \ \det(UV^T)]) V^T;$ **end****for** $i \leftarrow 1$ **to** n **do** $\Lambda_i \leftarrow \frac{1}{2} \sum_{j \sim i} \tilde{R}_{ij} R_j R_i^T +$ $\frac{1}{2} (\sum_{j \sim i} \tilde{R}_{ij} R_j R_i^T)^T + R_i R_i^T;$ **end** $\Lambda \leftarrow \text{blockdiag}(\Lambda_1, \dots, \Lambda_n);$ **if** $\min\{|\lambda_1|, |\lambda_2|, |\lambda_3|\} < \epsilon$ **then** \quad **return** $R, \Lambda;$ **end****end**

5. Cycle graphs: optimal closed-form solution

From a topological standpoint, cycle graphs are the simplest instance of rotation averaging. Nevertheless, they are usually solved via the same iterative methods that cater for the general synchronization problem [13, 12, 22]. In this section we show that rotation averaging problems with an underlying cycle graph topology have closed-form solutions for their stationary points, global optima included. We will first derive the closed-form solutions for one-parameter subgroups of $\text{SO}(3)$. We then show that, in the general case, there is a basis wherein the block matrix \tilde{R} can be written such that its non-null blocks lie in a one-parameter subgroup of $\text{SO}(3)$. This allows us to optimally solve problem (7).

Cycle error We start by defining the error $E \in \text{SO}(3)$ incurred while traversing the cycle graph starting and ending on the same node. Without loss of generality let

$$E := \prod_{k=0}^{n-1} \tilde{R}_{\text{mod}(k,n)+1, \text{mod}(k+1,n)+1}, \quad (17)$$

with the matrix product in (17) being defined from left to right. In a cycle with 3 nodes *e.g.*, $E = \tilde{R}_{12} \tilde{R}_{23} \tilde{R}_{31}$.

Further, let $\gamma \in [-\pi, \pi]$ be the angle of E , which we denote by $\gamma := \angle(E)$. We define the set of the n -th roots of E as

$$E^{\frac{1}{n}} := \{E_0, E_1, \dots, E_{n-1}\}, \quad (18)$$

with $E_k \in \text{SO}(3)$, $E_k^n = E$ and $\angle(E_k) = \gamma/n - 2k\pi/n$, for $k \in \{0, \dots, n-1\}$.

5.1. One-parameter subgroups of $\text{SO}(3)$

We consider for now one-parameter subgroups of $\text{SO}(3)$ by assuming that the pairwise rotation measurements $\tilde{R}_{12}, \tilde{R}_{23}, \dots, \tilde{R}_{n1}$ share a common axis.

Lemma 1. *For cycle graphs whose edge measurements lie in a one-parameter subgroup of $\text{SO}(3)$, the points*

$$R_i = \left(\prod_{s=1}^{i-1} \tilde{R}_{s,s+1} \right)^T E_k^{i-1}, \quad i \in \{2, \dots, n\} \quad (19)$$

with $R_1 = I_3$, indexed by $k \in \{0, \dots, n-1\}$, are stationary points of problem (7).

Proof. We rewrite the cost function f (6) as

$$f(R) = -3n - 2 \sum_{i \sim j} \text{tr}(\tilde{R}_{ij} R_j R_i^T). \quad (20)$$

Under the hypothesis that the rotations \tilde{R}_{ij} share a common axis, we can restrict our search for R_1, \dots, R_n to this subgroup. Thus, $\angle(\tilde{R}_{ij} R_j R_i^T) = \angle(\tilde{R}_{ij}) - \angle(R_i R_j^T)$. From $\text{tr}(R) = 1 + 2 \cos(\angle(R))$, the trace in (20) becomes

$$\text{tr}(\tilde{R}_{ij} R_j R_i^T) = 1 + 2 \cos(\angle(\tilde{R}_{ij}) - \angle(R_i R_j^T)). \quad (21)$$

Define the angles $\theta_{ij} := \angle(R_i R_j^T)$, $\tilde{\theta}_{ij} := \angle(\tilde{R}_{ij})$ and the set $\Theta = \{\theta_{12}, \dots, \theta_{n1}\}$. The optimization problem

$$\begin{aligned} & \underset{\Theta}{\text{maximize}} && \sum_{i \sim j} \cos(\tilde{\theta}_{ij} - \theta_{ij}) \\ & \text{subject to} && \theta_{12} + \dots + \theta_{n1} = 2k\pi, \end{aligned} \quad (22)$$

for $k \in \{0, \dots, n-1\}$, is equivalent to (7). Let the residuals be $\tilde{\theta}_{ij} - \theta_{ij} \in [-\pi, \pi]$ and let $y \in \mathbb{R}$ be a dual variable. The Lagrangian for (22) is

$$\mathcal{L}(\Theta, y) = \sum_{i \sim j} \cos(\tilde{\theta}_{ij} - \theta_{ij}) + y \left(\sum_{i \sim j} \theta_{ij} - 2k\pi \right) \quad (23)$$

From (23), we have the sufficient stationarity conditions

$$\exists w \in [-\pi, \pi], \forall_{i \sim j} \tilde{\theta}_{ij} - \theta_{ij} = w, \quad (24)$$

$$\exists k \in \{0, \dots, n-1\}, \sum_{i \sim j} \theta_{ij} = 2k\pi. \quad (25)$$

Summing (24) over all the edges of the cycle graph we have

$$\sum_{i \sim j} \tilde{\theta}_{ij} - \sum_{i \sim j} \theta_{ij} = nw. \quad (26)$$

Combining (25) and (26) with $\sum_{i \sim j} \tilde{\theta}_{ij} = \gamma$ yields

$$\theta_{ij} = \tilde{\theta}_{ij} - \gamma/n + 2k\pi/n. \quad (27)$$

From (19), we have $R_i R_j^\top = \tilde{R}_{ij} E_k^\top$, for $i \sim j$. Thus,

$$\angle(R_i R_j^\top) = \angle(\tilde{R}_{ij}) - \angle(E_k) \quad (28)$$

which is simply (27) since $\angle(E_k) = \gamma/n - 2k\pi/n$. \square

Theorem 2. For cycle graphs whose edge measurements lie in a one-parameter subgroup of $SO(3)$, the point

$$R_i^* = \left(\prod_{k=1}^{i-1} \tilde{R}_{k,k+1} \right)^\top E_0^{i-1}, \quad i \in \{2, \dots, n\} \quad (29)$$

with $R_1^* = I_3$ is a solution of problem (7).

Proof. Invoking Theorem 4.2 of [13] for cycle graphs, strong duality will hold and a solution will be globally optimal if $\forall i \sim j$ the residuals verify $|\tilde{\theta}_{ij} - \theta_{ij}^*| \leq \frac{\pi}{n}$. From (29), the optimal rotations verify

$$R_j^* = E_0 \tilde{R}_{ij}^\top R_i^*, \quad (30)$$

with $\angle(E_0) = \gamma/n$. From (30) we can write

$$\tilde{\theta}_{ij} - \theta_{ij}^* = \angle(\tilde{R}_{ij} R_j^* R_i^{*\top}) = \gamma/n, \quad (31)$$

Since

$$|\gamma/n| \leq \frac{\pi}{n} \quad (32)$$

due to $\gamma \in [-\pi, \pi]$, the solution in (29) is optimal. \square

In cycle graphs, rotation averaging problems in one-parameter subgroups of $SO(3)$ will redistribute the cycle error equitably over all of the edges. If we incur an error of E , with $\angle(E) = \gamma$, the optimal relative rotation $R_i^* R_j^{*\top}$ will have an angular residual of γ/n relative to the respective measurement \tilde{R}_{ij} . By increasing this figure by a multiple of $2\pi/n$ we obtain suboptimal stationary points of (7).

5.2. Optimization in $SO(3)$

We now show that any cycle graph problem in $SO(3)$ has the same expression for its stationary points (19) and global optimum (29) as derived for one-parameter subgroups. We accomplish this by rewriting \tilde{R} in a new basis.

Change-of-basis Define the matrix $U \in SO(3n)$ as

$$U := \text{blockdiag}(U_1, \dots, U_n), \quad (33)$$

with $U_i \in SO(3)$ for $i \in \{1, \dots, n\}$ computed according to

$$U_i := \begin{cases} I_3, & i = 1 \\ \tilde{R}_{i-1,i}^\top U_{i-1}, & i \in \{2, \dots, n\}. \end{cases} \quad (34)$$




Figure 3: The cycle graph problem on the left can be transformed into the problem on the right via a change-of-basis.

Lemma 3. Denote by \tilde{R}' , the matrix \tilde{R} written in the basis U i.e., $\tilde{R}' := U^\top \tilde{R} U$. Then,

$$\tilde{R}' = \begin{bmatrix} I_3 & I_3 & \dots & 0 & E^\top \\ I_3 & I_3 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & I_3 & I_3 \\ E & 0 & \dots & I_3 & I_3 \end{bmatrix}, \quad (35)$$

with $E = \tilde{R}_{12} \tilde{R}_{23} \dots \tilde{R}_{n1}$ being the cycle error (17).

Proof. From (34), the blocks on the lower triangular part of $\tilde{R}' := U^\top \tilde{R} U$ are given by

$$\tilde{R}'_{ij} = \begin{cases} U_n^\top \tilde{R}_{n1} U_1, & i = n, j = 1 \\ U_i^\top \tilde{R}_{ij} \tilde{R}_{ij}^\top U_i, & j = i - 1 \\ U_i^\top U_j, & i = j. \end{cases} \quad (36)$$

It is immediate that for $i = j$ and $j = i - 1$ we have $\tilde{R}'_{ij} = I_3$. It suffices to show that $U_n^\top \tilde{R}_{n1} U_1 = E$. Note that from (34) we have $U_1 = I_3$ and $U_n = \tilde{R}_{n-1,n}^\top \dots \tilde{R}_{23}^\top \tilde{R}_{12}^\top$. Thus,

$$U_n^\top \tilde{R}_{n1} U_1 = \tilde{R}_{12} \tilde{R}_{23} \dots \tilde{R}_{n-1,n} \tilde{R}_{n1}, \quad (37)$$

which equals E by definition of the latter. \square

We can visualize this result in Fig. 3. In cycle graphs, MLE rotation averaging (7) can be solved by concentrating the cycle error E at a single edge. Further, by changing basis, the pairwise measurements I_3 and E belong to a one-parameter subgroup of $SO(3)$. We can thus leverage the results from Section 5.1 to retrieve the global optimum and stationary points of problem (7) in closed-form.

Theorem 4. For cycle graphs with edge measurements in $SO(3)$, the point

$$R_i^* = \left(\prod_{s=1}^{i-1} \tilde{R}_{s,s+1} \right)^\top E_0^{i-1}, \quad i \in \{2, \dots, n\} \quad (38)$$

with $R_1^* = I_3$, is a solution of the problem (7).

Proof. We write f as

$$f(R) = \text{tr}((U^\top R)^\top U^\top \tilde{R}U(U^\top R)). \quad (39)$$

Using the change-of-variables $R' = U^\top R$ and the change-of-basis $\tilde{R}' = U^\top \tilde{R}U$ we have the equivalent problem

$$\begin{aligned} & \underset{R'}{\text{minimize}} && \text{tr}(R'^\top \tilde{R}' R') \\ & \text{subject to} && R' \in \text{SO}(3)^n, \end{aligned} \quad (40)$$

whose edge measurements are I_3 and E according to Lemma 3. These rotations belong to the one-parameter subgroup $t \mapsto \exp(t[\hat{n}]_\times)$, where \hat{n} is the axis of E . Theorem 2 is thus applicable and the solution of (40) is

$$R_i^* = E_0^{i-1}, \quad i \in \{1, \dots, n\}, \quad (41)$$

since $\tilde{R}'_{i,i+1} = I_3$, $\forall i \in \{1, \dots, n-1\}$. It suffices now to write (41) in the old basis vectors according to $R^* = UR'^*$. Since U is block-diagonal, $R_i^* = U_i R_i'^*$. From the definition of U_i (34) we have

$$R_i^* = \left(\prod_{s=1}^{i-1} \tilde{R}_{s,s+1} \right)^\top E_0^{i-1}, \quad i \in \{2, \dots, n\} \quad (42)$$

with $R_1^* = I_3$. \square

As a corollary of Theorem 4, we can take any stationary point of the problem in the new basis (see Lemma 1) and revert to the old basis vectors in order to obtain the corresponding stationary point of problem (7). Thus, the points

$$R_i = \left(\prod_{s=1}^{i-1} \tilde{R}_{s,s+1} \right)^\top E_k^{i-1}, \quad i \in \{2, \dots, n\} \quad (43)$$

with $R_1 = I_3$, are stationary points of (7) indexed by $k \in \{0, \dots, n-1\}$, where the cost function evaluates to

$$f(R) = -3n - 2n \text{tr}(E_k). \quad (44)$$

Since $\text{tr}(E_k) = 1 + 2 \cos(\gamma/n - 2k\pi/n)$ it follows that the greater the number of nodes, the greater the number of local minima near the global optimum. Hence the difficulty of solving rotation averaging optimally.

We conclude this section by showing that in cycle graphs the spectrum of \tilde{R} relates to the values of f at stationary points and can therefore be computed in closed-form.

Theorem 5. *Let $\sigma(\tilde{R})$ denote the spectrum of \tilde{R} . Then,*

$$\begin{aligned} \sigma(\tilde{R}) = & \{1 + 2 \cos(\gamma/n - 2k\pi/n)\}_{k=0, \dots, n-1} \\ & \cup \{1 + 2 \cos(2k\pi/n)\}_{k=0, \dots, n-1}. \end{aligned} \quad (45)$$

The proof is provided in the appendix (Section 8.1).

6. Experimental results

In this section, we evaluate the performance of our primal-dual update method (RAveSS) and our closed-form solution in pose graph datasets and synthetic rotation averaging problems in cycle graphs, respectively. Our algorithms were implemented in C++ and all the tests were conducted on a laptop computer with a 6-core Intel Core i7-9750H@2.6GHz running macOS Big Sur.

6.1. Pose graph datasets

Using seven datasets from the pose graph optimization literature available online [8], we extracted the pairwise rotation measurements from each one in order to generate rotation averaging problems. Some of these datasets contain multiple measurements per edge, from which only one was kept. We compare the performance of RAveSS (primal-dual update method in Algorithm 1) against Shonan Averaging (SA) [12] and SE-Sync [22]. The authors' implementations are available online and we tested them with their default parameters. Since SE-Sync is designed for solving pose graph optimization problems, we set the input translations to zero. The stopping criterion for our method was defined as $|\lambda_0| < 10^{-15}$, which corresponds to tolerance of the Krylov-based eigensolver used.

The results can be observed in Table 1. In order to juxtapose the three methods in terms of the positive semidefiniteness of $\Lambda - \tilde{R}$ *i.e.*, in order to verify optimality, we proceeded as follows. For each estimate computed, we obtained the Lagrange multiplier using the KKT condition in (10) and symmetrized it via the projection Ψ . The columns λ_0 in Table 1 correspond to the minimum eigenvalue of $\Psi(\Lambda) - \tilde{R}$, which is zero if a given solution is optimal and strong duality holds. In addition, we represent the cost function evaluated at the solution produced by RAveSS, denoted by f^* and the difference between this minimum and the minima computed by SA and SE-Sync, denoted by $\delta^* = f^* - f$. The CPU time is shown in seconds for all three algorithms.

The three methods benchmarked reach the global optimum in all seven datasets. While there may be disparities regarding precision, the differences in terms of the minimum attained and the set of rotations produced are negligible in the applications considered. We focus thus on the CPU time of each algorithm. Our primal-dual method attains machine precision of λ_0 , and therefore the global optimum, faster than the two other methods take to stop iterating. If we were to relax the upper bound on our stopping criterion, the CPU times could be brought down even further, without compromising the solution as far as geometric reconstructions are concerned. Plots showcasing the convergence of RAveSS in terms of the positive-semidefiniteness of $\Lambda - \tilde{R}$ for six of the datasets are available in the supplemental material.

Dataset	Graph		RAveSS (ours)			Shonan Averaging			SE-Sync		
	n	m	$ \lambda_0 $	f^*	$t(s)$	$ \lambda_0 $	δ^* (approx.)	$t(s)$	$ \lambda_0 $	δ^* (approx.)	$t(s)$
SmallGrid	125	297	10^{-15}	-2118.202	0.02	10^{-07}	-10^{-04}	0.55	10^{-08}	-10^{-05}	0.09
Garage	1661	6275	10^{-15}	-42632.998	0.06	10^{-05}	-10^{-01}	24.5	10^{-14}	-10^{-10}	0.99
Sphere	2200	8647	10^{-15}	-56981.692	0.36	10^{-07}	-10^{-03}	30.1	10^{-09}	-10^{-05}	2.79
Torus3D	5000	9048	10^{-15}	-69227.058	0.35	10^{-06}	-10^{-02}	98.8	10^{-09}	-10^{-05}	3.86
Cubicle	5750	12486	10^{-15}	-92163.079	0.46	10^{-05}	-10^{-02}	96.8	10^{-08}	-10^{-04}	2.49
Grid3D	8000	22236	10^{-15}	-157206.257	1.78	10^{-07}	-10^{-02}	154.54	10^{-09}	-10^{-04}	11.69
Rim	10195	22251	10^{-15}	-164037.930	5.92	10^{-05}	-10^{-01}	221.63	10^{-12}	-10^{-07}	8.73

Table 1: Comparison between RAveSS (Algorithm 1), Shonan Averaging [11] and SE-Sync [22]. Datasets from [8].

6.2. Cycle graphs

Borrowing the evaluation approach adopted in [12, 13], we tested our closed-form cycle graph solution, entitled C-RAveSS, in synthetic cycle graph datasets. These consisted of random rotation averaging problems with underlying cycle graph structures of different sizes, wherein the ground-truth absolute orientations R_i correspond to rotations around the z-axis, forming a circular trajectory. The synthetic pairwise measurements were simulated by perturbing the relative ground-truth orientations between adjacent nodes by an error matrix obtained from angle-axis representations. The axes were sampled uniformly over the unit sphere. The angles were drawn from a normal distribution with zero mean and standard deviation σ .

We benchmarked C-RAveSS against two baselines, the block coordinate descent method (BCD) [13] used to solve the dual of the dual of problem (7) and the SA algorithm which we also tested in our pose graph experiments. We implemented the former in MATLAB and used the author’s implementation of the latter. Both methods were initialized randomly. Results averaged over 20 simulations can be observed in Table 2. For our solution, we list the smallest eigenvalue of $\Lambda - \tilde{R}$, denoted by $|\lambda_0|$, which certifies that, as we have shown, our solution is optimal to machine precision in all the simulations we ran. In the two rightmost columns, we show the average difference $\bar{\delta}^*$ between our closed-form global minimum and the cost function evaluated at the set of rotations produced by SA and BCD.

Using its default settings, SA retrieved the global optimum in all the tests conducted. Nevertheless, not only does precision wane as the order of the cycle increases but also the average CPU time surges substantially as the number of variables increases. In order to test the BCD method, we first computed the global optimum in each simulation with C-RAveSS. We then used it to set the stopping criterion for the BCD as $\delta^* \leq 10^{-3}$. As evidenced by the average CPU time, convergence flatlined for the largest cycles. Attaining the global minimum to three decimal places using this algorithm took, on average, as much as 96 seconds for $n = 100$

Problem		C-RAveSS		Shonan		BCD	
n	σ (rad)	$ \lambda_0 $	\bar{t} (s)	$\bar{\delta}^*$	\bar{t} (s)	$\bar{\delta}^*$	\bar{t} (s)
20	0.2	10^{-15}	0.00007	10^{-4}	0.11	10^{-3}	0.18
	0.5	10^{-15}	0.00007	10^{-4}	0.12	10^{-3}	0.23
50	0.2	10^{-15}	0.00008	10^{-3}	0.26	10^{-3}	4.48
	0.5	10^{-15}	0.00008	10^{-2}	0.32	10^{-3}	6.80
100	0.2	10^{-15}	0.00009	10^{-2}	0.42	10^{-3}	51.75
	0.5	10^{-15}	0.00009	10^{-2}	0.50	10^{-3}	96.19
200	0.2	10^{-15}	0.00010	10^{-2}	0.74	10^{-3}	n.a.
	0.5	10^{-15}	0.00010	10^{-1}	1.10	10^{-3}	n.a.

Table 2: Comparison between our closed-form solution, Shonan Averaging [12] and the block coordinate descent method (BCD) [13] for random cycle graph problems.



Figure 4: Reconstruction using our closed-form solution for cycle graphs with 38 optimal rotations computed in 70 μs .

and longer than that would be required for $n = 200$. While this may be a shortcoming of our implementation, the orders of magnitude of the CPU time appear to be in accordance with those reported in [12].

These experiments validate our solution as both the fastest and most precise method to solve cycle graphs. As per Table 2, our optimal closed-form represents a performance gain over the state-of-the-art that may be as big as 10000-fold. Fig. 4 illustrates an application of this solution in geometric reconstruction using a cycle graph rotation averaging problem extracted from a larger 3D reconstruction dataset [10].

7. Conclusion

In this paper we presented two contributions to the problem of averaging multiple rotations. Considering the MLE formulation under the hypothesis of Langevin noise, we set forth a primal-dual update method and a closed-form solution for cycle graphs. As demonstrated by our empirical evaluation, the former produces optimal solutions to machine precision in a fraction of the time of existing solvers. Further, it verifies optimality by default at each iteration. The latter is, to our knowledge, the first optimal closed-form to be derived for this class of problems. The performance gain from optimally solving a problem that used to be tackled via iterative search methods, in closed-form, is substantial, as evidenced by the experiments conducted.

8. Appendix

8.1. Proof of Theorem 5

We will prove the result for the spectrum of \tilde{R}' , which is equal to that of \tilde{R} since the matrices are similar. We start by showing that the block-vectors $V^k \in \mathbb{R}^{3n \times 3}$, with

$$V^k := \begin{bmatrix} E_k^0 \\ E_k^1 \\ \vdots \\ E_k^{n-1} \end{bmatrix}, \quad (46)$$

indexed by $k \in \{0, \dots, n-1\}$ span invariant subspaces of \tilde{R}' i.e., $\exists H \in \mathbb{R}^{3 \times 3} : \tilde{R}'V^k = V^kH$. For this step, it suffices to compute $\tilde{R}'V^k$. From (35) and (46)

$$\begin{aligned} \tilde{R}' \begin{bmatrix} E_k^0 \\ \vdots \\ E_k^i \\ \vdots \\ E_k^{n-1} \end{bmatrix} &= \begin{bmatrix} E_k^0 + E_k^1 + E_k^\top E_k^{n-1} \\ \vdots \\ E_k^{i-1} + E_k^i + E_k^{i+1} \\ \vdots \\ EE_k^0 + E_k^{n-2} + E_k^{n-1} \end{bmatrix} \\ &= \begin{bmatrix} E_k^0 \\ \vdots \\ E_k^i \\ \vdots \\ E_k^{n-1} \end{bmatrix} (I_3 + E_k + E_k^\top). \end{aligned} \quad (47)$$

In compact notation, (47) reads as

$$\tilde{R}'V^k = V^k(I_3 + E_k + E_k^\top). \quad (48)$$

Since $I_3 + E_k + E_k^\top \in \mathbb{S}^3$, let its EVD be

$$I_3 + E_k + E_k^\top \stackrel{\text{EVD}}{=} J\Sigma J^\top, \quad (49)$$

with $J \in O(3)$ and $\Sigma \in \mathbb{R}^{3 \times 3}$ diagonal. From (48) we have

$$\tilde{R}'(V^k J) = (V^k J)\Sigma. \quad (50)$$

The diagonal of Σ contains thus three eigenvalues of \tilde{R}' . From (49), these eigenvalues are those of $I_3 + E_k + E_k^\top$ i.e., $\{3, 1+2\cos(\angle(E_k)), 1+2\cos(\angle(E_k))\}$. By definition, $\angle(E_k) = \gamma/n - 2k\pi/n$, thus

$$\{1 + 2\cos(\gamma/n - 2k\pi/n)\}_{k=0, \dots, n-1} \subset \sigma(\tilde{R}), \quad (51)$$

with each eigenvalue having multiplicity 2.

In order to identify the remaining n eigenvalues of \tilde{R} let \hat{n} denote the axis of the cycle error E i.e., $E\hat{n} = \hat{n}$ and $E^\top \hat{n} = \hat{n}$. Define the vectors $z^k \in \mathbb{R}^{3n}$

$$z^k := \begin{bmatrix} 1 \\ \cos\left(1\frac{2k\pi}{n}\right) \\ \cos\left(2\frac{2k\pi}{n}\right) \\ \vdots \\ \cos\left((n-1)\frac{2k\pi}{n}\right) \end{bmatrix} \otimes \hat{n}, \quad (52)$$

indexed by $k \in \{0, \dots, n-1\}$. We have

$$\begin{aligned} \tilde{R}'z^k &= \begin{bmatrix} (1 + \cos\left(\frac{2k\pi}{n}\right) + \cos\left((n-1)\frac{2k\pi}{n}\right))\hat{n} \\ \vdots \\ (\cos\left((i-1)\frac{2k\pi}{n}\right) + \cos\left(i\frac{2k\pi}{n}\right) + \cos\left((i+1)\frac{2k\pi}{n}\right))\hat{n} \\ \vdots \\ (1 + \cos\left((n-2)\frac{2k\pi}{n}\right) + \cos\left((n-1)\frac{2k\pi}{n}\right))\hat{n} \end{bmatrix} \\ &= \left(1 + 2\cos\left(\frac{2k\pi}{n}\right)\right) \begin{bmatrix} 1 \\ \vdots \\ \cos\left(i\frac{2k\pi}{n}\right) \\ \vdots \\ \cos\left((n-1)\frac{2k\pi}{n}\right) \end{bmatrix} \otimes \hat{n} \end{aligned} \quad (53)$$

In compact notation, (53) reads as

$$\tilde{R}'z^k = (1 + 2\cos(2k\pi/n))z^k, \quad (54)$$

for $k \in \{0, \dots, n-1\}$. It follows that

$$\{1 + 2\cos(2k\pi/n)\}_{k=0, \dots, n-1} \subset \sigma(\tilde{R}). \quad (55)$$

Finally, from (51) and (55) we have

$$\begin{aligned} \sigma(\tilde{R}) &= \{1 + 2\cos(\gamma/n - 2k\pi/n)\}_{k=0, \dots, n-1} \\ &\cup \{1 + 2\cos(2k\pi/n)\}_{k=0, \dots, n-1}. \end{aligned} \quad (56)$$

Acknowledgments The authors would like to thank the reviewers for their comments and suggestions. This work was funded by Fundação para a Ciência e Tecnologia, grant [UIDB/50009/2020]. João Paulo Costeira and Manuel Marques were also supported by the European Union's Horizon 2020 project (GA 825619, AI4EU).

References

- [1] Sameer Agarwal, Noah Snavely, Steven M Seitz, and Richard Szeliski. Bundle adjustment in the large. In *European Conference on Computer Vision*, pages 29–42. Springer, 2010. [1](#)
- [2] I. Aloise and G. Grisetti. Chordal Based Error Function for 3-D Pose-Graph Optimization. *IEEE Robotics and Automation Letters*, 5(1):274–281, 2020. [2](#)
- [3] Federica Arrigoni and Andrea Fusiello. Synchronization Problems in Computer Vision with Closed-Form Solutions. *International Journal of Computer Vision*, 128(1):26–52, 2020. [1](#), [2](#), [3](#)
- [4] Federica Arrigoni, Beatrice Rossi, and Andrea Fusiello. Spectral synchronization of multiple views in SE(3). *SIAM Journal on Imaging Sciences*, 9(4):1963–1990, 2016. [2](#)
- [5] N. Boumal. A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints. *arXiv preprint arXiv:1506.00575*, 2015. [2](#)
- [6] Nicolas Boumal, Amit Singer, P. A. Absil, and Vincent D. Blondel. Cramér-Rao bounds for synchronization of rotations. *Information and Inference*, 3(1):1–39, 2014. [1](#), [2](#)
- [7] Nicolas Boumal, Vlad Voroninski, and Afonso Bandeira. The non-convex Burer-Monteiro approach works on smooth semidefinite programs. *Advances in Neural Information Processing Systems*, 29:2757–2765, 2016. [2](#)
- [8] Luca Carlone, David M Rosen, Giuseppe Calafiore, John J Leonard, and Frank Dellaert. Lagrangian duality in 3D SLAM: Verification techniques and optimal solutions. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 125–132. IEEE, 2015. [1](#), [2](#), [6](#), [7](#)
- [9] Y. Chen, S. Huang, L. Zhao, and G. Dissanayake. Cramér-Rao Bounds and Optimal Design Metrics for Pose-Graph SLAM. *IEEE Transactions on Robotics*, pages 1–15, 2021. [1](#), [2](#)
- [10] Sungjoon Choi, Qian-Yi Zhou, Stephen Miller, and Vladlen Koltun. A large dataset of object scans. *arXiv:1602.02481*, 2016. [7](#)
- [11] Frank Dellaert. Factor graphs and GTSAM: A hands-on introduction. Technical report, Georgia Institute of Technology, 2012. [2](#), [7](#)
- [12] Frank Dellaert, David M Rosen, Jing Wu, Robert Mahony, and Luca Carlone. Shonan Rotation Averaging: Global Optimality by Surfing $SO(p)^n$. In *European Conference on Computer Vision*, pages 292–308. Springer, 2020. [1](#), [2](#), [4](#), [6](#), [7](#)
- [13] Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat Jun Chin. Rotation Averaging and Strong Duality. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 127–135, 2018. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)
- [14] 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, 2019. [2](#)
- [15] Bin Gao, Xin Liu, and Ya-xiang Yuan. Parallelizable algorithms for optimization problems with orthogonality constraints. *SIAM Journal on Scientific Computing*, 41(3):A1949–A1983, 2019. [1](#), [3](#)
- [16] R. Hartley, K. Aftab, and J. Trunpf. L1 rotation averaging using the Weiszfeld algorithm. In *CVPR 2011*, pages 3041–3048, 2011. [2](#)
- [17] Y. Kasten, A. Geifman, M. Galun, and R. Basri. GPSfM: Global Projective SFM Using Algebraic Constraints on Multi-View Fundamental Matrices. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3259–3267, 2019. [1](#)
- [18] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. g2o: A general framework for graph optimization. In *IEEE International Conference on Robotics and Automation*, pages 3607–3613, 2011. [2](#)
- [19] Daniel Martinec and Tomáš Pajdla. Robust rotation and translation estimation in multiview reconstruction. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2007. [2](#)
- [20] Gabriel Moreira, Manuel Marques, and Joao Paulo Costeira. Fast Pose Graph Optimization via Krylov-Schur and Cholesky Factorization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1898–1906, 2021. [1](#), [2](#), [3](#)
- [21] S. M. Nasiri, R. Hosseini, and H. Moradi. Novel Parameterization for Gauss–Newton Methods in 3-D Pose Graph Optimization. *IEEE Transactions on Robotics*, pages 1–18, 2020. [2](#)
- [22] 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. *International Journal of Robotics Research*, 38(2-3):95–125, 2019. [1](#), [2](#), [4](#), [6](#), [7](#)
- [23] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. [1](#)
- [24] Peter H Schönemann. A generalized solution of the orthogonal Procrustes problem. *Psychometrika*, 31(1):1–10, 1966. [3](#)
- [25] Roberto Tron and René Vidal. Distributed image-based 3-D localization of camera sensor networks. In *Proceedings of the IEEE Conference on Decision and Control*, pages 901–908, 2009. [1](#), [2](#)
- [26] R. Zhang, S. Zhu, T. Shen, L. Zhou, Z. Luo, T. Fang, and L. Quan. Distributed very large scale bundle adjustment by global camera consensus. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(02):291–303, feb 2020. [1](#)