

Shonan Rotation Averaging: Global Optimality by Surfing $SO(p)^n$

Frank Dellaert^{*1}[0000–0002–5532–3566], David M. Rosen^{*2}, Jing Wu¹, Robert Mahony³[0000–0002–7803–2868], and Luca Carlone²

¹ Georgia Institute of Technology, Atlanta, GA {fd27,jingwu}@gatech.edu

² Massachusetts Inst. of Technology, Cambridge, MA {dmrosen,lcarlone}@mit.edu

³ Australian National University, Canberra, Australia Robert.Mahony@anu.edu.au

Abstract. Shonan Rotation Averaging is a fast, simple, and elegant rotation averaging algorithm that is guaranteed to recover globally optimal solutions under mild assumptions on the measurement noise. Our method employs semidefinite relaxation in order to recover provably globally optimal solutions of the rotation averaging problem. In contrast to prior work, we show how to solve large-scale instances of these relaxations using manifold minimization on (only slightly) higher-dimensional rotation manifolds, re-using existing high-performance (but *local*) structure-from-motion pipelines. Our method thus preserves the speed and scalability of current SFM methods, while recovering *globally* optimal solutions.

1 Introduction

Rotation averaging is the problem of estimating a set of n unknown orientations $\mathbf{R}_1, \dots, \mathbf{R}_n \in SO(d)$ from noisy measurements $\bar{\mathbf{R}}_{ij} \in SO(d)$ of the *relative rotations* $\mathbf{R}_i^{-1}\mathbf{R}_j$ between them [18,17]. This problem frequently arises in geometric reconstruction; in particular, it occurs as a sub-problem in bundle adjustment [34,3], structure from motion [29], multi-camera rig calibration [25], and sensor network localization [35]. The development of *fast* and *reliable* algorithms for solving the rotation averaging problem is therefore of great practical interest.

While there are numerous (inequivalent) ways of formalizing the rotation averaging problem in common use [18], unfortunately all of them share the common features of (a) *high dimensionality*, due to the typically large number n of orientations \mathbf{R}_i to be estimated, and (b) *non-convexity*, due to the non-convexity of the space of rotations itself. In consequence, *all* of these approaches lead to optimization problems that are computationally hard to solve in general.

In this work, we address rotation averaging using *maximum likelihood estimation*, as this provides strong statistical guarantees on the quality of the resulting estimate [11,19]. We consider the maximum likelihood estimation problem:

$$\max_{\mathbf{R} \in SO(d)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \text{tr}(\mathbf{R}_i \bar{\mathbf{R}}_{ij} \mathbf{R}_j^T), \quad (1)$$

^{*} Equal contribution

where the $\kappa_{ij} \geq 0$ are concentration parameters for an assumed Langevin noise model [5,9,27]. Our goal in this paper is to develop a *fast* and *scalable* optimization method that is capable of computing *globally* optimal solutions of the rotation averaging problem (1) in practice, despite its non-convexity.

We propose a new, straightforward algorithm, *Shonan Rotation Averaging*, for finding *globally optimal* solutions of problem (1). At its core, our approach simply applies the standard Gauss-Newton or Levenberg-Marquardt methods to a *sequence* of successively *higher-dimensional* rotation averaging problems

$$\max_{\mathbf{Q} \in \text{SO}(p)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \text{tr}(\mathbf{Q}_i \mathbf{P} \bar{\mathbf{R}}_{ij} \mathbf{P}^\top \mathbf{Q}_j^\top), \quad (2)$$

for increasing $p \geq d$. Note that the only difference between (1) and (2) is the $p \times d$ projection matrix $\mathbf{P} \triangleq [\mathbf{I}_d; 0]$, which adapts the measurement matrix $\bar{\mathbf{R}}_{ij}$ to the higher-dimensional rotations \mathbf{Q}_i . We start by running local optimization on (2) with $p = d$, and if this fails to produce a globally-optimal solution, we increase the dimension p and try again. Under mild conditions on the noise, we prove that this simple approach succeeds in recovering a *globally* optimal solution of the rotation averaging problem (1).

A primary contribution of this work is to show how the fast optimization approach developed in [27,26] can be adapted to run directly on the manifold of rotations (rather than the Stiefel manifold), implemented using the venerable Gauss-Newton or Levenberg-Marquardt methods. This approach enables Shonan Averaging to be easily implemented in high-performance optimization libraries commonly used in robotics and computer vision [2,22,12].

2 Related Work

By far the most common approach to addressing smooth optimization problems in computer vision is to apply standard first- or second-order nonlinear optimization methods to compute a critical point of the objective function [24]; this holds in particular for the rotation averaging problem (see [18] generally). This approach is very attractive from the standpoint of computational efficiency, as the low per-iteration cost of these techniques (exploiting the sparsity present in real-world problems) enables these methods to scale gracefully to very large problem sizes; indeed, it is now possible to process reconstruction problems (of which rotation averaging is a crucial part) involving *millions* of images on a single machine [15,20]. However, this computational efficiency comes at the expense of *reliability*, as the use of *local* optimization methods renders this approach susceptible to convergence to bad (significantly suboptimal) local minima.

To address this potential pitfall, several recent lines of work have studied the convergence behavior of local search techniques applied to the rotation averaging problem. One thrust proposes various initialization techniques that attempt to start the local search in low-cost regions of the state space, thereby favoring convergence to the true (global) minimum [23,10,6,30]. Another direction investigates the size of the locally convex region around the global minimizer, in

order to understand when local search is likely to succeed [18,38]. A third class of approaches attempts to evaluate the *absolute* quality of a candidate solution $\hat{\mathbf{R}}$ by employing Lagrangian duality to derive an *upper bound* on $\hat{\mathbf{R}}$'s (global) suboptimality [16,9]. Interestingly, while these last two works employ different representations for rotations ([16] uses quaternions, whereas [9] uses rotation matrices), *both* of the resulting dual problems are semidefinite programs [37], and *both* of these duals are observed to be *tight* unless the measurements $\bar{\mathbf{R}}_{ij}$ are contaminated by large amounts of noise; this fact enables the *certification* of optimality of a global minimizer \mathbf{R}^* .

Motivated by the striking results reported in [16,9], a recent line of work proposes to recover *globally* optimal solutions of the rotation averaging problem from a solution of the Lagrangian dual. Both [16] and [9] propose to compute such solutions using an off-the-shelf SDP solver; however, as general-purpose SDP methods do not scale well with the problem size, this approach is limited to problems involving only a few hundred states. More recently, [13,14] proposed a block-coordinate-descent method specifically tailored to the dual of (1), and showed that this method was 1-2 orders of magnitude faster than the standard SDP algorithm SeDuMi [32]; however, the reported results were still limited to problems involving at most ≈ 300 states. Finally, [27] presents a global optimization approach for pose-graph SLAM based upon the *dual* of the Lagrangian dual (also an SDP), together with a fast optimization scheme that is capable of solving problems involving tens of thousands of poses in a few seconds. However, this optimization approach uses a truncated-Newton Riemannian optimization method [1] employing an *exact* model Hessian, and so cannot be deployed using the Gauss-Newton framework [24] that forms the basis of standard optimization libraries commonly used in computer vision applications [2,22,12].

Our method builds on the approach of [27], but adapts the optimization scheme to run directly on the manifold of rotations using the Gauss-Newton method or a trust-region variant like Levenberg-Marquardt. In this way, it is able to leverage the availability of high-performance software libraries [2,22,12] for performing *local* search on problems of the form (2) while preserving *global optimality* guarantees. In addition, working with the rotation manifold $\text{SO}(p)$, for $p \geq 3$ avoids introducing new, unfamiliar objects like Stiefel manifolds in the core algorithm. The result is a simple, intuitive method for globally optimal rotation averaging that improves upon the scalability of current global methods [13] by an order of magnitude.

3 Gauss-Newton for Rotation Averaging

This section reviews how the Gauss-Newton (GN) algorithm is applied to find a first-order critical point of the rotation averaging problem (1). We first rewrite (1) in terms of minimizing a sum of Frobenius norms:

$$\min_{\mathbf{R} \in \text{SO}(d)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \|\mathbf{R}_j - \mathbf{R}_i \bar{\mathbf{R}}_{ij}\|_{\text{F}}^2. \quad (3)$$

This can in turn be vectorized as:

$$\min_{\mathbf{R} \in \text{SO}(d)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \left\| \text{vec}(\mathbf{R}_j) - \text{vec}(\mathbf{R}_i \bar{\mathbf{R}}_{ij}) \right\|_2^2 \quad (4)$$

where “vec” is the column-wise vectorization of a matrix, and we made use of the fact that $\|\mathbf{M}\|_F^2 = \|\text{vec}(\mathbf{M})\|_2^2$.

Problem (4) does not admit a simple closed-form solution. Therefore, given a feasible point $\mathbf{R} = (\mathbf{R}_1, \dots, \mathbf{R}_n) \in \text{SO}(d)^n$, we will content ourselves with identifying a set of tangent vectors $\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_n \in \mathfrak{so}(d)$ along which we can *locally perturb* each rotation \mathbf{R}_i :

$$\mathbf{R}_i \leftarrow \mathbf{R}_i e^{[\boldsymbol{\omega}_i]} \quad (5)$$

to *decrease* the value of the objective; here $[\boldsymbol{\omega}_i]$ is the $d \times d$ skew-symmetric matrix corresponding to the hat operator of the Lie algebra $\mathfrak{so}(d)$ associated with the rotation group $\text{SO}(d)$. We can therefore reformulate problem (4) as:

$$\min_{\boldsymbol{\omega} \in \mathfrak{so}(d)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \left\| \text{vec}(\mathbf{R}_j e^{[\boldsymbol{\omega}_j]}) - \text{vec}(\mathbf{R}_i e^{[\boldsymbol{\omega}_i]} \bar{\mathbf{R}}_{ij}) \right\|_2^2. \quad (6)$$

In effect, equation (6) replaces the optimization over the *rotations* \mathbf{R}_i in (4) by an optimization over the *tangent vectors* $\boldsymbol{\omega}_i \in \mathfrak{so}(d)$. This is advantageous because $\mathfrak{so}(d)$ is a *linear* space, whereas $\text{SO}(d)$ is not. However, we still cannot solve (6) directly, as the $\boldsymbol{\omega}_i$ enter the objective through the (nonlinear) exponential map.

However, we can *locally approximate* the exponential map to first order as $e^{[\boldsymbol{\nu}]} \approx \mathbf{I} + [\boldsymbol{\nu}]$. Therefore, for any matrix \mathbf{A} we have

$$\text{vec}(\bar{\mathbf{R}} e^{[\boldsymbol{\nu}]} \mathbf{A}) \approx \text{vec}(\bar{\mathbf{R}} (\mathbf{I} + [\boldsymbol{\nu}]) \mathbf{A}) \quad (7)$$

$$= \text{vec}(\bar{\mathbf{R}} \mathbf{A}) + \text{vec}(\bar{\mathbf{R}} [\boldsymbol{\nu}] \mathbf{A}) \quad (8)$$

$$= \text{vec}(\bar{\mathbf{R}} \mathbf{A}) + (\mathbf{A}^\top \otimes \bar{\mathbf{R}}) \text{vec}([\boldsymbol{\nu}]), \quad (9)$$

where we made use of a well-known property of the Kronecker product \otimes . We can also decompose the skew-symmetric matrix $[\boldsymbol{\nu}]$ in terms of the coordinates ν^k of the vector $\boldsymbol{\nu}$ according to:

$$\text{vec}([\boldsymbol{\nu}]) = \text{vec}\left(\sum \nu^k \mathbf{G}_k\right) = \sum \nu^k \text{vec}(\mathbf{G}_k) = \bar{\mathbf{G}}_d \boldsymbol{\nu} \quad (10)$$

where \mathbf{G}_k is the k th generator of the Lie algebra $\mathfrak{so}(d)$, and $\bar{\mathbf{G}}_d$ is the matrix obtained by concatenating the vectorized generators $\text{vec}(\mathbf{G}_k)$ column-wise.

The (local) Gauss-Newton model of the rotation averaging problem (1) is obtained by substituting the linearizations (7)–(9) and decomposition (10) into (6) to obtain a *linear least-squares* problem in the tangent vectors $\boldsymbol{\omega}_i$:

$$\min_{\boldsymbol{\omega} \in \mathfrak{so}(d)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \left\| \mathbf{F}_j \boldsymbol{\omega}_j - \mathbf{H}_i \boldsymbol{\omega}_i - \mathbf{b}_{ij} \right\|_2^2, \quad (11)$$

where the Jacobians \mathbf{F}_j and \mathbf{H}_i and the right-hand side \mathbf{b}_{ij} can be calculated as

$$\mathbf{F}_j \doteq (\mathbf{I} \otimes \bar{\mathbf{R}}_j) \bar{\mathbf{G}}_d \quad (12)$$

$$\mathbf{H}_i \doteq (\bar{\mathbf{R}}_{ij}^\top \otimes \bar{\mathbf{R}}_i) \bar{\mathbf{G}}_d \quad (13)$$

$$\mathbf{b}_{ij} \doteq \text{vec}(\bar{\mathbf{R}}_i \bar{\mathbf{R}}_{ij} - \bar{\mathbf{R}}_j). \quad (14)$$

The *local model* problem (11) can be solved efficiently to produce an optimal *correction* $\boldsymbol{\omega}^* \in \mathfrak{so}(d)^n$. (To do so one can use either direct methods, based on sparse matrix factorization, or preconditioned conjugate gradient (PCG). As an example, to produce the results in Section 5 we use PCG with a block-Jacobi preconditioner and a Levenberg-Marquardt trust-region method.) This correction is then applied to *update* the state \mathbf{R} as in equation (5). Typically, this is done in conjunction with a trust-region control strategy to prevent taking a step that leaves the neighborhood of \mathbf{R} in which the local linear models (11)–(14) well-approximate the objective. The above process is then repeated to generate a *sequence* $\{\boldsymbol{\omega}\}$ of such corrections, each of which improves the objective value, until some termination criterion is satisfied.

We emphasize that while the Gauss-Newton approach is sufficient to *locally improve* an initial estimate, because of non-convexity the final iterate returned by this method is *not* guaranteed to be a minimizer of (1).

4 Shonan Rotation Averaging

4.1 A convex relaxation for rotation averaging

The main idea behind Shonan Averaging is to develop a *convex relaxation* of (1) (which can be solved *globally*), and then exploit this relaxation to search for good solutions of the rotation averaging problem. Following [27, Sec. 3], in this section we derive a convex relaxation of (1) whose minimizers in fact provide *exact, globally optimal* solutions of the rotation averaging problem subject to mild conditions on the measurement noise.

To begin, we rewrite problem (1) in a more compact, matricized form as:

$$f_{\text{MLE}}^* = \min_{\mathbf{R} \in \text{SO}(d)^n} \text{tr}(\bar{\mathbf{L}} \mathbf{R}^\top \mathbf{R}), \quad (15)$$

where $\mathbf{R} = (\mathbf{R}_1, \dots, \mathbf{R}_n)$ is the $d \times dn$ matrix of rotations $\mathbf{R}_i \in \text{SO}(d)$, and $\bar{\mathbf{L}}$ is a symmetric $(d \times d)$ -block-structured matrix constructed from the measurements \mathbf{R}_{ij} . The matrix $\bar{\mathbf{L}}$, known as the *connection Laplacian* [31], is the generalization of the standard (scalar) graph Laplacian to a graph having matrix-valued data \mathbf{R}_{ij} assigned to its edges.

Note that \mathbf{R} enters the objective in (15) only through the product $\mathbf{R}^\top \mathbf{R}$; this is a rank- d positive-semidefinite matrix (since it is an outer product of the rank- d matrix \mathbf{R}), and has a $(d \times d)$ -block-diagonal comprised entirely of identity matrices (since the blocks \mathbf{R}_i of \mathbf{R} are rotations). Our convex relaxation of (15)

is derived simply by replacing the rank- d product $\mathbf{R}^\top \mathbf{R}$ with a *generic* positive-semidefinite matrix \mathbf{Z} having identity matrices along its $(d \times d)$ -block-diagonal:

$$\begin{aligned} f_{\text{SDP}}^* &= \min_{\mathbf{Z} \succeq 0} \text{tr}(\bar{\mathbf{L}}\mathbf{Z}) \\ \text{s.t. } \text{BlockDiag}_{d \times d}(\mathbf{Z}) &= (\mathbf{I}_d, \dots, \mathbf{I}_d). \end{aligned} \quad (16)$$

Problem (16) is a *semidefinite program* (SDP) [37]: it requires the minimization of a linear function of a positive-semidefinite matrix \mathbf{Z} , subject to a set of linear constraints. Crucially, since the set of positive-semidefinite matrices is a convex cone, SDPs are *always* convex, and can therefore be solved *globally* in practice. Moreover, since by construction (15) and (16) share the same objective, and the feasible set of (16) contains every matrix of the form $\mathbf{R}^\top \mathbf{R}$ with \mathbf{R} feasible in (15), we can regard (16) as a convexification of (15) obtained by *expanding the latter's feasible set*. It follows immediately that $f_{\text{SDP}}^* \leq f_{\text{MLE}}^*$. Furthermore, if it so happens that a minimizer \mathbf{Z}^* of (16) admits a factorization of the form $\mathbf{Z}^* = \mathbf{R}^{*\top} \mathbf{R}^*$ with $\mathbf{R}^* \in \text{SO}(d)^n$, then it is clear that \mathbf{R}^* is also a (global) minimizer of (15), since it attains the lower-bound f_{SDP}^* for (15)'s optimal value f_{MLE}^* . The key fact that motivates our interest in the relaxation (16) is that this favorable situation *actually occurs*, provided that the noise on the observations $\bar{\mathbf{R}}_{ij}$ is not too large. More precisely, the following result is a specialization of [27, Proposition 1] to the rotation averaging problem (15).

Theorem 1. *Let \mathbf{L} denote the connection Laplacian for problem (15) constructed from the true (noiseless) relative rotations $\mathbf{R}_{ij} \triangleq \mathbf{R}_i^{-1} \mathbf{R}_j$. Then there exists a constant $\beta \triangleq \beta(\mathbf{L})$ (depending upon \mathbf{L}) such that, if $\|\bar{\mathbf{L}} - \mathbf{L}\|_2 \leq \beta$:*

- (i) *The semidefinite relaxation (16) has a unique solution \mathbf{Z}^* , and*
- (ii) *$\mathbf{Z}^* = \mathbf{R}^{*\top} \mathbf{R}^*$, where $\mathbf{R}^* \in \text{SO}(d)^n$ is a globally optimal solution of the rotation averaging problem (15).*

4.2 Solving the semidefinite relaxation: The Riemannian Staircase

In this section, we describe a specialized optimization procedure that enables the fast solution of large-scale instances of the semidefinite relaxation (16), following the approach of [27, Sec. 4.1].

Theorem 1 guarantees that in the (typical) case that (16) is exact, the *solution* \mathbf{Z}^* that we seek admits a concise description in the factored form $\mathbf{Z}^* = \mathbf{R}^{*\top} \mathbf{R}^*$ with $\mathbf{R}^* \in \text{SO}(d)^n$. More generally, even in those cases where exactness fails, minimizers \mathbf{Z}^* of (16) generically have a rank p not much greater than d , and therefore admit a symmetric rank decomposition $\mathbf{Z}^* = \mathbf{S}^{*\top} \mathbf{S}^*$ for $\mathbf{S}^* \in \mathbb{R}^{p \times dn}$ with $p \ll dn$. We exploit the existence of such low-rank solutions by adopting the approach of [8], and replacing the decision variable \mathbf{Z} in (16) with a symmetric rank- p product of the form $\mathbf{S}^\top \mathbf{S}$. This substitution has the effect of dramatically reducing the size of the search space, as well as rendering the positive-semidefiniteness constraint *redundant*, since $\mathbf{S}^\top \mathbf{S} \succeq 0$ for *any* \mathbf{S} . The

resulting *rank-restricted* version of (16) is thus a standard *nonlinear program* with decision variable the low-rank factor \mathbf{S} :

$$\begin{aligned} f_{\text{SDPLR}}^*(p) &= \min_{\mathbf{S} \in \mathbb{R}^{p \times dn}} \text{tr}(\bar{\mathbf{L}} \mathbf{S}^\top \mathbf{S}) \\ \text{s.t. } \text{BlockDiag}_{d \times d}(\mathbf{S}^\top \mathbf{S}) &= (\mathbf{I}_d, \dots, \mathbf{I}_d). \end{aligned} \quad (17)$$

The identity block constraints in (17) are equivalent to the $p \times d$ block-columns \mathbf{S}_i of \mathbf{S} being orthonormal d -frames in \mathbb{R}^p . The set of all orthonormal d -frames in \mathbb{R}^p is a matrix manifold, the **Stiefel manifold** $\text{St}(d, p)$ [1]:

$$\text{St}(d, p) \doteq \{\mathbf{M} \in \mathbb{R}^{p \times d} \mid \mathbf{M}^\top \mathbf{M} = \mathbf{I}_d\}. \quad (18)$$

We can therefore rewrite (17) as a low-dimensional *unconstrained* optimization over a product of n Stiefel manifolds:

$$f_{\text{SDPLR}}^*(p) = \min_{\mathbf{S} \in \text{St}(d, p)^n} \text{tr}(\bar{\mathbf{L}} \mathbf{S}^\top \mathbf{S}). \quad (19)$$

Now, let us compare the rank-restricted relaxation (19) with the original rotation averaging problem (15) and its semidefinite relaxation (16). Since any matrix \mathbf{R}_i in the (special) orthogonal group satisfies conditions (18) with $p = d$, we have the set of inclusions:

$$\text{SO}(d) \subset \text{O}(d) = \text{St}(d, d) \subset \dots \subset \text{St}(d, p) \subset \dots \quad (20)$$

Therefore, we can regard the rank-restricted problems (19) as a *hierarchy* of relaxations (indexed by the rank parameter p) of (15) that are intermediate between (15) and (16) for $d < p < r$, where r is the lowest rank of any solution of (16). Indeed, if (16) has a minimizer of rank r , then it is clear by construction that for $p \geq r$ any (global) minimizer \mathbf{S}^* of (19) corresponds to a minimizer $\mathbf{Z}^* = \mathbf{S}^{*\top} \mathbf{S}^*$ of (16). However, unlike (16), the rank-restricted problems (19) are no longer convex, since we have reintroduced the (nonconvex) orthogonality constraints in (18). It may therefore not be clear that anything has really been gained by relaxing problem (15) to problem (19), since it seems that we may have just replaced one difficult (nonconvex) optimization problem with another. The key fact that justifies our interest in the rank-restricted relaxations (19) is the following remarkable result [7, Corollary 8]:

Theorem 2. *If $\mathbf{S}^* \in \text{St}(d, p)^n$ is a rank-deficient second-order critical point of (19), then \mathbf{S}^* is a global minimizer of (19), and $\mathbf{Z}^* = \mathbf{S}^{*\top} \mathbf{S}^*$ is a global minimizer of the semidefinite relaxation (16).*

This result immediately suggests the following simple algorithm (the **Riemannian staircase** [7]) for recovering solutions \mathbf{Z}^* of (16) from (19): for some small relaxation rank $p \geq d$, find a second-order critical point \mathbf{S}^* of problem (19) using a local search technique. If \mathbf{S}^* is rank-deficient, then Theorem 2 proves that \mathbf{S}^* is a *global* minimizer of (19), and $\mathbf{Z}^* = \mathbf{S}^{*\top} \mathbf{S}^*$ is a solution of (16). Otherwise, increase the rank parameter p and try again. In the worst possible

case, we might have to take p as large as $dn + 1$ before finding a rank-deficient solution. However, in practice typically only one or two “stairs” suffice – just a *bit* of extra room in (19) vs. (15) is all one needs!

Many popular optimization algorithms only guarantee convergence to *first-order* critical points because they use only limited second-order information [24]. This is the case for the Gauss-Newton and Levenberg-Marquardt methods in particular, where the model Hessian is positive-semidefinite by construction. Fortunately, there is a simple procedure that one can use to test the global optimality of a *first-order* critical point \mathbf{S}^* of (19), and (if necessary) to construct a *direction of descent* that we can use to “nudge” \mathbf{S}^* away from stationarity before restarting local optimization at the next level of the Staircase [4].

Theorem 3. *Let $\mathbf{S}^* \in \text{St}(d, p)^n$ be a first-order critical point of (19), define*

$$\mathbf{C} \triangleq \bar{\mathbf{L}} - \frac{1}{2} \text{BlockDiag}_{d \times d} \left(\bar{\mathbf{L}} \mathbf{S}^{*\top} \mathbf{S}^* + \mathbf{S}^{*\top} \mathbf{S}^* \bar{\mathbf{L}} \right), \quad (21)$$

and let λ_{\min} be the minimum eigenvalue of \mathbf{C} , with corresponding eigenvector $\mathbf{v}_{\min} \in \mathbb{R}^{dn}$. Then:

- (i) *If $\lambda_{\min} \geq 0$, then \mathbf{S}^* is a global minimizer of (19), and $\mathbf{Z}^* = \mathbf{S}^{*\top} \mathbf{S}^*$ is a global minimizer of (16).*
- (ii) *If $\lambda_{\min} < 0$, then the higher-dimensional lifting $\mathbf{S}^+ \triangleq [\mathbf{S}^*; 0] \in \text{St}(d, p + 1)^n$ of \mathbf{S}^* is a stationary point for the rank-restricted relaxation (19) at the next level $p + 1$ of the Riemannian Staircase, and $\dot{\mathbf{S}}^+ \triangleq [0; \mathbf{v}_{\min}^\top] \in T_{\mathbf{S}^+}(\text{St}(d, p + 1)^n)$ is a second-order descent direction from \mathbf{S}^+ .*

Remark 1 (Geometric interpretation of Theorems 1–3). With \mathbf{C} defined as in (21), part (i) of Theorem 3 is simply the standard necessary and sufficient conditions for $\mathbf{Z}^* = \mathbf{S}^{*\top} \mathbf{S}^*$ to be a minimizer of (16) [37]. \mathbf{C} also gives the action of the Riemannian Hessian of (19) on tangent vectors; therefore, if $\lambda_{\min}(\mathbf{C}) < 0$, the corresponding direction of negative curvature $\dot{\mathbf{S}}^+$ defined in part (ii) of Theorem 3 provides a second-order descent direction from \mathbf{S}^+ . Theorem 2 is an analogue of Theorem 3(i). Finally, Theorem 1 follows from the fact that the certificate matrix \mathbf{C} depends continuously upon both the data $\bar{\mathbf{L}}$ and \mathbf{S}^* , and is *always* valid for $\mathbf{S}^* = \mathbf{R}^*$ in the noiseless case; one can then employ a continuity argument to show that \mathbf{C} *remains* a valid certificate for $\mathbf{Z}^* = \mathbf{R}^{*\top} \mathbf{R}^*$ as a minimizer of (16) for sufficiently small noise; see [28, Appendix C] for details.

4.3 Rounding the solution

Algorithm 1 provides a truncated SVD procedure to extract a feasible point $\hat{\mathbf{R}} \in \text{SO}(d)^n$ for the original rotation averaging problem (1) from the optimal factor $\mathbf{S}^* \in \text{St}(d, p)^n$ obtained via the rank-restricted relaxation (19). We need to ensure that $\hat{\mathbf{R}}$ is a *global minimizer* of (1) whenever the semidefinite relaxation (16) is *exact*, and is at least an *approximate* minimizer otherwise. The factor \mathbf{R} of a symmetric factorization $\mathbf{Z} = \mathbf{R}^\top \mathbf{R}$ is only unique up to left-multiplication

Algorithm 1 Rounding procedure for solutions of (19)

Input: A minimizer $\mathbf{S} \in \text{St}(d, p)^n$ of (19).
Output: A feasible point $\hat{\mathbf{R}} \in \text{SO}(d)^n$ for (1).

```

1: function ROUNDSOLUTION( $\mathbf{S}$ )
2:   Compute a rank- $d$  truncated SVD  $\mathbf{U}_d \mathbf{\Xi}_d \mathbf{V}_d^\top$  for  $\mathbf{S}$  and assign  $\hat{\mathbf{R}} \leftarrow \mathbf{\Xi}_d \mathbf{V}_d^\top$ .
3:   Set  $N_+ \leftarrow |\{\hat{\mathbf{R}}_i \mid \det(\hat{\mathbf{R}}_i) > 0\}|$ .
4:   if  $N_+ < \lceil \frac{n}{2} \rceil$  then
5:      $\hat{\mathbf{R}} \leftarrow \text{Diag}(1, \dots, 1, -1) \hat{\mathbf{R}}$ .
6:   end if
7:   for  $i = 1, \dots, n$  do
8:     Set  $\hat{\mathbf{R}}_i \leftarrow \text{NEARESTROTATION}(\hat{\mathbf{R}}_i)$  (see [36])
9:   end for
10:  return  $\hat{\mathbf{R}}$ 
11: end function

```

by some $\mathbf{A} \in \text{O}(d)$. The purpose of lines 3–6 is to choose a representative $\hat{\mathbf{R}}$ with a majority of $d \times d$ blocks $\hat{\mathbf{R}}_i$ satisfying $\det(\hat{\mathbf{R}}_i) > 0$, since these should *all* be rotations in the event (16) is exact [27, Sec. 4.2].

4.4 From Stiefel manifolds to rotations

In this section we show how to reformulate the low-rank optimization (19) as an optimization over the product $\text{SO}(p)^n$ of rotations of p -dimensional space. This is convenient from an implementation standpoint, as affordances for performing optimization over rotations are a standard feature of many high-performance optimization libraries commonly used in robotics and computer vision [2, 22, 12].

The main idea underlying our approach is the simple observation that since the columns of a rotation matrix $\mathbf{Q} \in \text{SO}(p)$ are orthonormal, then in particular the first d columns $\mathbf{Q}_{[1:d]}$ form an orthonormal d -frame in p -space, i.e., the submatrix $\mathbf{Q}_{[1:d]}$ is itself an element of $\text{St}(d, p)$. Let us define the following projection, which simply extracts these first d columns:

$$\begin{aligned} \pi: \text{SO}(p) &\rightarrow \text{St}(d, p) \\ \pi(\mathbf{Q}) &= \mathbf{Q}\mathbf{P} \end{aligned} \tag{22}$$

where $\mathbf{P} = [\mathbf{I}_d; 0]$ is the $p \times d$ projection matrix appearing in (2). It is easy to see that π is surjective for any $p > d$:⁴ given any element $\mathbf{S} = [s_1, \dots, s_d] \in \text{St}(d, p)$, we can construct a rotation $\mathbf{Q} \in \pi^{-1}(\mathbf{S})$ simply by extending $\{s_1, \dots, s_d\}$ to an orthonormal basis $\{s_1, \dots, s_d, v_1, \dots, v_{p-d}\}$ for \mathbb{R}^p using the Gram-Schmidt process, and (if necessary) multiplying the final element v_{p-d} by -1 to ensure that this basis has a positive orientation; the matrix $\mathbf{Q} = [s_1, \dots, s_d, v_1, \dots, v_{p-d}] \in \text{SO}(p)$ whose columns are the elements of this basis is then a rotation satisfying $\pi(\mathbf{Q}) = \mathbf{S}$. Conversely, if $\mathbf{Q} \in \pi^{-1}(\mathbf{S})$, then by (22) $\mathbf{Q} = [\mathbf{S}, \mathbf{V}]$ for some

⁴ In the case that $d = p$, $\text{St}(p, p) \supset \text{SO}(p)$, and it is impossible for π to be surjective.

$V \in \mathbb{R}^{p \times (p-d)}$; writing the orthogonality constraint $Q^\top Q = I_p$ in terms of this block decomposition then produces:

$$Q^\top Q = \begin{bmatrix} S^\top S & S^\top V \\ V^\top S & V^\top V \end{bmatrix} = \begin{bmatrix} I_d & 0 \\ 0 & I_{p-d} \end{bmatrix}. \quad (23)$$

It follows from (23) that the preimage of any $S \in \text{St}(d, p)$ under the projection π in (22) is given explicitly by:

$$\pi^{-1}(S) = \{[S, V] \mid V \in \text{St}(p-d, p), S^\top V = 0, \det([S, V]) = +1\}. \quad (24)$$

Equations (22) and (24) provide a means of representing $\text{St}(d, p)$ using $\text{SO}(p)$: given any $S \in \text{St}(d, p)$, we can represent it using one of the rotations in (24) in which it appears as the first d columns, and conversely, given any $Q \in \text{SO}(p)$, we can extract its corresponding Stiefel manifold element using π . We can extend this relation to *products* of rotations and Stiefel manifolds in the natural way:

$$\begin{aligned} \Pi: \text{SO}(p)^n &\rightarrow \text{St}(d, p)^n \\ \Pi(Q_1, \dots, Q_n) &= (\pi(Q_1), \dots, \pi(Q_n)) \end{aligned} \quad (25)$$

and Π is likewise surjective for $p > d$.

The projection Π enables us to “pull back” the rank-restricted optimization (19) on $\text{St}(d, p)^n$ to an equivalent optimization problem on $\text{SO}(p)^n$. Concretely, if $f: \text{St}(d, p)^n \rightarrow \mathbb{R}$ is the objective of (19), then we simply define the objective \tilde{f} of our “lifted” optimization over $\text{SO}(p)^n$ to be the pullback of f via Π :

$$\begin{aligned} \tilde{f}: \text{SO}(p)^n &\rightarrow \mathbb{R} \\ \tilde{f}(Q) &= f \circ \Pi(Q). \end{aligned} \quad (26)$$

Comparing (19), (25), and (26) reveals that the pullback of the low-rank optimization (19) to $\text{SO}(p)^n$ is exactly the Shonan Averaging problem (2).

Theorem 4. *Let $p > d$. Then:*

- (i) *The rank-restricted optimization (19) and the Shonan Averaging problem (2) attain the same optimal value.*
- (ii) *$Q^* \in \text{SO}(p)^n$ minimizes (2) if and only if $S^* = \Pi(Q^*)$ minimizes (19).*

Similarly, we also have the following analogue of Theorem 3 for problem (2):

Theorem 5. *Let $Q^* = (Q_1^*, \dots, Q_n^*) \in \text{SO}(p)^n$ be a first-order critical point of problem (2), and $S^* = (S_1^*, \dots, S_n^*) = \Pi(Q^*) \in \text{St}(d, p)^n$. Then:*

- (i) *S^* is a first-order critical point of the rank-restricted optimization (19).*
- (ii) *Let C be the certificate matrix defined in (21), λ_{\min} its minimum eigenvalue, and $v = (v_1, \dots, v_n) \in \mathbb{R}^{dn}$ a corresponding eigenvector. If $\lambda_{\min} < 0$, then the point $Q^+ \in \text{SO}(p+1)^n$ defined elementwise by:*

$$Q_i^+ = \begin{bmatrix} Q_i^* & 0 \\ 0 & 1 \end{bmatrix} \in \text{SO}(p+1) \quad (27)$$

Algorithm 2 Shonan Rotation Averaging**Input:** An initial point $\mathbf{Q} \in \text{SO}(p)^n$ for (2), $p \geq d$.**Output:** A feasible estimate $\hat{\mathbf{R}} \in \text{SO}(d)^n$ for the rotation averaging problem (1), and the lower bound f_{SDP}^* for problem (1)'s optimal value.

```

1: function SHONAN_AVERAGING( $\mathbf{Q}$ )
2:   repeat                                     ▷ Riemannian Staircase
3:      $\mathbf{Q} \leftarrow \text{LOCALOPT}(\mathbf{Q})$                ▷ Find critical point of (2)
4:     Set  $\mathbf{S} \leftarrow \Pi(\mathbf{Q})$                  ▷ Project to  $\text{St}(d, p)^n$ 
5:     Construct the certificate matrix  $\mathbf{C}$  in (21), and
       compute its minimum eigenpair  $(\lambda_{\min}, v_{\min})$ .
6:     if  $\lambda_{\min} < 0$  then                       ▷  $\mathbf{Z} = \mathbf{S}^\top \mathbf{S}$  is not optimal
7:       Set  $p \leftarrow p + 1$  and  $\mathbf{Q}_i \leftarrow \begin{pmatrix} \mathbf{Q}_i & 0 \\ 0 & 1 \end{pmatrix} \forall i$ .      ▷ Ascend to next level
8:       Construct descent direction  $\dot{\mathbf{Q}}$  as in (28).
9:        $\mathbf{Q} \leftarrow \text{LINESEARCH}(\mathbf{Q}, \dot{\mathbf{Q}})$            ▷ Nudge  $\mathbf{Q}$ 
10:    end if
11:  until  $\lambda_{\min} \geq 0$                            ▷  $\mathbf{Z} = \mathbf{S}^\top \mathbf{S}$  solves (16)
12:  Set  $f_{\text{SDP}}^* \leftarrow \text{tr}(\bar{\mathbf{L}} \mathbf{S}^\top \mathbf{S})$ .
13:  Set  $\hat{\mathbf{R}} \leftarrow \text{ROUND\_SOLUTION}(\mathbf{S})$ .
14:  return  $\{\hat{\mathbf{R}}, f_{\text{SDP}}^*\}$ 
15: end function

```

is a first-order critical point of problem (2) in dimension $p + 1$, and the tangent vector $\dot{\mathbf{Q}}^+ \in T_{\mathbf{Q}^+}(\text{SO}(p + 1)^n)$ defined blockwise by:

$$\dot{\mathbf{Q}}_i^+ = \mathbf{Q}_i^+ \begin{bmatrix} 0 & -v_i \\ v_i^\top & 0 \end{bmatrix} \in T_{\mathbf{Q}_i^+}(\text{SO}(p + 1)) \quad (28)$$

is a second-order descent direction from \mathbf{Q}^+ .

Theorems 4 and 5 are proved in Appendix A of the supplementary material.

4.5 The complete algorithm

Combining the results of the previous sections gives the complete *Shonan Averaging* algorithm (Algorithm 2). (We employ Levenberg-Marquardt to perform the fast local optimization required in line 3 – see Appendix B for details).

When applied to an instance of rotation averaging, Shonan Averaging returns a feasible point $\hat{\mathbf{R}} \in \text{SO}(d)^n$ for the maximum likelihood estimation (1), and a lower bound $f_{\text{SDP}}^* \leq f_{\text{MLE}}^*$ for problem (1)'s optimal value. This lower bound provides an *upper* bound on the suboptimality of *any* feasible point $\mathbf{R} \in \text{SO}(d)^n$ as a solution of problem (1) according to:

$$f(\bar{\mathbf{L}} \mathbf{R}^\top \mathbf{R}) - f_{\text{SDP}}^* \geq f(\bar{\mathbf{L}} \mathbf{R}^\top \mathbf{R}) - f_{\text{MLE}}^*. \quad (29)$$

Furthermore, if the relaxation (16) is exact, the estimate $\hat{\mathbf{R}}$ returned by Algorithm 2 *attains* this lower bound:

$$f(\bar{\mathbf{L}} \hat{\mathbf{R}}^\top \hat{\mathbf{R}}) = f_{\text{SDP}}^*. \quad (30)$$

	n	m	p	Shonan Averaging			
				λ_{\min}	f_{SDP}^*	Opt. time [s]	Min. eig. time [s]
smallGrid	125	297	5	-9.048×10^{-6}	4.850×10^2	2.561×10^{-1}	2.613×10^{-1}
sphere	2500	4949	6	-1.679×10^{-4}	5.024×10^0	1.478×10^1	1.593×10^1
torus	5000	9048	5	-2.560×10^{-4}	1.219×10^4	7.819×10^1	8.198×10^1
garage	1661	6275	5	1.176×10^{-4}	2.043×10^{-1}	1.351×10^1	1.420×10^1
cubicle	5750	16869	5	-9.588×10^{-5}	3.129×10^1	1.148×10^2	1.208×10^2

Table 1. Shonan Averaging results for the SLAM benchmark datasets

Consequently, verifying *a posteriori* that (30) holds provides a *certificate* of $\hat{\mathbf{R}}$'s *global* optimality as a solution of the rotation averaging problem (1).

5 Experimental Results

In this section we evaluate Shonan Averaging's performance on (a) several large-scale problems derived from standard pose-graph SLAM benchmarks, (b) small randomly generated synthetic datasets, and (c) several structure from motion (SFM) datasets used in earlier work. The experiments were performed on a desktop computer with a 6-core Intel i5-9600K CPU @ 3.70GHz running Ubuntu 18.04. Our implementation of Shonan Averaging (SA) is written in C++, using the GTSAM library [12] to perform the local optimization. The fast minimum-eigenvalue computation described in [26, Sec. III-C] is implemented using the symmetric Lanczos algorithm from the Spectra library⁵ to compute the minimum eigenvalue required in line 5. We initialize the algorithm with a randomly-sampled point $\mathbf{Q} \in \text{SO}(p)^n$ at level $p = 5$, and require that the minimum eigenvalue satisfy $\lambda_{\min} \geq -10^{-4}$ at termination of the Riemannian Staircase.

The results for the pose-graph SLAM benchmarks demonstrate that Shonan Averaging is already capable of recovering *certifiably globally optimal* solutions of large-scale real-world rotation averaging problems in tractable time. For these experiments, we have extracted the rotation averaging problem (1) obtained by simply ignoring the translational parts of the measurements. Table 1 reports the size of each problem (the number of unknown rotations n and relative measurements m), together with the relaxation rank p at the terminal level of the Riemannian Staircase, the optimal value of the semidefinite relaxation (16) corresponding to the recovered low-rank factor \mathbf{S}^* , and the total elapsed computation times for the local optimizations in line 3 and minimum-eigenvalue computations in line 5 of Algorithm 2, respectively. We remark that the datasets torus in Table 1 is *1-2 orders of magnitude larger* than the examples reported in previous work on direct globally-optimal rotation averaging algorithms [13,16].

We attribute Shonan Averaging's improved scalability to its use of super-linear *local* optimization in conjunction with the Riemannian Staircase; in effect, this strategy provides a means of "upgrading" a fast *local* method to a fast *global*

⁵ <https://Spectralib.org/>

one. This enables Shonan Averaging to leverage existing heavily-optimized, scalable, and high-performance software libraries for superlinear *local* optimization [2,22,12] as the main computational engine of the algorithm (cf. line 3 of Algorithm 2), while preserving the guarantee of *global* optimality.

n, σ	method	error	time (s)	success
n=20 $\sigma=0.2$	SA	0.000%	0.030	100%
	BD	1.397%	0.374	60%
	LM	-0.001%	0.003	80%
n=20 $\sigma=0.5$	SA	0.000%	0.037	100%
	BD	0.014%	0.446	40%
	LM	-0.001%	0.004	20%
n=50 $\sigma=0.2$	SA	0.000%	0.141	100%
	BD	0.291%	6.261	60%
	LM	-0.039%	0.011	20%
n=50 $\sigma=0.5$	SA	0.000%	0.194	100%
	BD	-7.472%	6.823	60%
	LM	-0.008%	0.014	40%
n=100 $\sigma=0.2$	SA	0.000%	0.403	100%
	BD	nan%	nan	0%
	LM	-0.014%	0.016	20%
n=100 $\sigma=0.5$	SA	0.000%	0.373	100%
	BD	3.912%	32.982	20%
	LM	-0.142%	0.023	60%
n=200 $\sigma=0.2$	SA	0.000%	1.275	100%
	BD	nan%	nan	0%
	LM	-0.102%	0.049	60%
n=200 $\sigma=0.5$	SA	0.000%	1.845	100%
	BD	nan%	nan	0%
	LM	-0.293%	0.046	40%

Table 2. Synthetic results for varying problem sizes n and noise levels σ . SA=Shonan, BD=block coordinate descent from [13], LM=Levenberg-Marquardt.

The synthetic results confirm that Shonan Averaging always converges to the true (*global*) minimizer, and significantly outperforms the block coordinate descent method from [13,14] as soon as the number of unknown rotations becomes large. We followed the approach for generating the data from [13], generating 4 sets of problem instances of increasing size, for two different noise levels. We generated 3D poses on a circular trajectory, forming a cycle graph. Relative rotation measurements were corrupted by composing with a random axis-angle perturbation, where the axis was chosen randomly and the angle was generated from a normal distribution with a standard deviation $\sigma = 0.2$ or $\sigma = 0.5$. Initialization for all three algorithms was done randomly, with angles uniform random over the range $(-\pi, \pi)$. We compare our results with two baselines: (LM) Levenberg-Marquardt on the $SO(3)^n$ manifold, also implemented in GT-

dataset	method	error	time	success
reichstag (n=71, m=2554)	SA	0.000%	0.197	100%
	BD	0.000%	0.247	100%
	LM	0.001%	0.085	80%
pantheon_interior (n=186, m=10000)	SA	0.000%	0.971	100%
	BD	0.000%	1.823	100%
	LM	0.001%	0.313	100%

Table 3. Results on SFM problems from the YFCC dataset.

SAM, which is fast but not expected to find the global minimizer; (BD): the block coordinate descent method from [13,14]. Because no implementation was available, we re-implemented it in Python, using a 3×3 SVD decomposition at the core. However, our implementation differs in that we establish optimality using the eigenvalue certificate from Shonan.

The synthetic results are shown in Table 2. In the table we show the percentage of the cases in which either method found a global minimizer, either as certified by the minimum eigenvalue $\lambda_{\min} \geq -10^{-4}$, or being within 5% of the optimal cost. Shonan Averaging (“SA” in the table) finds an optimal solution every time as certified by λ_{\min} . Levenberg-Marquardt is fast but finds global minima in only about 40% to 60% of the cases. Finally, the block coordinate descent method (BD) is substantially slower, and does not always converge within the allotted time. We found that in practice BD converges very slowly and it takes a long time for the algorithm to converge to the same value as found by SA or LM. In the table we have shown the relative error and running time (in seconds), both averaged over 5 runs. For $n = 100$ and $n = 200$ the BD method did often not converge to the global optimum within reasonable time.

Finally, Table 3 shows results on two datasets derived by Heinly et al. [21] from the large-scale YFCC-100M dataset [33,21]. The relative measurements for these were derived from the SFM solution provided with the data, and corrupted with noise as before, using $\sigma = 0.2$. All three methods agree on the globally optimal solution, and the timing data shows the same trend as in Table 2.

6 Conclusion

In this work we presented Shonan Rotation Averaging, a fast, simple, and elegant algorithm for rotation averaging that is *guaranteed* to recover globally optimal solutions under mild conditions on the measurement noise. Our approach applies a fast *local* search technique to a *sequence* of higher-dimensional lifts of the rotation averaging problem until a globally optimal solution is found. Shonan Averaging thus leverages the speed and scalability of existing high-performance *local* optimization methods already in common use, while enabling the recovery of *provably optimal* solutions.

References

1. P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, USA, 2007. [3](#), [7](#), [18](#), [19](#)
2. Sameer Agarwal and Keir Mierle. *Ceres Solver: Tutorial & Reference*. Google Inc., 2012. [2](#), [3](#), [9](#), [13](#)
3. S. Agarwal, N. Snavely, I. Simon, S. M. Seitz, and R. Szeliski. Bundle adjustment in the large. In *European Conf. on Computer Vision (ECCV)*, pages 29–42, 2010. [1](#)
4. N. Boumal. A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints. *arXiv preprint arXiv:1506.00575*, 2015. [8](#)
5. N. Boumal, A. Singer, and P.A. Absil. Cramer-rao bounds for synchronization of rotations. *Information and Inference*, 3(1):1–39, 2014. [2](#), [26](#)
6. N. Boumal, A. Singer, and P.-A. Absil. Robust estimation of rotations from relative measurements by maximum likelihood. In *52nd IEEE Conference on Decision and Control*, pages 1156–1161. IEEE, 2013. [2](#)
7. N. Boumal, V. Voroninski, and A. Bandeira. The non-convex Burer–Monteiro approach works on smooth semidefinite programs. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2757–2765. 2016. [7](#)
8. S. Burer and R. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003. [6](#)
9. L. Carlone, D.M. Rosen, G.C. Calafiore, J.J. Leonard, and F. Dellaert. Lagrangian duality in 3D SLAM: Verification techniques and optimal solutions. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 125–132, 2015. ([pdf](#)) ([code](#)) (datasets: ([web](#))) (supplemental material: ([pdf](#))). [2](#), [3](#)
10. L. Carlone, R. Tron, K. Daniilidis, and F. Dellaert. Initialization techniques for 3D SLAM: a survey on rotation estimation and its use in pose graph optimization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 4597–4604, 2015. ([pdf](#)) ([code](#)) (supplemental material: ([pdf](#))). [2](#)
11. T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley & Sons, New York, NY, 1991. [1](#)
12. Frank Dellaert. Factor graphs and GTSAM: A hands-on introduction. Technical Report GT-RIM-CP&R-2012-002, Georgia Institute of Technology, September 2012. [2](#), [3](#), [9](#), [12](#), [13](#)
13. A. Eriksson, C. Olsson, F. Kahl, and T.-J. Chin. Rotation averaging and strong duality. *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. [3](#), [12](#), [13](#), [14](#), [24](#), [25](#)
14. Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging with the chordal distance: Global minimizers and strong duality. *IEEE Trans. Pattern Anal. Machine Intell.*, 2019. [3](#), [13](#), [14](#), [24](#), [25](#), [28](#)
15. J.-M. Frahm, P. Fite-Georgel, D. Gallup, T. Johnson, R. Raguram, C. Wu, Y.-H. Jen, E. Dunn, B. Clipp, S. Lazebnik, et al. Building Rome on a cloudless day. In *European Conference on Computer Vision*, pages 368–381. Springer, 2010. [2](#)
16. J. Fredriksson and C. Olsson. Simultaneous multiple rotation averaging using lagrangian duality. In *Asian Conf. on Computer Vision (ACCV)*, 2012. [3](#), [12](#)
17. V.M. Govindu. Combining two-view constraints for motion estimation. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 2. IEEE, 2001. [1](#)
18. R. Hartley, J. Trumpf, Y. Dai, and H. Li. Rotation averaging. *IJCV*, 103(3):267–305, 2013. [1](#), [2](#), [3](#)

19. R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000. [1](#)
20. J. Heinly, J.L. Schonberger, E. Dunn, and J.-M. Frahm. Reconstructing the world* in six days*(as captured by the Yahoo 100 million image dataset). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3287–3295, 2015. [2](#)
21. Jared Heinly, Johannes L. Schonberger, Enrique Dunn, and Jan-Michael Frahm. Reconstructing the world* in six days *(as captured by the Yahoo 100 million image dataset). In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. [14](#), [24](#)
22. R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g2o: A general framework for graph optimization. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, May 2011. [2](#), [3](#), [9](#), [13](#)
23. D. Martinec and T. Pajdla. Robust rotation and translation estimation in multi-view reconstruction. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2007. [2](#)
24. Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer Series in Operations Research. Springer-Verlag, 1999. [2](#), [3](#), [8](#)
25. R. Pless. Using many cameras as one. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 587–593, 2003. [1](#)
26. D.M. Rosen and L. Carlone. Computational enhancements for certifiably correct SLAM. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2017. Workshop on “Introspective Methods for Reliable Autonomy”, ([pdf](#)). [2](#), [12](#)
27. D.M. Rosen, L. Carlone, A.S. Bandeira, and J.J. Leonard. SE-Sync: A certifiably correct algorithm for synchronization over the Special Euclidean group. In *Intl. Workshop on the Algorithmic Foundations of Robotics (WAFR)*, San Francisco, CA, December 2016. extended arxiv preprint: 1611.00128, ([pdf](#)) ([pdf](#)) ([code](#)). [2](#), [3](#), [5](#), [6](#), [9](#), [26](#)
28. D.M. Rosen, L. Carlone, A.S. Bandeira, and J.J. Leonard. SE-Sync: a certifiably correct algorithm for synchronization over the Special Euclidean group. *Intl. J. of Robotics Research*, 2018. accepted, arxiv preprint: 1611.00128, ([pdf](#)). [8](#)
29. J.L. Schonberger and J.-M. Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. [1](#)
30. A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.*, 30:20–36, 2010. [2](#)
31. A. Singer and H.-T. Wu. Vector diffusion maps and the connection Laplacian. *Comm. Pure Appl. Math.*, 65:1067–1144, 2012. [5](#)
32. J.F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization methods and software*, 11(1-4):625–653, 1999. [3](#)
33. Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. YFCC100M: the new data in multimedia research. *Commun. ACM*, 59(2):64–73, 2016. [14](#), [24](#)
34. B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. Bundle adjustment – a modern synthesis. In W. Triggs, A. Zisserman, and R. Szeliski, editors, *Vision Algorithms: Theory and Practice*, volume 1883 of *LNCS*, pages 298–372. Springer Verlag, 2000. [1](#), [23](#), [24](#)
35. R. Tron, R. Vidal, and A. Terzis. Distributed pose averaging in camera networks via consensus on SE(3). In *2008 Second ACM/IEEE International Conference on Distributed Smart Cameras*, pages 1–10. IEEE, 2008. [1](#)
36. S. Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Anal. Machine Intell.*, 13(4):376–380, April 1991. [9](#)

- 37. L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Rev.*, 38(1):49–95, 1996. [3](#), [6](#), [8](#)
- 38. K. Wilson, D. Bindel, and N. Snavely. When is rotations averaging hard? In *European Conference on Computer Vision*, pages 255–270. Springer, 2016. [3](#), [24](#), [26](#)

Supplementary material

A Proof of Theorems 4 and 5

A.1 Proof of Theorem 4

Proof. For part (i), let $\mathbf{Q}^* \in \text{SO}(p)^n$ be a minimizer of (2), with corresponding optimal value \tilde{f}^* , and suppose for contradiction that $\mathbf{S}^* = \Pi(\mathbf{Q}^*)$ is *not* a minimizer of (19). Then there exists some $\mathbf{S}' \in \text{St}(d, p)^n$ with $f(\mathbf{S}') < f(\mathbf{S}^*) = \tilde{f}^*$. However, since Π is surjective for $p > d$, then there exists some $\mathbf{Q}' \in \Pi^{-1}(\mathbf{S}')$, and by definition of \tilde{f} in (26) we have that $\tilde{f}(\mathbf{Q}') = f(\mathbf{S}') < \tilde{f}^*$, contradicting the optimality of \mathbf{Q}^* . We conclude that \mathbf{S}^* is indeed a minimizer of f over $\text{St}(d, p)^n$, and therefore the optimal values of (19) and (2) coincide.

Part (ii) follows immediately from part (i) since $\tilde{f}(\mathbf{Q}) = f(\mathbf{S})$ for any $\mathbf{Q} \in \text{SO}(p)^n$ and $\mathbf{S} = \Pi(\mathbf{Q})$ by (26).

A.2 Proof of Theorem 5

In this section we derive Theorem 5 as a consequence of Theorem 3 and the equivalence of the rank-restricted optimization (19) and the Shonan Averaging problem (2) (Theorem 4). To do so, we need to understand how the local geometry of the lifted objective \tilde{f} relates to that of f near a critical point $\mathbf{Q}^* \in \text{SO}(p)^n$. Recall that the tangent spaces of the rotational and Stiefel manifolds can be expressed as:

$$T_{\mathbf{Q}}(\text{SO}(p)) = \left\{ \mathbf{Q}\dot{\Delta} \mid \dot{\Delta} \in \text{Skew}(p) \right\} \quad (31)$$

and

$$T_{\mathbf{S}}(\text{St}(d, p)) = \left\{ \mathbf{S}\dot{\Omega} + \mathbf{V}\dot{\mathbf{K}} \mid \begin{array}{l} \dot{\Omega} \in \text{Skew}(d), \\ \dot{\mathbf{K}} \in \mathbb{R}^{(p-d) \times d}, \\ \mathbf{V} \in \mathbb{R}^{p \times (p-d)}, \\ \mathbf{S}^\top \mathbf{V} = 0 \end{array} \right\} \quad (32)$$

respectively (cf. Example 3.5.2 of [1]). If we rewrite the elements of (31) in a block-partitioned form compatible with the action of the projection (22):

$$T_{\mathbf{Q}}(\text{SO}(p)) = \left\{ \mathbf{Q} \begin{bmatrix} \dot{\Omega} - \dot{\mathbf{K}}^\top \\ \dot{\mathbf{K}} \\ \dot{\mathbf{I}} \end{bmatrix} \mid \begin{array}{l} \dot{\Omega} \in \text{Skew}(d), \\ \dot{\mathbf{K}} \in \mathbb{R}^{(p-d) \times d}, \\ \dot{\mathbf{I}} \in \text{Skew}(p-d) \end{array} \right\} \quad (33)$$

then the derivative of the projection π at $\mathbf{Q} = [\mathbf{S} \mathbf{V}] \in \text{SO}(p)$ is:

$$\begin{aligned} d\pi_{\mathbf{Q}}: T_{\mathbf{Q}}(\text{SO}(p)) &\rightarrow T_{\mathbf{S}}(\text{St}(d, p)) \\ d\pi_{\mathbf{Q}} \left(\mathbf{Q} \begin{bmatrix} \dot{\Omega} - \dot{\mathbf{K}}^\top \\ \dot{\mathbf{K}} \\ \dot{\mathbf{I}} \end{bmatrix} \right) &= \mathbf{S}\dot{\Omega} + \mathbf{V}\dot{\mathbf{K}}. \end{aligned} \quad (34)$$

The next result is a direct consequence of (31)–(34):

Lemma 1. *The projection $\pi: \text{SO}(p) \rightarrow \text{St}(d, p)$ is a submersion.*⁶

Proof. Let $\mathbf{Q} = [\mathbf{S} \ \mathbf{V}] \in \text{SO}(p)$, $\mathbf{S} = \pi(\mathbf{Q})$, and $\dot{\mathbf{Y}} \in T_{\mathbf{S}}(\text{St}(d, p))$; we must show that there is some $\dot{\mathbf{X}} \in T_{\mathbf{Q}}(\text{SO}(p))$ such that $\dot{\mathbf{Y}} = d\pi_{\mathbf{Q}}(\dot{\mathbf{X}})$. Since $\mathbf{S}^\top \mathbf{V} = 0$ (as the columns of \mathbf{Q} are orthonormal), by (32) there exist some $\dot{\mathbf{\Omega}} \in \text{Skew}(d)$ and $\dot{\mathbf{K}} \in \mathbb{R}^{(p-d) \times d}$ such that $\dot{\mathbf{Y}} = \mathbf{S} \dot{\mathbf{\Omega}} + \mathbf{V} \dot{\mathbf{K}}$. Letting $\dot{\mathbf{I}} \in \text{Skew}(p-d)$, and defining

$$\dot{\mathbf{X}} = \mathbf{Q} \begin{bmatrix} \dot{\mathbf{\Omega}} & -\dot{\mathbf{K}}^\top \\ \dot{\mathbf{K}} & \dot{\mathbf{I}} \end{bmatrix} \in T_{\mathbf{Q}}(\text{SO}(p)) \quad (35)$$

as in (33), it follows from (34) that $\dot{\mathbf{Y}} = d\pi_{\mathbf{Q}}(\dot{\mathbf{X}})$.

Corollary 1. *The projection $\Pi: \text{SO}(p)^n \rightarrow \text{St}(d, p)^n$ is a submersion.*

Proof of Theorem 5:

Part(i): Applying the Chain Rule to (26) produces:

$$d\tilde{f}_{\mathbf{Q}^*} = df_{\mathbf{S}^*} \circ d\Pi_{\mathbf{Q}^*}. \quad (36)$$

The left-hand side of (36) is 0 because \mathbf{Q}^* is a first-order critical point of \tilde{f} by hypothesis, and Corollary 1 shows that $d\Pi_{\mathbf{Q}^*}$ is full-rank (i.e., its image is all of $T_{\mathbf{S}^*}(\text{St}(d, p)^n)$). It follows that $df_{\mathbf{S}^*} = 0$, and therefore \mathbf{S}^* is a first-order critical point of (19).

Part (ii): Observe that $\Pi(\mathbf{Q}^+) = \mathbf{S}^+$ as defined in Theorem 3(ii); since \mathbf{S}^+ is a stationary point for (19), another application of the Chain Rule as in (36) shows that $d\tilde{f}_{\mathbf{Q}^+} = 0$, and therefore \mathbf{Q}^+ is also a first-order critical point of the lifted optimization (2) in dimension $p+1$. Theorem 3(ii) also provides the second-order descent direction $\dot{\mathbf{S}}^+$ from \mathbf{S}^+ . The tangent vector $\dot{\mathbf{Q}}^+$ defined in (28) is constructed as a *lift* of $\dot{\mathbf{S}}^+$ to $T_{\mathbf{Q}^+}(\text{SO}(p+1)^n)$; that is, so that $\dot{\mathbf{Q}}^+$ satisfies $d\Pi_{\mathbf{Q}^+}(\dot{\mathbf{Q}}^+) = \dot{\mathbf{S}}^+$. Indeed, using (27), (28), and (34) we can compute each block of $d\Pi_{\mathbf{Q}^+}(\dot{\mathbf{Q}}^+)$ as:

$$\begin{aligned} d\Pi_{\mathbf{Q}^+}(\dot{\mathbf{Q}}^+)_i &= d\pi_{\mathbf{Q}_i^+} \left(\begin{bmatrix} \mathbf{Q}_i^* & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -v_i \\ v_i^\top & 0 \end{bmatrix} \right) \\ &= \begin{bmatrix} 0 \\ v_i^\top \end{bmatrix} = \dot{\mathbf{S}}_i^+ \end{aligned} \quad (37)$$

for all $i \in [n]$, verifying that $\dot{\mathbf{Q}}^+$ is a lift of $\dot{\mathbf{S}}^+$.

Our goal now is to show that $\dot{\mathbf{Q}}^+$ is likewise a second-order descent direction from \mathbf{Q}^+ . We do this using a proof technique similar to that of Proposition 5.5.6 in [1]. Let $R_{\mathbf{Q}^+}: T_{\mathbf{Q}^+}(\text{SO}(p+1)^n) \rightarrow \text{SO}(p+1)^n$ be any (first-order) retraction on the tangent space of $\text{SO}(p+1)^n$ at \mathbf{Q}^+ , and let $\epsilon > 0$ be sufficiently small that the curve:

$$\begin{aligned} \tilde{\gamma}: (-\epsilon, \epsilon) &\rightarrow \text{SO}(p+1)^n \\ \tilde{\gamma}(t) &= R_{\mathbf{Q}^+}(t\dot{\mathbf{Q}}^+) \end{aligned} \quad (38)$$

⁶ A smooth mapping $\varphi: \mathcal{X} \rightarrow \mathcal{Y}$ between manifolds is called a *submersion* at a point $x \in \mathcal{X}$ if its derivative $d\varphi_x: T_x(\mathcal{X}) \rightarrow T_{\varphi(x)}(\mathcal{Y})$ is surjective. It is called a *submersion* (unqualified) if it is a submersion at *every* $x \in \mathcal{X}$.

obtained by moving through the point $\mathbf{Q}^+ \in \text{SO}(p+1)^n$ along the direction $\dot{\mathbf{Q}}^+$ is well-defined. Our approach simply involves examining the behavior of the lifted objective \tilde{f} at points along the curve $\tilde{\gamma}$ in a neighborhood of $\tilde{\gamma}(0) = \mathbf{Q}^+$. To do so, we define:

$$\begin{aligned}\tilde{f}: (-\epsilon, \epsilon) &\rightarrow \mathbb{R} \\ \tilde{f}(t) &= \tilde{f} \circ \tilde{\gamma}(t)\end{aligned}\tag{39}$$

and then consider its first- and second-order derivatives. Once again using the Chain Rule (and unwinding the definition of \tilde{f}), we compute:

$$\begin{aligned}\tilde{f}'(t) &= d\tilde{f}_{\tilde{\gamma}(t)} \circ \tilde{\gamma}'(t) \\ &= df_{\Pi \circ \tilde{\gamma}(t)} \circ d\Pi_{\tilde{\gamma}(t)} \circ \tilde{\gamma}'(t) \\ &= \langle \text{grad } f(\Pi \circ \tilde{\gamma}(t)), d\Pi_{\tilde{\gamma}(t)} \circ \tilde{\gamma}'(t) \rangle.\end{aligned}\tag{40}$$

Note that at $t = 0$, $\text{grad } f(\Pi \circ \tilde{\gamma}(0)) = \text{grad } f(\Pi(\mathbf{Q}^+)) = \text{grad } f(\mathbf{S}^+) = 0$ since \mathbf{S}^+ is a stationary point. It follows from (40) that $\tilde{f}'(0) = 0$ (as expected, since we know that \mathbf{Q}^+ is a stationary point for the lifted optimization (2)). Continuing, we compute the second derivative of $\tilde{f}(t)$ by applying the Product and Chain Rules to differentiate the inner product on the final line of (40):

$$\begin{aligned}\tilde{f}''(t) &= \frac{d}{dt} [\langle \text{grad } f(\Pi \circ \tilde{\gamma}(t)), d\Pi_{\tilde{\gamma}(t)} \circ \tilde{\gamma}'(t) \rangle] \\ &= \left\langle \frac{d}{dt} [\text{grad } f(\Pi \circ \tilde{\gamma}(t))], d\Pi_{\tilde{\gamma}(t)} \circ \tilde{\gamma}'(t) \right\rangle \\ &\quad + \left\langle \text{grad } f(\Pi \circ \tilde{\gamma}(t)), \frac{d}{dt} [d\Pi_{\tilde{\gamma}(t)} \circ \tilde{\gamma}'(t)] \right\rangle.\end{aligned}\tag{41}$$

Now, we just saw that $\text{grad } f(\Pi \circ \tilde{\gamma}(0)) = 0$, so the second term of the final line of (41) is zero for $t = 0$. Moreover, the derivative in the first term can be further developed as:

$$\frac{d}{dt} [\text{grad } f(\Pi \circ \tilde{\gamma}(t))] = \text{Hess } f(\Pi \circ \tilde{\gamma}(t)) [d\Pi_{\tilde{\gamma}(t)} \circ \tilde{\gamma}'(t)].\tag{42}$$

Therefore, at $t = 0$ (41) simplifies as:

$$\begin{aligned}\tilde{f}''(0) &= \langle \text{Hess } f(\Pi \circ \tilde{\gamma}(0)) [d\Pi_{\tilde{\gamma}(0)} \circ \tilde{\gamma}'(0)], d\Pi_{\tilde{\gamma}(0)} \circ \tilde{\gamma}'(0) \rangle \\ &= \langle \text{Hess } f(\Pi(\mathbf{Q}^+)) [d\Pi_{\mathbf{Q}^+}(\dot{\mathbf{Q}}^+)], d\Pi_{\mathbf{Q}^+}(\dot{\mathbf{Q}}^+) \rangle \\ &= \langle \text{Hess } f(\mathbf{S}^+) [\dot{\mathbf{S}}^+], \dot{\mathbf{S}}^+ \rangle \\ &< 0\end{aligned}\tag{43}$$

where the final line of (43) follows from the fact that $\dot{\mathbf{S}}^+$ is a second-order direction of descent from \mathbf{S}^+ . We conclude from (43) that $\dot{\mathbf{Q}}^+$ is a second-order direction of descent from \mathbf{Q}^+ , as desired. \square

B Gauss-Newton for Shonan Averaging

We can implement the local search for (first-order) critical points of (2) required in line 3 of the Shonan Averaging algorithm using the same Gauss-Newton approach described in Section 3.

B.1 Linearization

As before, we first rewrite (2) more explicitly as the minimization of the sum of the individual measurement residuals, in a vectorized form:

$$\min_{\mathbf{Q} \in \text{SO}(p)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \left\| \text{vec}(\mathbf{Q}_j \mathbf{P} - \mathbf{Q}_i \mathbf{P} \bar{\mathbf{R}}_{ij}) \right\|_2^2, \quad (44)$$

and reparameterize this minimization in terms of the Lie algebra $\mathfrak{so}(p)$:

$$\min_{\boldsymbol{\delta} \in \mathfrak{so}(p)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \left\| \text{vec} \left(\mathbf{Q}_j e^{[\boldsymbol{\delta}_j]} \mathbf{P} - \mathbf{Q}_i e^{[\boldsymbol{\delta}_i]} \mathbf{P} \bar{\mathbf{R}}_{ij} \right) \right\|_2^2. \quad (45)$$

Once again, we approximate (45) to first order as the linear least-squares objective:

$$\min_{\boldsymbol{\delta} \in \mathfrak{so}(p)^n} \sum_{(i,j) \in \mathcal{E}} \kappa_{ij} \left\| \mathbf{F}_j \boldsymbol{\delta}_j - \mathbf{H}_i \boldsymbol{\delta}_i - \mathbf{b}_{ij} \right\|_2^2, \quad (46)$$

where now the Jacobians \mathbf{F}_j and \mathbf{H}_i and the right-hand side \mathbf{b}_{ij} are calculated as

$$\mathbf{F}_j \doteq (\mathbf{P}^\top \otimes \mathbf{Q}_j) \bar{\mathbf{G}}_p \quad (47)$$

$$\mathbf{H}_i \doteq ((\mathbf{P} \bar{\mathbf{R}}_{ij})^\top \otimes \mathbf{Q}_i) \bar{\mathbf{G}}_p \quad (48)$$

$$\mathbf{b}_{ij} \doteq \text{vec}(\mathbf{Q}_i \mathbf{P} \bar{\mathbf{R}}_{ij} - \mathbf{Q}_j \mathbf{P}) \quad (49)$$

with $\bar{\mathbf{G}}_p$ the matrix of vectorized $\mathfrak{so}(p)^n$ generators. Note that the right-hand side \mathbf{b}_{ij} in (49) in fact involves only the Stiefel manifold elements $\mathbf{S}_i = \pi(\mathbf{Q}_i)$:

$$\mathbf{b}_{ij} = \text{vec}(\mathbf{S}_i \bar{\mathbf{R}}_{ij} - \mathbf{S}_j). \quad (50)$$

B.2 The structure of the Lie algebra

In this section we investigate the structure of the Lie algebra $\mathfrak{so}(p)$ as it pertains to the linearized objective in (46). Recall that $\mathfrak{so}(p)$ is identified with the tangent space $T_{\mathbf{I}_p}(\text{SO}(p)) = \text{Skew}(p)$, the space of $p \times p$ skew-symmetric matrices. Once again writing these matrices in a block-partitioned form compatible with the projection π in (22) (as in (33)) produces:

$$\mathfrak{so}(p) = \left\{ \begin{bmatrix} [\boldsymbol{\omega}] & -\mathbf{K}^\top \\ \mathbf{K} & [\boldsymbol{\gamma}] \end{bmatrix} \mid \begin{array}{l} \boldsymbol{\omega} \in \mathfrak{so}(d), \\ \mathbf{K} \in \mathbb{R}^{(p-d) \times d}, \\ \boldsymbol{\gamma} \in \mathfrak{so}(p-d) \end{array} \right\}, \quad (51)$$

and it follows from (22) the derivative of π at \mathbf{I}_p is:

$$d\pi_{\mathbf{I}_p} \left(\begin{bmatrix} [\boldsymbol{\omega}] - \mathbf{K}^\top \\ \boldsymbol{\gamma} \end{bmatrix} \right) = \begin{bmatrix} [\boldsymbol{\omega}] \\ \mathbf{K} \end{bmatrix}. \quad (52)$$

Note that $d\pi_{\mathbf{I}_p}$ does *not* depend upon the $(p-d)$ -dimensional vector $\boldsymbol{\gamma}$. In particular, let us define:

$$V_{\mathbf{I}_p} \doteq \ker d\pi_{\mathbf{I}_p} = \left\{ \begin{pmatrix} 0 & 0 \\ 0 & [\boldsymbol{\gamma}] \end{pmatrix} \mid \boldsymbol{\gamma} \in \mathfrak{so}(p-d) \right\} \subset \mathfrak{so}(p). \quad (53)$$

Geometrically, the set $V_{\mathbf{I}_p}$ defined in (53) consists of those directions of motion $\dot{\boldsymbol{\Omega}} \in \mathfrak{so}(p)$ at \mathbf{I}_p along which the projection π is *constant*. Equivalently:

$$V_{\mathbf{I}_p} = T_{\mathbf{I}_p}(\pi^{-1}(\mathbf{P})); \quad (54)$$

that is, $V_{\mathbf{I}_p}$ is the set of vectors that are *tangent* to the preimage $\pi^{-1}(\mathbf{P})$ of $\mathbf{P} = \pi(\mathbf{I}_p)$ at \mathbf{I}_p . If we think of the elements of the preimage $\pi^{-1}(\mathbf{P})$ as being vertically “stacked” above their common projection $\mathbf{P} \in \text{St}(d, p)$, then the subspace $V_{\mathbf{I}_p}$ of the Lie algebra is precisely the set of tangent vectors at \mathbf{I}_p that correspond to *vertical motions*. Consequently, $V_{\mathbf{I}_p}$ is referred to as the *vertical space*. We may define a corresponding *horizontal space* in the natural way, i.e., as the orthogonal complement of the vertical space:

$$H_{\mathbf{I}_p} \doteq V_{\mathbf{I}_p}^\perp = \left\{ \begin{pmatrix} [\boldsymbol{\omega}] - \mathbf{K}^\top \\ \mathbf{K} \end{pmatrix} \mid \begin{array}{l} \boldsymbol{\omega} \in \mathfrak{so}(d), \\ \mathbf{K} \in \mathbb{R}^{(p-d) \times d} \end{array} \right\} \subset \mathfrak{so}(p). \quad (55)$$

The significance of (51)–(55) is that, to first order, the exponential map (or any retraction) can be written as $\mathbf{I} + [\boldsymbol{\delta}]$ for $\boldsymbol{\delta} \in \mathfrak{so}(p)$. In conjunction with the projection map π (equivalently, with \mathbf{P}), this implies:

$$\mathbf{Q}_i e^{[\boldsymbol{\delta}]} \mathbf{P} \approx \mathbf{Q}_i (\mathbf{I} + [\boldsymbol{\delta}]) \mathbf{P} = \mathbf{S}_i + \mathbf{Q}_i \begin{bmatrix} [\boldsymbol{\omega}] \\ \mathbf{K} \end{bmatrix}. \quad (56)$$

From this we see that the derivative of the cost function (45) will not depend on the $(p-d)$ -dimensional vector $\boldsymbol{\gamma}$, i.e., on the component of $[\boldsymbol{\delta}]$ lying in the vertical space $V_{\mathbf{I}_p}$. This makes intuitive sense, since the Shonan Averaging objective $\tilde{f}(\mathbf{Q})$ in (2) is defined in terms of the projection $\Pi(\mathbf{Q})$ (recall (26)), and moving along vertical directions leaves this projection unchanged.

This in turn enables us to characterize the Jacobians in more detail. If we split $\mathbf{Q}_i = [\mathbf{S}_i \ \mathbf{V}_i]$, the Jacobians \mathbf{F}_j and \mathbf{H}_i can be shown to be:

$$\mathbf{F}_j \doteq [(\mathbf{I}_d \otimes \mathbf{S}_j) \bar{\mathbf{G}}_d (\mathbf{I}_d \otimes \mathbf{V}_j) \ 0] \quad (57)$$

$$\mathbf{H}_i \doteq [(\bar{\mathbf{R}}_{ij}^\top \otimes \mathbf{S}_i) \bar{\mathbf{G}}_d (\bar{\mathbf{R}}_{ij}^\top \otimes \mathbf{V}_i) \ 0] \quad (58)$$

where $\bar{\mathbf{G}}_d$ is the matrix of vectorized generators for the Lie algebra $\mathfrak{so}(d)$ (as in Section 3). Again we see that the last columns, corresponding to the vertical directions, are zero.

The astute reader may now wonder whether the rank-deficiency of these Jacobians poses any numerical difficulties when solving the linear systems needed to compute the update step δ . In fact there are several straightforward ways to address this. One approach is simply to employ the Levenberg-Marquardt method directly in conjunction with the Jacobians (58); in this case, the Tikhonov regularization applied by the LM algorithm itself will ensure that all of the linear systems to be solved are nonsingular. Moreover, this regularization will additionally encourage update steps to lie in the horizontal subspace, since any update with a nonzero vertical component will have to “pay” for the length of that component (via regularization), while “gaining” nothing for it (in terms of reducing the local model of the objective).

Alternatively, one can remove the final $(p - d)$ all-0’s columns from the Jacobians in (58), and solve the resulting reduced linear system in the variables (ω, \mathbf{K}) . Geometrically, this corresponds to minimizing the local quadratic model of the objective (46) over the horizontal subspace; a horizontal full-space update $[\delta]$ can then be obtained by simply taking $\gamma = 0$. It is straightforward to see that this procedure corresponds to computing a pseudoinverse (minimum-norm) minimizer of the quadratic model (46).

Finally, a third approach is to regularize the original Shonan Averaging problem (2) by adding a prior term on the Karcher mean of the rotations \mathbf{Q}_i in \mathbf{Q} ; this has the effect of fixing the gauge for the underlying estimation problem, similarly to the use of “inner constraints” in photogrammetry [34].

C More Experimental Results on the YFCC Datasets

In this section we present more extensive results on the datasets derived by Heinly et al. [21] from the large-scale YFCC-100M dataset [33,21]. As in the main paper, the relative measurements for these were derived from the SFM solution provided with the data, and corrupted with noise as before, using $\sigma = 0.2$. For all results below, minimum, average, and maximum running times (in seconds) are computed over 10 random initializations for each dataset. Also shown is the fraction of cases in which the method converges to a global minimizer. All Shonan Averaging variants examined below use the same Levenberg-Marquardt non-linear optimizer, with a Jacobi-preconditioned conjugate gradient method as the linear system solver.

C.1 Small Datasets ($n < 50$)

Table 4 shows additional experimental results on the small YFCC datasets (with $n < 50$) with a more systematic exploration of the Shonan parameters. The intent is to provide more quantitative results for Shonan Averaging and its convergence properties at different levels of p , as well as explore parameter settings for optimal performance in practical settings.

In particular, we compared

- SA: Shonan Averaging with $p_{min} = 5$ and $p_{max} = 30$.
- BD: the block coordinate descent method from [13,14].
- SL: Same as SA, but starting (L)ow, from $SO(3)$: $p_{min} = 3$ and $p_{max} = 30$.
- S3: Only run with $p = 3$, which corresponds to LM in the main paper: we only optimize at the base $SO(3)$ level. Note that this approach has no guarantee of converging to a global minimizer.
- S4: Similar to S3, but with $P = 4$: assesses in what percentage of cases we converge to global minima for $p = 4$.
- S5: Similar to S3 and S4, but for $p = 5$.
- SK: Shonan Averaging with $p_{min} = 5$ and $p_{max} = 30$, i.e., the same as SA, but with a different prior to fix the gauge freedom.

The last version, SK, was inspired by [38], who stressed the importance of fixing the gauge symmetry to make the rotation averaging problem better-behaved. In particular, we fixed the Karcher mean of all rotations (for any level p) to remain at its initial value, similar to the “inner constraints” often used in photogrammetry [34].

In Table 4, we have indicated the best performing method out of SA, SL, and SK. The specialized solvers S3, S4, and S5 that only optimize at one level are not considered in the comparison, as they are not guaranteed to converge to a global minimum, and are only there to provide insight into the relative performance in those different $SO(p)$ spaces.

dataset	method	error	min	avg	max	success
statue_of_liberty_1 (n=19, m=54)	SA	0.000%	0.038	0.313	1.211	100%
	SL	0.000%	0.010	0.219	0.901	100%
	S3	nan%	nan	nan	nan	0%
	S4	0.001%	0.009	0.016	0.019	30%
	S5	0.001%	0.010	0.018	0.022	40%
	SK	-0.000%	0.008	0.108	0.459	100%
natural_history_museum_london (n=30, m=274)	SA	0.000%	0.019	0.036	0.049	100%
	SL	0.000%	0.011	0.021	0.068	100%
	S3	0.000%	0.010	0.013	0.015	60%
	S4	0.000%	0.011	0.016	0.021	100%
	S5	0.000%	0.009	0.014	0.018	100%
	SK	-0.000%	0.021	0.022	0.024	100%
statue_of_liberty_2 (n=39, m=156)	SA	0.000%	0.030	0.063	0.094	100%
	SL	0.001%	0.011	0.034	0.060	100%
	S3	0.000%	0.010	0.028	0.046	40%
	S4	0.000%	0.011	0.024	0.047	100%
	S5	0.000%	0.010	0.030	0.057	100%
	SK	0.000%	0.019	0.050	0.113	100%
taj_mahal_entrance (n=42, m=1272)	SA	0.000%	0.071	0.117	0.165	100%
	SL	0.000%	0.032	0.043	0.062	100%
	S3	0.000%	0.037	0.046	0.063	80%
	S4	0.000%	0.033	0.042	0.051	100%
	S5	0.000%	0.032	0.037	0.039	100%
	SK	-0.000%	0.070	0.081	0.092	100%
sistine_chapel_ceiling_1 (n=49, m=1754)	SA	0.000%	0.102	0.173	0.246	100%
	SL	0.000%	0.064	0.085	0.108	100%
	S3	0.000%	0.057	0.073	0.087	60%
	S4	0.000%	0.072	0.121	0.293	100%
	S5	0.000%	0.064	0.083	0.095	100%
	SK	-0.000%	0.102	0.116	0.130	100%

Table 4. More results on YFCC datasets with $n < 50$. In this table, we compare 6 methods (see text for details). For each, we show the relative error with respect to SA, we give the minimum, average, and maximum running times (in seconds), and the fraction of cases in which the method converges to a global minimizer.

Conclusions For these small datasets, block-coordinate descent [13,14] performs very well. Even so, Shonan averaging with Karcher mean is faster in several cases. One such example in this Table is the Statue of Liberty dataset with $n = 39$ and $m = 156$. Interestingly, SK appears to be the fastest Shonan Averaging variant for some datasets, despite the fact that it contains an additional *dense* term in the objective that involves *all* of the poses. We conjecture that for these relatively small datasets, the inclusion of the prior on the Karcher mean helps to promote faster convergence of the manifold optimization by penalizing components of the update step that lie in the subspace of (global) gauge symmetry directions for the rotation averaging problem. Intuitively, it discourages the step from having a

component that does not “actually change” the solution. This can be important in the context of a trust-region method like ours, where the *total* length of the step is restricted at each iteration. It would also be interesting to investigate what (if any) effect the inclusion of the Karcher mean has on the presence of suboptimal critical points at each level of the Riemannian Staircase, although we leave these questions for future work.

Also clear is that starting Shonan Averaging with $p = 3$, shown as SL in the table, is always either on par or much faster than SA. There is a simple explanation for this: in many instances, it is possible to recover a global minimizer from the optimization at the lowest level $p = 3$. This can be appreciated by comparing with the results of S3: it rarely finds global minima *every* time, but when it does it is obviously the fastest of all methods. In SL, we only move to the next $SO(p)$ level if that does *not* happen, and **hence we get the best of both worlds: fast convergence if we happened to pick a lucky initial estimate, and upgrade to global optimality if not.**

The S3, S4, and S5 lines are shown to indicate at what level this occurs, and for these datasets it is almost always at $p = 4$. However, the results reveal that there are indeed no guarantees, so it is not a good idea to run Shonan Averaging at a single level: the Riemannian Staircase provides the global guarantee but at minimal extra cost, as it is *only* triggered when we converge to a suboptimal critical point at a lower level p .

C.2 Intermediate-size Datasets ($50 \leq n < 150$)

In Tables 5 and 6 we show additional results on increasingly larger YFCC datasets, with exactly the same parameters as in the previous section. However, here we omit the S3-S5 variants *unless* they do not converge to global minima in *all* tested cases.

Conclusion From these results **it is clear that SL starts to emerge as the best among the global optimization methods in terms of average running time.** Again, we observe that for these larger datasets it is rare *not* to converge to a global minimizer at $p = 3$, which is interesting in its own right. Of course, there are some exceptions, e.g., the Big Ben dataset in Table 6 with $n = 101$ and $m = 1880$, for which global minimizers were *not* found using only local search with $p \in \{3, 4, 5\}$. The BD method is still competitive in cases where the number of measurements is large, as BD’s running time is dominated by the number of images n , given that it is optimizing for a $3n \times 3n$ PSD matrix.

Results on these larger datasets suggest that finding global minima is actually *harder* in general for small datasets than for larger, well connected datasets. We can gain some theoretical insight into this empirical finding in light of several recent works [38, 27, 5] that have studied the connection between graph-theoretic properties of the measurement network $G = (n, \mathcal{E})$ that underpins the rotation averaging problem, and the statistical and geometric/computational properties of the resulting maximum-likelihood estimation (1). In a nutshell, these investigations indicate that *both* the statistical properties of the maximum-likelihood

dataset	method	error	min	avg	max	success
sistine_chapel_ceiling_2 (n=51, m=1670)	SA	0.000%	0.101	0.162	0.234	100%
	SL	0.000%	0.065	0.115	0.332	100%
	S3	0.000%	0.076	0.081	0.084	50%
	SK	-0.000%	0.097	0.127	0.211	100%
milan_cathedral (n=69, m=2782)	SA	0.000%	0.175	0.203	0.220	100%
	SL	0.000%	0.078	0.092	0.105	100%
	SK	-0.000%	0.170	0.190	0.206	100%
reichstag (n=71, m=2554)	SA	0.000%	0.166	0.291	0.399	100%
	SL	0.000%	0.073	0.084	0.108	100%
	SK	-0.000%	0.179	0.193	0.215	100%
piazza_dei_miracoli (n=74, m=3456)	SA	0.000%	0.261	0.353	0.593	100%
	SL	0.000%	0.092	0.119	0.147	100%
	S3	0.000%	0.102	0.119	0.146	90%
	SK	-0.000%	0.260	0.344	0.413	100%
ruins_of_st_pauls (n=82, m=4998)	SA	0.000%	0.337	0.426	0.735	100%
	SL	0.000%	0.149	0.173	0.259	100%
	SK	-0.000%	0.372	0.450	0.544	100%
mount_rushmore (n=83, m=4012)	SA	0.000%	0.264	0.366	0.511	100%
	SL	0.000%	0.105	0.123	0.193	100%
	SK	-0.000%	0.263	0.306	0.352	100%
london_bridge_3 (n=88, m=1500)	SA	0.000%	0.113	0.141	0.243	100%
	SL	0.000%	0.050	0.070	0.095	100%
	S3	0.000%	0.052	0.066	0.084	90%
	SK	-0.000%	0.115	0.128	0.143	100%
palace_of_versailles_chapel (n=91, m=3964)	SA	0.000%	0.296	0.319	0.379	100%
	SL	0.000%	0.102	0.118	0.138	100%
	SK	-0.000%	0.291	0.313	0.348	100%
pieta_michelangelo (n=93, m=7728)	SA	0.000%	0.540	0.709	1.051	100%
	SL	0.000%	0.194	0.220	0.228	100%
	SK	-0.000%	0.538	0.559	0.614	100%
blue_mosque_interior_2 (n=95, m=2288)	SA	0.000%	0.169	0.220	0.348	100%
	SL	0.000%	0.075	0.149	0.510	100%
	S3	0.000%	0.087	0.106	0.140	90%
	SK	-0.000%	0.172	0.200	0.222	100%
st_vitus_cathedral (n=97, m=8334)	SA	0.000%	0.592	0.909	1.306	100%
	SL	0.000%	0.250	0.291	0.326	100%
	S3	0.000%	0.252	0.298	0.387	90%
	SK	-0.000%	0.571	0.650	0.805	100%

Table 5. YFCC results for $50 \leq n \leq 100$. Same format as Table 4 but only showing S3-S5 results if they do not converge to global minima in *all* tested cases.

estimation (1) and its computational hardness are controlled by the *algebraic connectivity* $\lambda_2(L(G))$, i.e., the *second*-smallest eigenvalue of the weighted graph Laplacian $L(G)$ associated with the graph G . Larger values imply *both* a better (lower-uncertainty) estimate *and* that the resulting relaxation (16) is stronger. For densely-connected measurement networks (of the kind that frequently appear

dataset	method	error	min	avg	max	success
big_ben.1 (n=101, m=1880)	SA	0.000%	0.196	4.084	30.563	100%
	SL	0.000%	0.080	0.348	1.285	100%
	S3	0.000%	0.105	0.127	0.199	70%
	S4	0.000%	0.090	0.118	0.172	60%
	S5	0.000%	0.077	0.140	0.197	80%
	SK	0.000%	0.185	1.001	5.572	100%
london_bridge.2 (n=106, m=2742)	SA	0.000%	0.220	0.351	1.218	100%
	SL	0.000%	0.091	0.114	0.136	100%
	SK	-0.000%	0.222	0.243	0.264	100%
palazzo_pubblico (n=112, m=4420)	SA	0.000%	0.331	0.373	0.427	100%
	SL	0.000%	0.155	0.262	0.926	100%
	SK	-0.000%	0.340	0.382	0.442	100%
london_bridge.1 (n=118, m=5690)	SA	0.000%	0.416	0.472	0.543	100%
	SL	0.000%	0.168	0.203	0.285	100%
	SK	-0.000%	0.423	0.469	0.536	100%
national_gallery_london (n=124, m=2160)	SA	0.000%	0.186	0.233	0.351	100%
	SL	0.000%	0.089	0.107	0.131	100%
	SK	-0.000%	0.179	0.210	0.246	100%
lincoln_memorial (n=127, m=3516)	SA	0.000%	0.324	0.374	0.548	100%
	SL	0.000%	0.119	0.135	0.185	100%
	SK	-0.000%	0.259	0.301	0.329	100%
grand_central_terminal_new_york (n=132, m=5880)	SA	0.000%	0.414	0.584	1.043	100%
	SL	0.000%	0.187	0.220	0.291	100%
	SK	-0.000%	0.528	0.586	0.661	100%
paris_opera.2 (n=133, m=10778)	SA	0.000%	0.809	0.888	0.982	100%
	SL	0.000%	0.341	0.395	0.478	100%
	SK	-0.000%	0.837	0.906	1.058	100%

Table 6. YFCC results for $100 \leq n \leq 150$. Same layout as Table 5.

in structure-from-motion applications), it is an elementary result from algebraic graph theory that this quantity can grow at a rate of up to $O(n)$. This provides insight into the observation that problems with larger measurement networks appear easier to solve, assuming a reasonably dense set of measurements.

C.3 Larger Datasets ($n \geq 150$)

Finally, in Table 7 we show additional results on the largest YFCC datasets with $n \geq 150$. The block-coordinate descent method from [14] did not converge in reasonable time for many of the larger datasets, which is because we use the minimum eigenvalue optimality certificate threshold λ_{min} to establish convergence. The threshold we used in all experiments was 10^{-5} , and for these datasets it takes a long time for BD to reach that level, in contrast to Shonan Averaging.

Conclusion For these large datasets convergence to the global minimizer occurs almost always at $p = 3$ and hence the SL and S3 methods are basically identical in

terms of operation and performance. However, SL comes with a global guarantee: in the rare case that S3 does not converge to the global minimizer at the $SO(3)$ level, SL will simply move up to $SO(4)$ and up, and thereby still recover the true maximum likelihood estimate.

dataset	method	error	min	avg	max	success
st_peters_basilica_interior_2 (n=173, m=11688)	SA	0.000%	0.948	1.165	1.738	100%
	SL	0.000%	0.557	0.713	0.897	100%
	S3	0.000%	0.531	0.587	0.715	80%
pantheon_interior (n=186, m=10000)	SA	0.000%	0.852	1.001	1.570	100%
	SL	0.000%	0.449	0.509	0.623	100%
	S3	0.000%	0.455	0.530	0.642	100%
florence_cathedral_dome_interior_1 (n=213, m=31040)	SA	0.000%	2.435	2.848	3.226	100%
	SL	0.000%	1.869	2.170	2.499	100%
	S3	0.000%	1.856	2.209	2.814	100%
paris_opera_1 (n=254, m=45754)	SA	0.000%	3.673	4.145	4.467	100%
	SL	0.000%	1.686	2.024	2.677	100%
	S3	0.000%	1.683	2.104	2.510	100%
pike_place_market (n=265, m=53242)	SA	0.000%	4.437	5.255	7.285	100%
	SL	0.000%	2.019	2.322	3.400	100%
	S3	0.000%	2.122	2.246	2.384	100%
blue_mosque_interior_1 (n=272, m=40292)	SA	0.000%	3.581	4.153	4.964	100%
	SL	0.000%	2.268	2.708	2.991	100%
	S3	0.000%	2.511	2.997	3.645	100%
notre_dame_rosary_window (n=326, m=93104)	SA	0.000%	8.102	8.652	10.515	100%
	SL	0.000%	3.478	4.074	4.690	100%
	S3	0.000%	3.885	4.326	5.012	100%
british_museum (n=344, m=45450)	SA	0.000%	4.297	4.851	7.846	100%
	SL	0.000%	1.927	2.558	3.400	100%
	S3	0.000%	2.009	2.558	3.911	100%
palace_of_westminster (n=345, m=11522)	SA	0.000%	1.334	1.480	1.692	100%
	SL	0.000%	0.598	0.837	1.242	100%
	S3	0.000%	0.731	1.038	2.397	100%
louvre (n=367, m=26656)	SA	0.000%	2.872	3.141	3.709	100%
	SL	0.000%	1.606	2.284	3.265	100%
	S3	0.000%	1.648	2.058	2.495	100%
st_peters_basilica_interior_1 (n=365, m=55024)	SA	0.000%	5.283	5.654	6.295	100%
	SL	0.000%	3.255	4.030	4.995	100%
	S3	0.000%	3.662	4.169	4.728	100%
st_pauls_cathedral (n=370, m=83060)	SA	0.000%	7.385	7.861	8.384	100%
	SL	0.000%	3.385	3.792	4.373	100%
	S3	0.000%	3.420	4.176	6.689	100%
westminster_abbey_1 (n=501, m=38863)	SA	0.000%	4.238	4.582	4.957	100%
	SL	0.000%	2.059	2.396	2.710	100%
	S3	0.000%	2.209	2.693	3.385	100%
pantheon_exterior (n=720, m=49256)	SA	0.000%	6.922	7.876	11.202	100%
	SL	0.000%	3.798	5.247	6.765	100%
	S3	0.000%	4.091	4.945	6.156	100%

Table 7. Comparing SA, SL and S3 (see text) on YFCC Datasets with $n \geq 150$. In many cases S3 has a lower average computation time than SL, since it performs only local optimization at the *lowest* level $p = 3$ of the Riemannian Staircase. However, in contrast to SL, S3 has *no* guarantees regarding convergence to global minima.