



## **Making Rotation Averaging Fast and Robust with Anisotropic Coordinate Descent**

Downloaded from: <https://research.chalmers.se>, 2025-12-31 07:25 UTC

Citation for the original published paper (version of record):

Lochman, Y., Olsson, C., Zach, C. (2025). Making Rotation Averaging Fast and Robust with Anisotropic Coordinate Descent. Proceedings of the 36th British Machine Vision Conference 2025

N.B. When citing this work, cite the original published paper.

# Making Rotation Averaging Fast and Robust with Anisotropic Coordinate Descent

Yaroslava Lochman<sup>1</sup>  
lochman@chalmers.se

Carl Olsson<sup>2</sup>  
carl.olsson@math.lth.se

Christopher Zach<sup>1</sup>  
christopher.m.zach@gmail.com

<sup>1</sup> Chalmers University of Technology

<sup>2</sup> Lund University

## Abstract


Anisotropic rotation averaging has recently been explored as a natural extension of respective isotropic methods. In the anisotropic formulation, uncertainties of the estimated relative rotations—obtained via standard two-view optimization—are propagated to the optimization of absolute rotations. The resulting semidefinite relaxations are able to recover global minima but scale poorly with the problem size. Local methods are fast and also admit robust estimation but are sensitive to initialization. They usually employ minimum spanning trees and therefore suffer from drift accumulation and can get trapped in poor local minima. In this paper, we attempt to bridge the gap between optimality, robustness and efficiency of anisotropic rotation averaging. We analyze a family of block coordinate descent methods initially proposed to optimize the standard chordal distances, and derive a much simpler formulation and an anisotropic extension obtaining a fast general solver. We integrate this solver into the extended anisotropic large-scale robust rotation averaging pipeline. The resulting algorithm achieves state-of-the-art performance on public structure-from-motion datasets. Project page: <https://ylochman.github.io/acd>

## 1 Introduction

Rotation averaging is an important building block in global [23, 24, 28] and hybrid [9, 10] structure-from-motion (SfM) pipelines. The goal is to find absolute camera rotations that are globally aligned with the pre-estimated (via two-view optimization) noisy pairwise relative rotations. Traditional methods for rotation averaging use isotropic error measures [11, 12, 17, 29, 37, 38], where deviations between the observed and predicted relative rotations are penalized equally in all directions. A common approach is to minimize the so-called chordal distance [11, 12, 29], which is the least squares distance between the rotation matrices. Since observed relative rotations in SfM are typically not direct measurements, but rather the result of an estimation from image data, covariance/precision matrices describing the certainty of the estimated rotation in all directions, can be extracted.

It is well known [27, 41] that objective functions based on reprojection error in two-view-relative-pose problems can be insensitive to changes in certain directions, making the exact location of the best rotation hard to pinpoint. Since it is desirable to propagate such information to the later stages of the reconstruction pipeline, several recent works have looked

Figure 1: Comparison of the solutions to the isotropic versus anisotropic rotation averaging (RA) using the synthetic loop scene (left). Comparison between the state-of-the-art methods and the proposed **ACD** algorithm on the “boulders” dataset from ETH3D [57] (right).



into integrating uncertainties, through anisotropic objective functions, into rotation averaging. Zhang *et al.* [41] propose an approach to optimize the distances induced by two-view Hessians. Their results illustrate that rotation averaging can be made much more accurate than the previous isotropic versions, achieving camera poses that are closer to ground truth by taking uncertainties into account. The approach is based on local refinement of angle-axis representations of rotations. To achieve a global method, Olsson *et al.* [47] show how to integrate anisotropic cost with chordal distances into semidefinite programs (SDP) [65]. Solving such problems is equivalent to finding the maximum likelihood estimation (MLE) under an anisotropic Langevin distribution, where the anisotropic terms approximate uncertainties in the input rotations obtained as solutions of two-view SfM problems. The resulting SDP turns out to be tight significantly less often than the corresponding isotropic variant. To address that, a stronger convex relaxation is proposed in [47] that is generally capable of finding the global optimum and leads to more accurate solutions than the isotropic counterpart. On the downside, the SDP formulation has a large number of additional constraints, and therefore the employed general purpose solvers suffer from poor scalability, restricting the applicability of the method to small-scale reconstruction problems.

The design of efficient algorithms therefore has been a key challenge for making large scale certifiable anisotropic rotation averaging feasible. In this paper, we try to address this problem. We study the block coordinate descent methods [13, 49] used for optimizing chordal distance objective in isotropic rotation averaging, and extend this framework to incorporate uncertainties. In summary, our contributions are: (1) we propose a simple and efficient algorithm for anisotropic rotation averaging called anisotropic coordinate descent (ACD); (2) we analyze ACD and show empirically that it has a high success rate from different initializations; (3) we propose an efficient pipeline for accurate, fast and robust anisotropic rotation averaging that achieves state-of-the-art performance on the tested structure-from-motion datasets. Figure 1 illustrates the advantages of ACD in accuracy/runtime compared to the recent methods of Parra *et al.* [49] (RCD) and Olsson *et al.* [47] (cSO(3)).

## 2 Related Work

A common formulation of rotation averaging starts with minimizing the sum of chordal distances,  $\sum_{i,j} \|\tilde{\mathbf{R}}_{ij} - \mathbf{R}_i \mathbf{R}_j^\top\|_F^2$ , subject to  $\mathbf{R}_i \in \text{SO}(3)$  [11, 12, 13, 14, 16, 19, 27]. Under the isotropic Langevin noise [8] assumption, this optimization problem corresponds to MLE. Martinec *et al.* [40] relaxes the rotation constraints and proposes a sparse homogeneous linear system to obtain an approximate solution. Convex relaxations using Lagrangian duality

are able to retain orthogonality constraints and have been shown to admit certifiably correct solutions [10, 13, 14, 16]. While the resulting SDP programs can be solved with off-the-shelf solvers, much more efficient dedicated algorithms have also been developed for isotropic rotation averaging. Moreira *et al.* [22] designed a primal-dual sequence algorithm that is very efficient. Another highly efficient method was proposed by Rosen *et al.* [33] for pose graph optimization and further developed by Dellaert *et al.* [11]. It uses a rank-restricted semidefinite relaxation equivalent to the unconstrained Riemannian optimization on a product of Stiefel manifolds. The employed solution method is known as the Riemannian staircase algorithm [9]. Spectral relaxations lack optimality certificates but generally have comparable accuracy to SDP formulations for isotropic problems [10, 12, 33].

To address outliers among the input data, robustified chordal distances have been employed [9, 36]. Wang *et al.* [36] proposed a semidefinite relaxation for the least unsquared deviations (LUD) that also guarantees exact and stable recovery under a specific noise model. The resulting relaxation is solved using an alternating direction method [36]. A similar approach is used in [9] to derive convex relaxations and a posteriori guarantees for robust SE(2) synchronization with the LUD,  $l_1$  and Huber loss. Lie-algebraic rotation averaging [17] locally minimizes  $\sum_{i,j} \|\log(\tilde{\mathbf{R}}_{ij}\mathbf{R}_i\mathbf{R}_j^\top)\|_F^2$ . Chatterjee *et al.* [5, 6] use the first-order distance approximation to obtain an iterative refinement scheme. To make the approach robust, they propose an iterative  $l_1$  minimization followed by iteratively reweighted least squares.

Local optimizers rely on accurate initialization to ensure convergence to a good optimum. Exact inference along the edges of a minimum spanning tree (MST) of the camera graph is a common initialization strategy [8, 9, 29, 40]. Since MSTs lack redundancy, measurement noise can not be suppressed resulting in accumulated drift that may lead to failure in large problems [9]. Several works propose different heuristics to robustify the initial estimation of absolute rotations. Govindu *et al.* [18] samples multiple MSTs from the view-graph using RANSAC [15], while Zach *et al.* [39] proposes a Bayesian inference approach to identify likely perceptual aliases from sampled cycles of the view graph. Lee *et al.* [20] constructs a hierarchical spanning tree based on the consistency and reliability of the triplet constraints.

### 3 Anisotropic Rotation Averaging

**Preliminaries** All mono-spaced matrices, *e.g.*  $\mathbf{R}$ , are  $3 \times 3$  matrices. All bold matrices, *e.g.*  $\mathbf{R}$ , are formed by stacking the respective mono-spaced matrices either vertically ( $\mathbf{R}$ ,  $\mathbf{S}$ ) or both vertically and horizontally ( $\mathbf{N}$ ,  $\tilde{\mathbf{R}}$ ,  $\mathbf{X}$ ,  $\mathbf{B}$ ). Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  be the view graph with  $n$  vertices in  $\mathcal{V}$  representing absolute rotations  $\{\mathbf{R}_1, \dots, \mathbf{R}_n\}$ , and edges  $\mathcal{E}$  representing estimated two-view relations, *i.e.*  $(i, j) \in \mathcal{E}$  if an estimate of the relative rotation  $\tilde{\mathbf{R}}_{ij}$  exists.

We remark that a common scenario is that  $\tilde{\mathbf{R}}_{ij}$  is obtained from an optimization of a two-view SfM problem. Assuming that  $\tilde{\mathbf{R}}_{ij}$  has been optimized to a minimum, the local behavior of the objective function is described by a Hessian, which can also be seen as the inverse of the covariance matrix. Hence such estimates usually come with uncertainty estimates, in this case describing the change in reprojection error of the two-view objective when deviating from the estimate  $\tilde{\mathbf{R}}_{ij}$ .

We follow the formulation from [27] and seek to optimize anisotropic chordal distances

$$\min_{\mathbf{R} \in \text{SO}(3)^n} - \sum_{(i,j) \in \mathcal{E}} \left\langle \mathbf{M}_{ij} \tilde{\mathbf{R}}_{ij}, \mathbf{R}_j \mathbf{R}_i^\top \right\rangle, \quad (1)$$

where  $M_{ij} = \frac{1}{2} \text{tr}(H_{ij})\mathbf{I} - H_{ij}$ , and  $H_{ij}$  are Hessians of the optimized two-view objective functions. Further, we will equivalently write the objective in (1) as  $-\langle \mathbf{N}, \mathbf{R}\mathbf{R}^\top \rangle$ , where  $\mathbf{N}$  is a symmetric block matrix containing the blocks  $N_{ij} = M_{ij}\tilde{\mathbf{R}}_{ij}$  for  $(i, j) \in \mathcal{E}^1$ , and zeros for  $(i, j) \notin \mathcal{E}$ . An isotropic version of  $\mathbf{N}$  is  $\tilde{\mathbf{R}}$  containing the blocks  $\tilde{\mathbf{R}}_{ij}$ .

### 3.1 Block coordinate descent

To derive an algorithm able to solve (1), we take inspiration from the block coordinate descent methods [13, 14] tackling the isotropic objective. The isotropic SDP formulation analyzed in [14] is as follows,

$$\min_{\mathbf{X} \succeq 0} -\langle \tilde{\mathbf{R}}, \mathbf{X} \rangle \quad \text{s.t. } X_{ii} = \mathbf{I}. \quad (2)$$

Under strong duality the optimal matrix  $\mathbf{X}$  can be factorized into  $\mathbf{X} = \mathbf{R}\mathbf{R}^\top$  giving a certifiably optimal solution  $\mathbf{R}$  to (1). A sub-problem is created by fixing all blocks except for the  $k^{\text{th}}$  block-column/row of  $\mathbf{X}$ , and re-ordering the block-columns and respective block-rows in  $\tilde{\mathbf{R}}$  and  $\mathbf{X}$  such that the  $k^{\text{th}}$  block-column/row becomes the first one

$$\min_{\mathbf{S}} -\langle \mathbf{W}, \mathbf{S} \rangle \quad \text{s.t. } \begin{pmatrix} \mathbf{I} & \mathbf{S}^\top \\ \mathbf{S} & \mathbf{B} \end{pmatrix} \succeq 0, \quad (3)$$

where  $\mathbf{W} = \tilde{\mathbf{R}}_{k,k} \in \mathbb{R}^{3(n-1) \times 3}$  is the  $k^{\text{th}}$  block-column of  $\tilde{\mathbf{R}}$  excluding the  $k^{\text{th}}$  block-row,  $\mathbf{S} \in \mathbb{R}^{3(n-1) \times 3}$  is the block-column matrix of unknowns,  $\mathbf{B}$  contains all fixed blocks. The problem in (3) has a closed-form solution  $\mathbf{S}^* = -\mathbf{B}\mathbf{W}((\mathbf{W}^\top \mathbf{B}\mathbf{W})^{1/2})^\dagger$  which constitutes a single iteration of the block coordinate descent method in [13].

Instead of maintaining the  $3n \times 3n$  matrix of unknowns  $\mathbf{X}$ , the rotation coordinate descent (RCD) algorithm [14] only keeps the “effective”  $3n \times 3$  matrix  $\mathbf{R}$ , motivated by the fact that the solution  $\mathbf{X}$  admits factorization  $\mathbf{R}\mathbf{R}^\top$ . However, as a result, RCD does not increase the rank of the sought matrix  $\mathbf{X}$  above 3 at any iteration, and the method therefore reaches a local minimum over the set of rank-3 matrices. Hence, the solution returned by RCD is not necessarily a global solution of the underlying semidefinite program. Moreover, RCD is practically initialized with rotations, which, as we found, limits the solution of each iteration to  $\text{SO}(3)^n$ . To prove this, we need to show that for any  $3 \times 3$  matrix  $\mathbf{A}$  the following holds

$$\mathbf{A}^\top \left( (\mathbf{A}\mathbf{A}^\top)^{1/2} \right)^\dagger \in \text{SO}(3). \quad (4)$$

This is easy to see if we let  $\mathbf{A} = \mathbf{U}\mathbf{Q}\mathbf{V}^\top$  be the SVD of  $\mathbf{A}$ . Then  $(\mathbf{A}\mathbf{A}^\top)^{1/2} = \mathbf{U}\mathbf{Q}\mathbf{U}^\top$  and  $\mathbf{A}^\top \left( (\mathbf{A}\mathbf{A}^\top)^{1/2} \right)^\dagger = \mathbf{V}\mathbf{Q}\mathbf{U}^\top \mathbf{U}\mathbf{Q}^{-1}\mathbf{U}^\top = \mathbf{V}\mathbf{U}^\top$ . Consequently, if  $\mathbf{R} \in \text{SO}(3)^n$  then

$$\mathbf{R}\mathbf{R}^\top \mathbf{W} \left( (\mathbf{W}^\top \mathbf{R}\mathbf{R}^\top \mathbf{W})^{1/2} \right)^\dagger \in \text{SO}(3)^n. \quad (5)$$

The expression in (5) is the key operation of an RCD iteration (see Algorithm 2 in [14]), and it restricts the search space to  $\text{SO}(3)^n$ . This, however, does not appear to be an obstacle in obtaining accurate reconstructions [14] which we also confirm experimentally. We therefore make use of this finding as shown in the next paragraph.

<sup>1</sup> Note that  $M_{ji}\tilde{\mathbf{R}}_{ji} = (M_{ij}\tilde{\mathbf{R}}_{ij})^\top$ . For  $i < j$ ,  $N_{ij} = (M_{ij}\tilde{\mathbf{R}}_{ij})^\top$  and  $N_{ji} = M_{ij}\tilde{\mathbf{R}}_{ij}$ .

### 3.2 Anisotropic coordinate descent

Now, we present how to leverage block coordinate descent in anisotropic rotation averaging. It was shown in [27] that anisotropic costs require a stronger SDP of the form

$$\min_{\mathbf{X} \succeq 0} -\langle \mathbf{N}, \mathbf{X} \rangle \quad \text{s.t. } X_{ii} = \mathbf{I} \text{ and } X_{ij} \in \text{convhull}(\text{SO}(3)). \quad (6)$$

The additional constraint  $\text{convhull}(\text{SO}(3))$  turns out to be significant. For both real and synthetic cases tested in [27] optimal solutions could only be extracted when using the extra constraints. In contrast, the SDP turned out not to be tight when the convex hull constraints were omitted. Note that while  $\text{convhull}(\text{SO}(3))$  is used in the SDP to achieve convexity, the optimal solution will still have  $X_{ij} \in \text{SO}(3)$ . Based on these findings and the analysis in the previous paragraph, we propose to incorporate  $\text{SO}(3)$  membership as a constraint which leads to a much simpler formulation. We therefore consider the following sub-problem

$$\min_{\mathbf{S} \in \text{SO}(3)^{n-1}} -\langle \mathbf{N}_{\bar{k},k}, \mathbf{S} \rangle \quad \text{s.t. } \begin{pmatrix} \mathbf{I} & \mathbf{S}^\top \\ \mathbf{S} & \mathbf{B} \end{pmatrix} \succeq 0, \quad (7)$$

where  $\mathbf{N}_{\bar{k},k} \in \mathbb{R}^{3(n-1) \times 3}$  is the  $k^{\text{th}}$  block-column of  $\mathbf{N}$  excluding the  $k^{\text{th}}$  block-row. We note that, in addition to allowing anisotropic weights in  $\mathbf{N}$ , the problem in (7) is different from (3) due to the added  $\text{SO}(3)$  constraints. Using the generalized Schur complement (e.g. [9, 40]), we rewrite the PSD constraint as  $\mathbf{I} - \mathbf{S}^\top \mathbf{B}^\dagger \mathbf{S} \succeq 0$  and  $(\mathbf{I} - \mathbf{B}\mathbf{B}^\dagger)\mathbf{S} = 0$ . Matrix  $\mathbf{B}$  can be factorized as  $\mathbf{B} = \mathbf{R}_{\bar{k}}\mathbf{R}_{\bar{k}}^\top$ , where  $\mathbf{R}_{\bar{k}}$  is the column-wise stacked rotations excluding the  $k^{\text{th}}$  one. Similarly,  $\mathbf{S}$  can be decomposed via  $\mathbf{S} = \mathbf{R}_{\bar{k}}\mathbf{R}_k^\top$ . Using the properties of rotations,  $\mathbf{R}_{\bar{k}}^\dagger = \mathbf{R}_{\bar{k}}^\top$  and  $\mathbf{R}_{\bar{k}}^\top \mathbf{R}_{\bar{k}} = \mathbf{I}$ , we rewrite the constraints as follows

$$\mathbf{I} - (\mathbf{R}_{\bar{k}}^\top \mathbf{S})^\top \mathbf{R}_{\bar{k}}^\top \mathbf{S} \succeq 0, \quad (\mathbf{I} - \mathbf{R}_{\bar{k}}\mathbf{R}_{\bar{k}}^\top)\mathbf{S} = 0. \quad (8)$$

Let  $\mathbf{A} = \mathbf{R}_{\bar{k}}^\top \mathbf{S} \in \mathbb{R}^{3 \times 3}$ . From  $(\mathbf{I} - \mathbf{R}_{\bar{k}}\mathbf{R}_{\bar{k}}^\top)\mathbf{S} = 0$  we deduce  $\mathbf{S} = \mathbf{R}_{\bar{k}}\mathbf{A}$ . On the other hand, inserting  $\mathbf{S} = \mathbf{R}_{\bar{k}}\mathbf{R}_k^\top$  yields  $\mathbf{A} = \mathbf{R}_k^\top$ , and  $\mathbf{A}$  is simply the sought-after (inverted) absolute rotation. This implies that by reparametrizing (7) in terms of  $\mathbf{R}_k$  the PSD constraint is automatically satisfied—hence redundant—and we arrive at the following reformulation of (7),

$$\min_{\mathbf{R}_k} -\langle \mathbf{N}_{\bar{k}}^\top \mathbf{R}, \mathbf{R}_k \rangle \quad \text{s.t. } \mathbf{R}_k \in \text{SO}(3), \quad (9)$$

where  $\mathbf{N}_{\bar{k}} \in \mathbb{R}^{3n \times 3}$  is the  $k^{\text{th}}$  block-column of  $\mathbf{N}$ .<sup>2</sup> The solution to (9) is determined as

$$\mathbf{R}_k^* = \text{project}_{\text{SO}(3)}(\mathbf{N}_{\bar{k}}^\top \mathbf{R}) \quad (10)$$

where  $\text{project}_{\text{SO}(3)}(\mathbf{M}) = \mathbf{U} \text{diag}(1, 1, \det(\mathbf{U}\mathbf{V}^\top))\mathbf{V}^\top$  with  $[\mathbf{U}, \sim, \mathbf{V}] = \text{svd}(\mathbf{M})$ . Further, the solution of (7) is given by  $\mathbf{S}^* = \mathbf{R}_{\bar{k}}\mathbf{R}_k^{*\top}$ .<sup>3</sup>

A simple block coordinate descent method for anisotropic rotation averaging (1) called anisotropic coordinate descent (ACD) is presented in Algorithm 1. Note that in isotropic setting ACD arrives at the same solution as RCD [49] while performing much simpler computations. This is because at each iteration, RCD (1) does not leverage the properties of  $\text{SO}(3)$  and instead performs additional matrix multiplications, and (2) rotates the world coordinate system so that  $\mathbf{R}_k$  is  $\mathbf{I}$ , which is possible because of the gauge ambiguity, but is not necessary—ACD directly modifies  $\mathbf{R}_k$  instead.

<sup>2</sup>Note that now we use full block-columns in the coefficient matrix since  $\mathbf{N}_{k,k}$  is zeros.

<sup>3</sup>We would like to point out that (10) exactly coincides with minimizing (1) w.r.t.  $\mathbf{R}_k$  for a fixed  $k$ .

---

**Algorithm 1:** ACD algorithm. Inputs:  $\mathbf{N}, \mathbf{R}^{(0)}$ .

---

```

for  $t \in \{1, \dots\}$  do
  for  $k \in \text{shuffle}(\{1, \dots, n\})$  do
     $\mathbf{R}_i^{(t+1)} \leftarrow \mathbf{R}_i^{(t)} \quad (i \neq k); \quad \mathbf{R}_k^{(t+1)} \leftarrow \text{project}_{\text{SO}(3)}(\mathbf{N}_k^\top \mathbf{R}^{(t)});$ 
  end
end
return  $\mathbf{R}$ 

```

---

Due to the non-communicating blocks of constraints, the method converges to at least a stationary point [60], often a good optimum<sup>4</sup> (see Section 4.5). Finally, note that global SfM assumes an unordered image collection. Inspired by stochastic gradient descent, we shuffle the sequence of absolute rotations at each iteration  $t$ , and empirically verify the slightly improved performance.

## 4 Experiments

We evaluate the proposed and state-of-the-art rotation averaging algorithms on a number of synthetic and real datasets detailed below. We use the following metrics for evaluation: RMS angular error (Angl. err.), area under the cumulative rotation error curve at  $n^\circ$  (AUC @  $n^\circ$ ), and average accuracy (AA) over the set of angular error thresholds  $\{0.1^\circ, 0.2^\circ, \dots, 19.9^\circ, 20^\circ\}$ . To remove gauge freedom, we post-multiply the estimated absolute rotations  $\{\mathbf{R}_i\}_i$  with  $\mathbf{Q} = \text{project}_{\text{SO}(3)}(\sum_i \mathbf{R}_i^\top \mathbf{R}_i^*)$ , where  $\{\mathbf{R}_i^*\}_i$  are the ground-truth absolute rotations.

### 4.1 Synthetic experiments


To generate synthetic scenes, we mainly follow the pipeline of [27] with the difference that the Hessian eigenvalues are sampled from  $\mathcal{U}(a, b)$ ,  $b \sim \mathcal{U}(2a, 100a)$ ,  $a \sim \mathcal{U}(10, 100)$ , and we also synthesize the “loop” scenes (Figure 2, left) in addition to the general random scenes (Figure 2, right). For the “loop” scenes, we generate a set of 100 cameras evenly spread out on a circle and connect each camera with its two neighbors to form relative rotation estimates. In the noiseless setting, this problem has multiple good local minima. We found that all block coordinate descent methods converge to any of these if started from random rotations. Initializing Algorithm 1 with zeros is however sufficient to find a good minimum. In Figure 2 (left), we present the convergence results for 1000 noisy “loop” scenes. We compare ACD to a simple rotation chaining as a baseline<sup>5</sup> and isotropic rotation averaging using RCD. We show median and interquartile range (IQR) of the angular error with respect to ground truth over the iterations. We also report the mean (over scenes) average (over cameras) accuracy, mAA. In a few iterations, ACD converges a more accurate solution.

For the general scenes, we randomly sample absolute orientations of 100 cameras and a fraction  $p \sim \mathcal{U}(0.1, 1)$  of observed relative rotations. In Figure 2 (right), we studied how the Hessians’ eigenvector perturbations—rotating around a random axis by  $\mathcal{N}(0, \sigma)$  degrees—and eigenvalue perturbations—*i.e.* adding uniform noise  $\mathcal{U}(0, \gamma \cdot \text{avg. EV})$ —affect the performance. We show median and IQR of the error reduction in % compared to the isotropic

<sup>4</sup>Similarly to RCD, it has a large convergence basin, but convergence to a global optimum is not guaranteed.

<sup>5</sup>MST followed by chaining is often used as an initialization strategy in robust rotation averaging [8, 9, 60].

Figure 2: Synthetic results. Convergence of ACD from zero to a better solution as compared to RCD and chaining baseline on "loop" scenes (left). Error reduction of ACD versus RCD with the increasing perturbation scale in Hessian estimates on general scenes (right).



RCD method. We see that small perturbations slightly affect the accuracy, however the error reduction is still quite significant (from around 30% down to around 10%).

## 4.2 Experiments on real data

We compare the proposed ACD to the following methods: (i) two SDP solvers optimizing isotropic and anisotropic chordal distances—O(3) [13] and cSO(3) [27], respectively; (ii) two dedicated solvers optimizing isotropic chordal distances—RCD [29] and Shonan [11], (iii) a local robust rotation averaging method of Chatterjee *et al.* [8] optimizing isotropic distances in the axis-angle representation space using L1 averaging followed by iteratively reweighted least squares (L1IRLS), and (iv) a local robust anisotropic rotation averaging method of Zhang *et al.* [41] anisotropic distances in the axis-angle representation space using MAGSAC (Zhang). We use the author’s implementations of their methods, where RCD, L1IRLS, and Zhang are initialized using MST, while Shonan, O(3), and cSO(3) start from random initialization. Our method is initialized with zeros. Since ACD employs camera shuffling (see Algorithm 1), we run it 10 times on each dataset.

We evaluate rotation averaging methods on structure from motion (SfM) datasets [26] and ETH3D MVS (DSLr) datasets [82]. We obtain pairwise relative rotations and Hessian estimates via matched SIFT points for the SfM datasets and data provided as part of the ETH3D datasets. We then follow a standard two-view estimation pipeline, where we (i) obtain the initial relative pose with RANSAC and the 5 point solver [25], (ii) refine the relative pose together with the 3D points using two-view bundle adjustment [84], i.e. by minimizing the sum-of-squared reprojection errors, and (iii) use the Jacobian  $J$  of the residual

Table 1: Accuracy of non-robust rotation averaging on structure from motion datasets [26].

	Angl. err. (°) ↓			AUC @ 1° (%) ↑			AUC @ 5° (%) ↑		
	RCD	Shonan	ACD	RCD	Shonan	ACD	RCD	Shonan	ACD
LU Sphinx	0.46	0.46	0.36	97.14	97.14	98.57	100.00	100.00	100.00
Round Church	0.59	0.59	0.49	89.13	89.13	94.78	100.00	100.00	100.00
UWO	1.19	1.19	0.88	70.17	70.53	87.72	99.12	99.12	100.00
Nikolai I	0.48	0.48	0.22	96.94	96.94	100.00	100.00	100.00	100.00
Vercingetorix	1.54	1.56	1.41	55.07	50.87	60.29	98.55	98.55	98.55
Eglise Du Dome	0.24	0.24	0.18	100.00	100.00	100.00	100.00	100.00	100.00
King’s College	0.78	0.77	0.67	79.22	78.70	82.86	100.00	100.00	100.00
Kronan	0.76	0.76	1.28	94.66	94.66	82.44	99.24	99.24	100.00
Alcatraz	0.62	0.62	0.52	90.23	90.23	93.23	100.00	100.00	100.00
Museum Barcelona	0.79	0.79	0.46	79.70	79.70	99.25	100.00	100.00	100.00
Temple Singapore	0.85	0.86	0.52	75.80	75.80	95.67	100.00	100.00	100.00
Average	0.75	0.76	0.64	84.37	83.97	90.44	99.72	99.72	99.87



Table 2: Accuracy of non-robust rotation averaging on ETH3D MVS (DSLR) [52].

	Angl. err. (°) ↓			AUC @ 1° (%) ↑			AUC @ 5° (%) ↑		
	RCD	Shonan	ACD	RCD	Shonan	ACD	RCD	Shonan	ACD
courtyard	0.28	0.28	0.18	97.37	97.37	100.00	100.00	100.00	100.00
delivery area	0.46	0.51	0.14	97.73	98.86	100.00	100.00	100.00	100.00
electro	0.33	0.33	0.61	97.78	97.78	88.89	100.00	100.00	100.00
facade	0.19	0.19	0.12	100.00	100.00	100.00	100.00	100.00	100.00
meadow	0.96	0.96	1.32	73.33	73.33	80.00	100.00	100.00	100.00
office	1.24	1.24	0.69	88.46	88.46	92.31	96.15	96.15	100.00
pipes	0.93	0.93	0.66	85.71	85.71	92.86	100.00	100.00	100.00
playground	0.11	0.11	0.05	100.00	100.00	100.00	100.00	100.00	100.00
relief	0.38	0.38	0.32	93.55	93.55	96.77	100.00	100.00	100.00
relief 2	0.44	0.44	0.19	93.55	93.55	100.00	100.00	100.00	100.00
terrace	0.79	0.79	0.76	91.30	91.30	95.65	100.00	100.00	100.00
terrains	0.39	0.38	0.27	95.24	95.24	100.00	100.00	100.00	100.00
botanical garden	0.93	1.09	0.63	73.33	55.33	96.67	100.00	100.00	100.00
boulders	0.27	0.27	0.07	100.00	100.00	100.00	100.00	100.00	100.00
bridge	0.16	0.16	0.35	100.00	100.00	98.91	100.00	100.00	100.00
door	0.01	0.01	0.01	100.00	100.00	100.00	100.00	100.00	100.00
exhibition hall	0.19	0.19	0.10	100.00	100.00	100.00	100.00	100.00	100.00
lecture room	0.23	0.23	0.16	100.00	100.00	100.00	100.00	100.00	100.00
living room	0.83	0.83	0.94	84.61	84.61	87.69	100.00	100.00	100.00
observatory	0.05	0.05	0.04	100.00	100.00	100.00	100.00	100.00	100.00
old computer	0.22	0.22	0.19	100.00	100.00	100.00	100.00	100.00	100.00
statue	0.07	0.07	0.01	100.00	100.00	100.00	100.00	100.00	100.00
terrace 2	0.01	0.01	0.01	100.00	100.00	100.00	100.00	100.00	100.00
Average	0.41	0.42	0.34	94.43	93.70	96.95	99.83	99.83	100.00

function to obtain the full Hessian and marginalize out the 3D points and relative position via Schur complement [26]. For a fair comparison (*i.e.* to alleviate potential differences due to different two-view costs) we use the same Hessians in both ACD and Zhang *et al.* [41].

Tables 1 and 2 compare ACD to the other non-robust rotation averaging solvers, RCD [29] and Shonan [40]. These methods are specifically designed to optimize isotropic chordal distances. Results for the SfM datasets are shown in Table 1, where the proposed method outperforms RCD and Shonan (that are both producing very similar results) in all but one dataset. These results are in line with those obtained for the corresponding SDP solvers in [27]. Note that there is a large improvement in the proportion of highly accurate solutions (*i.e.* giving less than 1 degree error). Table 2 presents results for the ETH3D datasets, where ACD gives better estimates in the majority of the scenes. Similarly, we see a particular increase in the percentage of highly accurate solutions. Overall, the tendency is that anisotropic rotation averaging gives more accurate solutions (see also Section A in the Supplementary).

### 4.3 Robust rotation averaging

We test the proposed ACD solver as an initialization strategy for robust refinement and obtain two algorithms, ACD<sub>IRLS</sub> and ACD<sub>AIRLS</sub>. ACD<sub>IRLS</sub> refers to applying iteratively reweighted least squares (see Algorithm 2 in [5]). The robust estimation stage of ACD<sub>AIRLS</sub>

Table 3: Accuracy of robust rotation averaging on structure from motion datasets [26].

	Angl. err. (°) ↓				AUC @ 1° (%) ↑				AUC @ 5° (%) ↑			
	LIRLS	Zhang	ACD <sub>IRLS</sub>	ACD <sub>AIRLS</sub>	LIRLS	Zhang	ACD <sub>IRLS</sub>	ACD <sub>AIRLS</sub>	LIRLS	Zhang	ACD <sub>IRLS</sub>	ACD <sub>AIRLS</sub>
LU Sphinx	0.41	0.37	0.41	0.37	98.57	98.57	98.57	97.14	100.00	100.00	100.00	100.00
Round Church	0.59	0.53	0.59	0.55	92.39	93.48	92.39	94.56	100.00	100.00	100.00	100.00
UWO	1.08	13.52	1.08	0.55	41.23	52.63	41.23	99.12	100.00	95.61	100.00	100.00
Nikolai I	0.33	5.17	0.33	0.27	100.00	98.98	100.00	100.00	100.00	98.98	100.00	100.00
Vercingetorix	0.47	38.20	0.47	0.54	94.20	0.00	94.20	92.75	100.00	0.00	100.00	100.00
Eglise Du Dome	0.22	0.18	0.22	0.17	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
King's College	0.79	0.60	0.79	0.73	76.62	83.82	76.62	80.52	100.00	100.00	100.00	100.00
Kronan	0.40	7.44	0.40	0.35	97.71	87.02	97.71	100.00	100.00	98.47	100.00	100.00
Alcatraz	0.58	0.49	0.58	0.53	91.73	93.23	91.73	91.73	100.00	100.00	100.00	100.00
Museum Barcelona	0.43	8.29	0.43	0.43	100.00	99.25	100.00	97.74	100.00	99.25	100.00	100.00
Temple Singapore	0.59	0.57	0.59	0.56	91.72	92.99	91.72	96.18	100.00	100.00	100.00	100.00
Average	0.53	6.86	0.53	0.46	89.47	81.63	89.47	95.43	100.00	90.21	100.00	100.00

Table 4: Accuracy of robust rotation averaging on ETH3D MVS (DSLRL) [52].

	Angl. err. (°) ↓				AUC @ 1° (%) ↑				AUC @ 5° (%) ↑			
	L1IRLS	Zhang	ACD <sub>IRLS</sub>	ACD <sub>AIRLS</sub>	L1IRLS	Zhang	ACD <sub>IRLS</sub>	ACD <sub>AIRLS</sub>	L1IRLS	Zhang	ACD <sub>IRLS</sub>	ACD <sub>AIRLS</sub>
courtyard	0.17	0.17	0.17	0.15	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
delivery area	0.11	0.10	0.11	0.09	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
electro	0.20	0.40	0.20	0.77	100.00	97.78	100.00	93.33	100.00	100.00	100.00	100.00
facade	0.14	0.11	0.14	0.12	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
meadow	0.47	0.70	0.47	0.68	86.67	93.33	86.67	86.67	100.00	100.00	100.00	100.00
office	0.26	4.50	0.26	0.73	100.00	69.23	100.00	96.15	100.00	96.15	100.00	100.00
pipes	0.18	5.13	0.18	0.69	100.00	0.00	100.00	92.86	100.00	92.86	100.00	100.00
playground	0.08	0.05	0.08	0.04	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
relief	0.07	0.16	0.07	0.14	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
relief 2	0.07	3.36	0.07	0.15	100.00	96.77	100.00	100.00	100.00	96.77	100.00	100.00
terrace	0.12	0.43	0.12	0.07	100.00	95.65	100.00	100.00	100.00	100.00	100.00	100.00
terrains	0.13	0.10	0.13	0.26	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
botanical garden	0.84	0.61	0.84	0.54	90.00	96.67	90.00	100.00	100.00	100.00	100.00	100.00
boulders	0.11	0.06	0.11	0.03	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
bridge	0.09	0.13	0.09	0.10	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
door	0.01	0.02	0.01	0.01	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
exhibition hall	0.07	0.07	0.07	0.06	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
lecture room	0.16	0.10	0.16	0.16	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
living room	0.15	4.59	0.15	0.40	100.00	83.08	100.00	100.00	100.00	95.39	100.00	100.00
observatory	0.05	0.04	0.05	0.04	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
old computer	0.10	0.11	0.10	0.09	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
statue	0.07	0.01	0.07	0.02	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
terrace 2	0.01	0.01	0.01	0.01	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
Average	0.16	0.91	0.16	0.23	98.99	92.72	98.99	98.65	100.00	99.18	100.00	100.00

incorporates the uncertainties into the objective (similarly to [44]). Let  $\omega(\cdot)$  extract the exponential map parameters,  $\{\Delta\hat{\omega}_k\}$  encode the incremental updates of rotations  $\{R_k\}$ , and  $\rho(x) = \frac{x^2}{x^2 + \tau^2}$  be the Geman-McClure kernel function. The following objective is iteratively optimized by ACD<sub>AIRLS</sub>


$$\min_{\{\Delta\hat{\omega}_k\}} \sum_{(i,j) \in \mathcal{E}} \rho \left( \left\| (\Delta\hat{\omega}_j - \Delta\hat{\omega}_i) - \omega(R_j^\top \tilde{R}_{ij} R_i) \right\|_{H_{ij}} \right), \quad (11)$$

followed by updating  $R_k \leftarrow R_k e^{[\Delta\hat{\omega}_k]_\times}$ . In Tables 3 and 4, we compare these methods to the other robust rotation averaging approaches on SfM and ETH3D datasets, respectively. We see that the solution accuracy of ACD<sub>IRLS</sub> is identical to L1IRLS [5], and is higher than that of Zhang *et al.* [44]. ACD<sub>AIRLS</sub> outperforms the other methods on the majority of the SfM datasets and many ETH3D datasets. The method of Zhang *et al.* [44] gives poor solutions in a several cases, *e.g.* UW0, Vercingetorix. We think that might be related to MAGSAC’s hyperparameter that may need to be tuned for each dataset. On the other hand, ACD does not require hyperparameter tuning, and for the refinement step we used the same threshold parameter  $\tau = 5^\circ$  of the robust kernel as in the author’s implementation of [5].

## 4.4 Runtime

We further present the runtime of rotation averaging algorithms. Both absolute and relative tolerance in ACD are set to  $10^{-12}$ . As indicated in Figure 3 (left), ACD is as fast and sometimes faster than the state-of-the-art non-robust solvers RCD [49] and Shonan [41]. More

Figure 3: Runtimes of rotation averaging algorithms.



importantly, ACD is—not surprisingly—orders of magnitude faster than the anisotropic cSO(3) [27]. Further, Figure 3 (right) shows that ACD<sub>IRLS</sub> is on par with the state-of-the-art robust method L<sub>1</sub>IRLS [6] in terms of runtime, while ACD<sub>AIRLS</sub> has comparable runtimes with the method of Zhang *et al.* [40].

## 4.5 Ablation studies

Table 5 compares different initializations of ACD: random rotation matrices (rand), all identities (id), all zeros (0), MST + chaining (mst), and presents corresponding results for the SDP solver of [27] (cSO(3)). We found that in practice all initializations performed similarly well on real data. Results for two datasets—King’s College and facade—indicate that possibly ACD converged to a different solution than cSO(3). Those are however reasonably good judging by the obtained accuracies. Overall, in all runs, ACD gave accurate solutions.

Table 5: Comparison to the SDP solver of [27] (cSO(3)) and effects of different initializations on ACD solver: random rotations (rand), all identities (id), all zeros (0), MST+chaining (mst)

		Angl. err. (°) ↓					AA (%) ↑				
		cSO(3)	ACD-rand	ACD-id	ACD-0	ACD-mst	cSO(3)	ACD-rand	ACD-id	ACD-0	ACD-mst
SIM	LU Sphinx	0.36	0.36	0.36	0.36	0.36	98.76	98.75	98.76	98.76	98.73
	Round Church	0.54	0.48	0.48	0.49	0.48	98.70	98.50	98.53	98.42	98.52
	UWO	0.86	0.89	0.89	0.88	0.86	97.22	97.04	97.03	97.08	97.24
	Nikolai I	0.22	0.22	0.23	0.22	0.22	99.31	99.32	99.28	99.31	99.34
	Vercingetorix	1.42	1.41	1.40	1.41	1.41	94.46	94.54	94.57	94.55	94.57
	Eglise Du Dome	0.21	0.18	0.18	0.18	0.18	99.39	99.52	99.51	99.52	99.54
	King’s College	0.37	0.68	0.67	0.67	0.68	98.69	97.18	97.27	97.24	97.19
	Kronan	1.38	1.25	1.26	1.28	1.27	96.10	96.46	96.44	96.38	96.41
	Alcatraz	0.45	0.52	0.52	0.52	0.52	98.50	98.27	98.27	98.28	98.29
	Museum Barcelona	0.46	0.46	0.45	0.46	0.46	98.81	98.83	98.89	98.81	98.74
	Temple Singapore	0.55	0.55	0.52	0.52	0.59	98.34	97.99	98.16	98.14	97.83
ETH3D	courtyard	0.18	0.18	0.18	0.18	0.18	99.54	99.52	99.53	99.52	99.53
	delivery area	0.13	0.16	0.16	0.14	0.12	99.77	99.53	99.52	99.62	99.76
	electro	0.71	0.61	0.61	0.61	0.61	98.21	98.45	98.46	98.44	98.44
	facade	8.31	0.12	0.13	0.12	0.12	95.10	99.80	99.81	99.81	99.81
	meadow	1.43	1.32	1.32	1.32	1.32	95.10	95.45	95.44	95.46	95.47
	office	0.69	0.69	0.69	0.69	0.69	98.08	98.08	98.08	98.08	98.06
	pipes	0.66	0.66	0.66	0.66	0.66	97.79	97.79	97.79	97.79	97.79
	playground	0.05	0.05	0.05	0.05	0.05	99.99	99.98	99.99	99.97	99.99
	relief	0.33	0.32	0.32	0.32	0.32	99.26	99.42	99.43	99.43	99.43
	relief 2	0.20	0.19	0.19	0.19	0.19	99.48	99.54	99.53	99.53	99.58
	terrace	0.73	0.76	0.76	0.76	0.76	98.63	98.54	98.54	98.53	98.54
	terrains	0.34	0.27	0.27	0.27	0.27	98.80	99.19	99.21	99.20	99.20
	botanical garden	0.58	0.58	0.58	0.63	0.56	97.48	97.48	97.49	97.22	97.55
	boulders	0.09	0.06	0.07	0.07	0.07	99.92	99.97	99.94	99.96	99.95
	bridge	0.52	0.38	0.33	0.35	0.33	98.82	98.49	98.78	98.66	98.74
	door	0.02	0.01	0.02	0.01	0.01	100.00	100.00	100.00	100.00	100.00
	exhibition hall	0.10	0.10	0.10	0.10	0.10	99.82	99.83	99.83	99.83	99.83
	lecture room	0.16	0.16	0.16	0.16	0.16	99.72	99.71	99.71	99.71	99.72
	living room	1.22	1.09	0.92	0.94	0.94	96.16	96.58	97.12	97.06	97.07
	observatory	0.07	0.04	0.05	0.04	0.04	99.96	100.00	100.00	100.00	100.00
	old computer	0.22	0.18	0.20	0.19	0.20	99.31	99.57	99.35	99.48	99.37
	statue	0.01	0.01	0.01	0.01	0.01	100.00	100.00	100.00	100.00	100.00
	terrace 2	0.01	0.01	0.01	0.01	0.01	100.00	100.00	100.00	100.00	100.00
Average		0.69	0.44	0.43	0.44	0.43	98.51	98.63	98.65	98.64	98.65

## 5 Conclusions

In this paper, we address the problem of fast, accurate and robust anisotropic rotation averaging. The proposed ACD solver is the simple extension to the family of block coordinate descent methods tackling the isotropic objective. It is the first dedicated solver that optimizes anisotropic chordal distances. We show that ACD has a large improvement in runtime compared to the state-of-the-art anisotropic SDP solver while delivering as accurate solutions.

## 6 Acknowledgments

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 (grant no. 2023-05341).

## References

- [1] Mica Arie-Nachimson, Shahar Z Kovalsky, Ira Kemelmacher-Shlizerman, Amit Singer, and Ronen Basri. Global motion estimation from point matches. In *2012 Second international conference on 3D imaging, modeling, processing, visualization & transmission*, pages 81–88. IEEE, 2012.
- [2] Nicolas Boumal. A riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints. *arXiv preprint arXiv:1506.00575*, 2015.
- [3] Fennell Burns, David Carlson, Emilie Haynsworth, and Thomas Markham. Generalized inverse formulas using the schur complement. *SIAM Journal on Applied Mathematics*, 26(2):254–259, 1974.
- [4] Luca Carlone and Giuseppe C Calafiore. Convex relaxations for pose graph optimization with outliers. *IEEE Robotics and Automation Letters*, 3(2):1160–1167, 2018.
- [5] Avishek Chatterjee and Venu Madhav Govindu. Efficient and robust large-scale rotation averaging. In *Proceedings of the IEEE international conference on computer vision*, pages 521–528, 2013.
- [6] Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):958–972, 2017.
- [7] 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*, pages 10358–10367, 2021.
- [8] Alessandro Chiuso, Giorgio Picci, and Stefano Soatto. Wide-sense estimation on the special orthogonal group. *Commun. Inf. Syst.*, 8(3):185–200, 2008.
- [9] 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.
- [10] Hainan Cui, Xiang Gao, Shuhan Shen, and Zhanyi Hu. Hsfm: Hybrid structure-from-motion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1212–1221, 2017.
- [11] Frank Dellaert, David M. Rosen, Jing Wu, Robert E. Mahony, and Luca Carlone. Shonan rotation averaging: Global optimality by surfing  $\text{so}(p)^n$ . In *European Conference on Computer Vision*, 2020.
- [12] Kevin J Doherty, David M Rosen, and John J Leonard. Performance guarantees for spectral initialization in rotation averaging and pose-graph slam. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 5608–5614. IEEE, 2022.

- [13] Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging and strong duality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 127–135, 2018.
- [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, 2021. doi: 10.1109/TPAMI.2019.2930051.
- [15] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [16] Johan Fredriksson and Carl Olsson. Simultaneous multiple rotation averaging using lagrangian duality. In Kyoung Mu Lee, Yasuyuki Matsushita, James M. Rehg, and Zhanyi Hu, editors, *Computer Vision – ACCV 2012*, pages 245–258, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-37431-9.
- [17] Venu Madhav Govindu. Lie-algebraic averaging for globally consistent motion estimation. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–684–I–691, Los Alamitos, CA, USA, Jul 2004. IEEE Computer Society. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2004.147>.
- [18] Venu Madhav Govindu. Robustness in motion averaging. In *Asian conference on computer vision*, pages 457–466. Springer, 2006.
- [19] Richard Hartley, Jochen Trumpf, Yuchao Dai, and Hongdong Li. Rotation averaging. *International Journal of Computer Vision*, 103(3):267 – 305, 2013. ISSN 09205691.
- [20] Seong Hun Lee and Javier Civera. Hara: A hierarchical approach for robust rotation averaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15777–15786, 2022.
- [21] Daniel Martinec and Tomas Pajdla. Robust rotation and translation estimation in multiview reconstruction. In *2007 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2007.
- [22] Gabriel Moreira, Manuel Marques, and Joao 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*, pages 5452–5460, 2021.
- [23] Pierre Moulon and Pascal Monasse. Unordered feature tracking made fast and easy. In *CVMP 2012*, page 1, 2012.
- [24] Pierre Moulon, Pascal Monasse, Romuald Perrot, and Renaud Marlet. Openmvg: Open multiple view geometry. In Bertrand Kerautret, Miguel Colom, and Pascal Monasse, editors, *Reproducible Research in Pattern Recognition*, pages 60–74, Cham, 2017. Springer International Publishing. ISBN 978-3-319-56414-2.

- [25] David Nistér. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:756–770, 2004. URL <https://api.semanticscholar.org/CorpusID:886598>.
- [26] Carl Olsson and Olof Enqvist. Stable structure from motion for unordered image collections. In Anders Heyden and Fredrik Kahl, editors, *Image Analysis*, pages 524–535, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. ISBN 978-3-642-21227-7.
- [27] Carl Olsson, Yaroslava Lochman, Johan Malmport, and Christopher Zach. Certifiably optimal anisotropic rotation averaging. *arXiv preprint arXiv:2503.07353*, 2025.
- [28] Linfei Pan, Dániel Baráth, Marc Pollefeys, and Johannes L. Schönberger. Global structure-from-motion revisited. In *European Conference on Computer Vision (ECCV)*, 2024.
- [29] 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, June 2021.
- [30] Liangzu Peng and René Vidal. Block coordinate descent on smooth manifolds: Convergence theory and twenty-one examples. *arXiv preprint arXiv:2305.14744*, 2023.
- [31] 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.
- [32] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3260–3269, 2017.
- [33] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and computational harmonic analysis*, 30(1):20–36, 2011.
- [34] 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, ICCV '99*, page 298–372, Berlin, Heidelberg, 1999. Springer-Verlag. ISBN 3540679731.
- [35] Lieven Vandenbergh and Stephen Boyd. Semidefinite programming. *SIAM review*, 38(1):49–95, 1996.
- [36] 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.
- [37] 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. doi: 10.1109/CVPR42600.2020.00607.
- [38] Kyle Wilson, David Bindel, and Noah Snively. When is rotations averaging hard? In *Proceedings of ECCV 2016*, October 2016.

- [39] Christopher Zach, Manfred Klopschitz, and Marc Pollefeys. Disambiguating visual relations using loop constraints. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1426–1433. IEEE, 2010.
- [40] Fuzhen Zhang. *The Schur complement and its applications*, volume 4. Springer Science & Business Media, 2006.
- [41] 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.