

Rotation averaging and weak convexity

Richard Hartley, Jochen Trumpf and Yuchao Dai

Abstract— We generalize the concept of geodesic convexity in the Special Orthogonal Group $SO(3)$ and apply the generalization to the discussion of rotation averaging. As a result we are able to derive strong and new theorems about the location of global minima of the rotation averaging cost function. A brief discussion of the relationship of our results to previous results from the literature will be provided, as well as an application to camera rig calibration in computer vision.

I. INTRODUCTION

Rotation averaging refers to a particular optimisation problem on the Special Orthogonal Group $SO(3)$ in real three dimensional space that generalises the notion of mean in vector spaces. Given a (finite) set of rotations, the goal is to find a rotation that minimizes the sum of distances or sum of squared distances to the given rotations. Such a minimizing rotation is called a mean or averaged rotation.

There are several natural and non-equivalent choices for distances on $SO(3)$, including the geodesic distance with respect to the standard biinvariant Riemannian metric, the chordal distance in the embedding 9-dimensional matrix space and the quaternion distance, leading to a large variety of means with different sets of properties.

We focus on issues such as uniqueness and convexity, in particular for the mean with respect to the squared geodesic distance. Among the previous work in this area are contributions by Karcher [7], Pennec [12], Corcuera and Kendall [1], Gramkow [5], Moakher [11], Manton [10], Krakowski et al. [8], Sarlette and Sepulchre [13] and Fletcher et al. [4].

Our main contribution is a generalisation of the concept of geodesic convexity which we call weak convexity. The generalisation works by giving up one of the usual axioms of geodesic convexity, namely that a connecting geodesic arc necessarily has minimal length. It turns out that this new concept allows us to link the Riemannian and convex geometries of $SO(3)$ very closely to those of three dimensional real projective space which in turn leads to nice characterizations of the convexity properties of the cost functions involved in rotation averaging. As a result, we are able to prove strong and new theorems about the location of global minima of the averaging cost functions. Note that our concept of weak convexity is different from Corcuera and Kendall's concept of generalized convexity where the sacrificed axiom is that of connecting geodesic arcs being unique [1].

We briefly touch on an application of this theory to camera rig calibration, a problem in computer vision [2].

R. Hartley and J. Trumpf are with the School of Engineering, The Australian National University, Canberra ACT 0200, Australia (Richard.Hartley|Jochen.Trumpf)@anu.edu.au

Y. Dai is with the School of Electronics and Information, Northwestern Polytechnical University, P.R.China daiyuchao@gmail.com

II. DISTANCE MEASURES FOR ROTATIONS

In this paper we mainly use the following standard matrix representation of the Special Orthogonal group $SO(3)$ of rotations in real three dimensional space.

$$SO(3) = \{R \in \mathbb{R}^{3 \times 3} \mid R^T R = I, \det R = 1\}$$

Any rotation in $SO(3)$ can be expressed as a rotation through a given angle θ about some axis. We define the *angular distance* between two rotations S and R to be the angle of the rotation SR^T . Thus,

$$d_{\angle}(S, R) = \frac{1}{\sqrt{2}} \|\log(SR^T)\|_{\text{Fro}}$$

where the norm is the Frobenius norm. The logarithm can be computed using the formula

$$\log(R) = \begin{cases} \arcsin(\|y\|_2) \frac{Y}{\|y\|_2}, & Y \neq 0 \\ 0, & Y = 0 \end{cases}$$

where $y = \text{vex}(Y) = (y_1, y_2, y_3)$ is computed from

$$\frac{1}{2}(R - R^T) = Y = \begin{pmatrix} 0 & -y_3 & y_2 \\ y_3 & 0 & -y_1 \\ -y_2 & y_1 & 0 \end{pmatrix}$$

and $\hat{y} = y/\|y\|_2$ is the rotation axis of R . Note that by this definition, the angular distance between two rotations is at most equal to π and is in fact equal to the geodesic distance under the standard biinvariant Riemannian metric on $SO(3)$. The logarithm map above is an inverse of the Riemannian exponential map given by the matrix exponential $\exp: \mathfrak{so}(3) \rightarrow SO(3)$ in this representation where $\mathfrak{so}(3)$ is the set of skew-symmetric matrices, the Lie algebra associated with the Lie group $SO(3)$. The Lie bracket is given by the matrix commutator $[A, B] = AB - BA$.

The *chordal distance* between two rotations S and R is the Euclidean distance between them in the embedding space $\mathbb{R}^{3 \times 3}$. Thus,

$$d_{\text{chord}}(S, R) = \|S - R\|_{\text{Fro}} = 2\sqrt{2} \sin(\theta/2)$$

where $\theta = d_{\angle}(S, R)$. The relationship between the angular distance and the chordal distance is most easily established using Rodrigues' formula.

Another distance measure derives from the Euclidean distance between two quaternions in the embedding space \mathbb{R}^4 . Any rotation R may be represented by a unit quaternion r as follows. If \hat{y} is the axis of the rotation and θ is the angle of the rotation about that axis, then r is defined as

$$r = (\cos(\theta/2), \hat{y} \sin(\theta/2)).$$

The converse mapping from the set of unit quaternions Q to $SO(3)$ given by

$$(\cos(\theta/2), \hat{y} \sin(\theta/2)) \mapsto \exp \text{vex}(\theta \hat{y})$$

is a Lie group homomorphism since the quaternion multiplication corresponds to matrix multiplication under this mapping. Both r and $-r$ represent the same rotation, that is, the homomorphism from Q to $SO(3)$ is a 2-to-1 mapping. Topologically, the unit quaternions form a unit sphere S^3 in \mathbb{R}^4 , and hence the above identifies $Q \simeq S^3$ as a two-fold covering space (or double cover) of $SO(3)$. If we restrict ourselves to rotations through angles less than π then these are in 1-to-1 correspondence to points of the upper quaternion hemisphere with the “north pole” $(1, 0, 0, 0)$ corresponding to the identity rotation (rotation through an angle of 0). In this picture, the “equator” of the quaternion sphere corresponds exactly to the rotations through an angle of π with opposite points on the equator representing the same rotation. We may think to define a *quaternion distance* $d_{\text{quat}}(S, R)$ between two rotations to be $d_{\text{quat}}(S, R) = \|s - r\|_2$, where s and r are quaternion representations of S and R , respectively. Unfortunately, this simple equation will not do, since both r and $-r$ represent the same rotation, and it is not clear which one to choose (and analogous for s and $-s$, of course). However, this is resolved by defining

$$d_{\text{quat}}(S, R) = \min\{\|s - r\|_2, \|s + r\|_2\} = 2 \sin(\theta/4)$$

where once again $\theta = d_{\perp}(S, R)$. This follows from a simple geometric argument on $Q \simeq S^3$.

Plots of the three different distance functions discussed so far, plotted as functions of the angular distance are shown in Figure 1. While the distances themselves are not equivalent, they all yield the same *path metric* on $SO(3)$ in that the shortest curves connecting two points on $SO(3)$ are the same under all three distance measures, namely the geodesics with respect to the standard biinvariant Riemannian metric. This follows from an infinitesimal comparison argument. These shortest paths can hence be pictured as great circle arcs on the quaternion unit hemisphere.

Starting from the representation of $SO(3)$ as the quaternion sphere visualized as the unit sphere S^3 embedded in \mathbb{R}^4 , the *gnomonic projection* to \mathbb{R}^3 is the projection from the centre of the sphere, $(0, 0, 0, 0)$, onto a tangent (3-dimensional) hyper-plane. For simplicity, we may consider this to be the tangent hyper-plane passing through the point $(1, 0, 0, 0)$ on S^3 , representing the identity rotation. Clearly, this is a 2-to-1 projection of S^3 , since opposite points on the sphere project to the same point. Since a great circle on S^3 is the intersection of S^3 with a (2-dimensional) plane passing through the centre point $(0, 0, 0, 0)$, namely the plane spanned by the radius vector of any point on the great circle and a tangent vector along the great circle at that point, we easily see that the projection of a great circle is the intersection between this plane and the projection hyper-plane. This shows that the projection of a great circle on S^3 is a straight line in the projection hyper-plane. This type of map is sometimes also called a Beltrami map in the literature.




Fig. 1. **Distance metrics.** On the left (top to bottom) are angular, quaternion and chordal distances plotted as a function of rotation angle. On the right the squared distances. Plots are shown for rotation angles from 0 to π . The plots of the chordal and quaternion metrics are scaled to be comparable with the angle metric.

In S^3 , the “equator” is the intersection of the “equatorial




Fig. 2. **Gnomonic projection of a sphere.**

hyper-plane” consisting of points $(0, x, y, z)$, with the sphere. Projecting from the origin, we see that the equator maps to the “plane at infinity” in \mathbb{R}^3 . More exactly, we see that the gnomonic projection maps S^3 to $\mathbb{R}^3 \cup \pi_{\infty}$, which is a 3-dimensional projective space, topologically homeomorphic to $SO(3)$. Geodesics in $SO(3)$ correspond to straight-lines in \mathbb{R}^3 along with straight lines in the plane at infinity. We will see later that this representation of SO is particularly useful when it comes to concepts like geodesics and convexity. The above described the gnomonic projection localized at the identity rotation, since the tangent hyper-plane was chosen to pass through the point in the quaternion sphere representing

the identity rotation. One may equally well construct a gnomonic projection, with similar properties, localized about any other rotation (point on the quaternion sphere).

III. WEAK CONVEXITY

The general definition of convexity of a function in \mathbb{R}^n is as follows. Given a convex region $U \subset \mathbb{R}^n$ a function f defined on U is convex if for any two points x_0 and x_1 in U , and any point y lying on the line segment bounded by x_0 and x_1 , given by $y = (1 - \lambda)x_0 + \lambda x_1$ with $0 \leq \lambda \leq 1$, we have

$$f(y) \leq (1 - \lambda)f(x_0) + \lambda f(x_1).$$

In adapting this definition to $SO(3)$, or indeed to any Riemannian manifold, the role of a line is naturally taken by a geodesic.

As discussed in the last section the geodesics on $SO(3)$ are doubly-covered by great circles on S^3 and there is a uniform length scaling by a factor of 2 between the geodesics on $SO(3)$ and those on S^3 . In particular, we see that the geodesics on $SO(3)$ are closed curves with a total length of 2π . There are exactly two geodesic segments between any two points in $SO(3)$ (without exception). Given two points (rotations) R_0 and R_1 in $SO(3)$, we call the shorter of the two geodesic segments from R_0 to R_1 the *short geodesic segment* between these points. If R_0 and R_1 differ by a rotation through π , then which of the two geodesic segments is the shorter one is ambiguous and hence there is no short geodesic segment between such points. We can now define two slightly different notions of geodesic convexity of sets in $SO(3)$.

Definition 1: A non-empty region $U \subset SO(3)$ is called *weakly convex* if for any two points R_0 and R_1 in U exactly one geodesic segment from R_0 to R_1 lies entirely inside U . A weakly convex region $U \subset SO(3)$ is called *convex* if the geodesic segment from R_0 to R_1 in U is always the short geodesic segment between these points.

A closed ball of radius $r \geq 0$ in $SO(3)$ is a set

$$B(R, r) = \{S \in SO(3) \mid d_{\angle}(S, R) \leq r\}$$

for some R in $SO(3)$. An open ball of radius $r > 0$ in $SO(3)$, denoted $\dot{B}(R, r)$, is the interior of the closed ball, consisting of rotations at distance strictly less than r from R . We emphasize for clarity that the balls $B(R, r)$ or $\dot{B}(R, r)$ are defined in terms of the angular distance on $SO(3)$.

Lemma 1: A closed ball in $SO(3)$ is convex if and only if its radius is less than $\pi/2$. Similarly, an open ball in $SO(3)$ is convex if and only if its radius is less than or equal to $\pi/2$. A closed ball in $SO(3)$ is weakly convex if and only if its radius is less than π , and an open ball on $SO(3)$ is weakly convex if and only if its radius is less than or equal to π .

We will not include any proofs of results in this short paper. Proofs of all the statements made will be reported elsewhere and can be obtained from the authors by request. Convex and weakly convex subsets of $SO(3)$ can not be arbitrarily “large”, in the following precise sense.

Theorem 1 (Size of convex sets): Any convex subset of $SO(3)$ is contained in a closed ball of radius $2\pi/3$. Any weakly convex subset of $SO(3)$ is contained in an open ball of radius π .

The proof of this theorem turns out to be surprisingly difficult (particularly the second part). As a consequence of this result we may picture any weakly convex subset of $SO(3)$ simply as a convex set in \mathbb{R}^3 under a suitably chosen gnomonic projection, namely the one mapping the boundary of the containing ball of radius π to the plane at infinity. This is because the gnomonic projection maps geodesics to geodesics, and hence weakly convex sets to convex sets.

Some of the notions from classical convex geometry in \mathbb{R}^3 carry over to weakly convex or convex sets in $SO(3)$ via this mapping. It follows immediately that the closure of a convex set in $SO(3)$ is convex. This is not always true for weakly convex sets. We can define the *affine hull* $\text{aff}(C)$ of a weakly convex set $C \neq \emptyset$ to be the preimage of the \mathbb{R}^3 based affine hull under the gnomonic map completed by the missing “points at infinity”, i.e. the points in $SO(3)$ closing the geodesics contained in this preimage. Equivalently, $\text{aff}(C)$ is the preimage of the projective closure of the \mathbb{R}^3 based affine hull under the gnomonic map. It is clear that this “affine” hull is a compact submanifold of $SO(3)$ of the same dimension as the classical affine hull of the gnomonic picture of C that consists entirely of whole, closed geodesics. It is hence in particular independent of where the gnomonic projection was centered. This means it is well defined as an intrinsic geometric object in $SO(3)$. In dimension 2, we will refer to such an “affine” submanifold as a *geodesic plane*. In dimension 0 it consists of a single point, in dimension 1 of a single whole, closed geodesic and in dimension 3 it is equal to $SO(3)$ itself. As usual, the *relative interior* $\text{ri}(C)$ of C is now the interior with respect to the topology of $\text{aff}(C)$, and in general the relative interior $\text{ri}(C)$ is different from the interior \dot{C} , unless $\dim \text{aff}(C) = 3$ whence $\text{aff}(C) = SO(3)$. Also, the *relative boundary* $\bar{C} \setminus \text{ri}(C)$ is in general different from the boundary $\bar{C} \setminus \dot{C}$.

Proposition 1: The intersection of two convex sets in $SO(3)$ is convex. This is not true of weakly convex sets in general. However, the intersection of a weakly convex set with a geodesic plane is weakly convex. Moreover, the intersection of two weakly convex sets consists of at most two weakly convex and pairwise disjoint components.

The previous proposition implies that there is no natural concept of a (unique) “weakly convex hull” of a given subset of $SO(3)$. Hence some of the machinery of convex analysis does not immediately generalize to the weakly convex setting. However, convex functions can be defined as in \mathbb{R}^n , except that geodesics in $SO(3)$ take the place of straight lines joining two points in \mathbb{R}^n .

Definition 2: Consider a function $f : U \rightarrow \mathbb{R}$ defined on a weakly convex subset U of $SO(3)$. Let $R_0, R_1 \in U$ and let $g : [0, 1] \rightarrow U$ be a geodesic path (constant speed parametrization of a geodesic arc) in U , such that $g(0) = R_0$ and $g(1) = R_1$. The function f is called *convex*, if for any

such R_0, R_1 and g , we have an inequality

$$f(g(\lambda)) \leq (1 - \lambda)f(R_0) + \lambda f(R_1)$$

for all $\lambda \in [0, 1]$. The function is called *strictly convex* if this inequality is strict for all $\lambda \in (0, 1)$ whenever $R_0 \neq R_1$.

Various properties of convex functions hold true, just as with convex functions in \mathbb{R}^n .

Proposition 2: The sum of convex (or strictly convex) functions defined on a weakly convex region U is convex (respectively, strictly convex). A strictly convex function defined on a weakly convex set has at most a single global minimum; for convex functions, any local minimum is a global minimum and the minima form a weakly convex set on which the function is constant.

The proof is the same as for convex functions in \mathbb{R}^n . Convexity of functions can be defined locally through computing the second derivative of their restriction along geodesic paths through a point.

Definition 3: A function $f : \text{SO}(3) \rightarrow \mathbb{R}$ is *locally convex* at a point $R_0 \in \text{SO}$ if for any constant speed geodesic path $s : [-1, 1] \rightarrow \text{SO}(3)$, with $s(0) = R_0$ the function $f \circ s(t) = f(s(t))$ has non-negative second derivative at $t = 0$. It is *locally strictly convex* at R_0 if any such $f \circ s(t)$ has positive second derivative at $t = 0$.

The connection between local convexity and convexity is as follows.

Proposition 3: If $f : \text{SO}(3) \rightarrow \mathbb{R}$ is smooth and locally convex (or strictly convex) at each point in a weakly convex set U , except possibly at isolated global minima of f , then it is convex (respectively, strictly convex) in U . If $f : \text{SO}(3) \rightarrow \mathbb{R}$ is smooth but not locally convex at some point then it is not convex in any non-trivial ball around that point.

Next we investigate when the distance function $d(S, R)$ defined for two rotations is a convex function of S (for fixed R).

Theorem 2 (Convexity of distances): Consider the function $f(S) = d(S, R)^p$ for a fixed rotation R , a distance $d(\cdot, \cdot)$, and an exponent p . The function is convex, or strictly convex, as a function of S in the following cases.

- 1) $d_{\angle}(\cdot, R)$ is convex on the set $\mathring{B}(R, \pi)$.
- 2) $d_{\text{chord}}(\cdot, R)$ is not convex on any non-trivial ball around R .
- 3) $d_{\text{quat}}(\cdot, R)$ is not convex on any non-trivial ball around R .
- 4) $d_{\angle}(\cdot, R)^2$ is strictly convex on the set $\mathring{B}(R, \pi)$.
- 5) $d_{\text{chord}}(\cdot, R)^2$ is strictly convex on the set $\mathring{B}(R, \pi/2)$ and convex on the set $B(R, \pi/2)$.
- 6) $d_{\text{quat}}(\cdot, R)^2$ is strictly convex on the set $\mathring{B}(R, \pi)$.

This theorem is proved by calculating the eigenvalues of the Riemannian Hessian of these distance functions. Compare these results to the graphs in Figure 1.

IV. THE MAIN RESULTS

Given an exponent $p \geq 1$ and a set of $n \geq 1$ rotations $\{R_1, \dots, R_n\} \subset \text{SO}(3)$ we define the *L^p -mean rotation with*

respect to d as

$$d^p - \text{mean}(\{R_1, \dots, R_n\}) = \underset{R \in \text{SO}(3)}{\text{argmin}} \sum_{i=1}^n d(R_i, R)^p.$$

Since $\text{SO}(3)$ is compact, a minimum will exist as long as the distance function is continuous. In this problem, we will be given a set of rotations R_i , and the task will be to find the mean of these rotations, as defined previously. This problem has been much studied in the literature, but there are still open problems, some of which are resolved here. We consider the L^2 -mean with respect to the angular distance d_{\angle} . This mean is also known as the *Karcher mean* [7] or the *geometric mean* [11].

Theorem 3: Given rotations $R_i, i = 1, \dots, n$, the cost function $C(R) = \sum_{i=1}^n d_{\angle}(R_i, R)^2$ is strictly convex, except on the union of sets

$$\pi_i = \{S \in \text{SO}(3) \mid d_{\angle}(R_i, S) = \pi\}$$

in the following sense. These sets π_i divide $\text{SO}(3)$ into at most $\binom{n}{3} + n$ regions whose interior is weakly convex. $C(R)$ is strictly convex on the interior of each of these regions and is non-differentiable on the boundary, that is, on the union of the sets π_i . The cost function $C(R)$ has at most one minimum on each of the closed regions and hence there are at most $\binom{n}{3} + n$ minima.

This theorem indicates that $\text{SO}(3)$ may be divided into a large number of individual weakly convex regions, each with its own local minimum. It may seem, therefore, that the problem of finding the global minimum is quite challenging. The following discussion shows that if the rotations R_i are not too widely separated, one of the weakly convex regions may be quite large, and will contain the global minimum.

We consider the case where all the rotations R_i lie inside a closed ball of radius $r < \pi$. Thus, let S be a rotation such that $d_{\angle}(R_i, S) \leq r < \pi$. Now, let T be a rotation in the open ball $\mathring{B}(S, \pi - r)$. Then $d_{\angle}(R_i, T) \leq d_{\angle}(R_i, S) + d_{\angle}(S, T) \leq r + d_{\angle}(S, T) < \pi$. Hence $\mathring{B}(S, \pi - r)$ is contained in all the $\mathring{B}(R_i, \pi)$ and hence also in the component of their intersection containing S . Applying this reasoning to the case where $r < \pi/2$ and noting that closed balls are compact, allows us to assert the following theorem.

Theorem 4: Let R_i be rotations satisfying $d_{\angle}(R_i, S) \leq r < \pi$ for some rotation S and for all i , then

$$C(R) = \sum_{i=1}^n d_{\angle}(R_i, R)^2$$

is strictly convex on $\mathring{B}(S, \pi - r)$, and hence has a single isolated minimum on any closed ball contained in this set.

The most interesting case is where $r < \pi/2$, so all rotations are contained in $\mathring{B}(S, \pi/2)$. The theorem then implies that the cost function is strictly convex and has a single isolated minimum on $B(S, \pi/2)$.¹ This is a useful result, since it means that there will be a single (local)

¹An alternative proof of the unique existence of this minimum (in a much more general setting) can be found in [6, Theorem 3.7].

minimum of $C(R)$ in this closed ball and that minimum may be found by convex minimization techniques. However, it is not clear that the minimum will be the global minimum of $C(R)$ on $\text{SO}(3)$. Nor is it clear whether the minimum can lie on the boundary of the closed ball. We settle these questions now with the following general theorem, our main contribution. The proof uses the $\text{SO}(3)$ -version of the law of cosines.

Theorem 5 (Location of global minima): Let C be a convex subset of $\text{SO}(3)$ and let the rotations R_i , $i = 1, \dots, n$, $n \geq 2$ be contained in the closure of C . Let f be any function strictly increasing on the interval $[0, \pi]$. Then any global minimum $R_{\min} = \operatorname{argmin}_{R \in \text{SO}(3)} \sum_{i=1}^n f(d_{\angle}(R_i, R))$ lies in the relative interior of C .

Since d_{\angle}^2 is obviously a monotonic function of d_{\angle} , we immediately get the following strong existence and uniqueness result for the global Karcher mean.

Corollary 1: Let R_i be two or more rotations satisfying $d_{\angle}(R_i, S) \leq \delta < \pi/2$ for some rotation S and for all i , then the unique global L_2 -mean with respect to d_{\angle} lies in the open ball $\bar{B}(S, \delta)$ and moreover the cost function $C(R) = \sum_{i=1}^n d_{\angle}(R_i, R)^2$ is strictly convex on the larger ball $B(S, \pi/2)$.

Note that means on manifolds can in general lie “far away” from the point set being averaged, so the above corollary places $\text{SO}(3)$ into the category of “nice geometries”; see Corcuera and Kendall [1] for a detailed discussion of this issue, but note that the notion of convexity used in that paper is subtly different from the two notions we have used in this paper.

For the situation of Corollary 1, Manton [10] has provided the following convergent algorithm where the inner loop of the algorithm is computing the average in the tangent space and then projecting back onto the manifold $\text{SO}(3)$ via the exponential map.

- 1: Set $R := R_1$. Choose a tolerance $\varepsilon > 0$.
- 2: **loop**
- 3: Compute $Y := \frac{1}{n} \sum_{i=1}^n \log(R^{\top} R_i)$.
- 4: **if** $\|Y\| < \varepsilon$ **then**
- 5: **return** R
- 6: **end if**
- 7: Update $R := R \exp(Y)$.
- 8: **end loop**

Algorithm 1: Computing the angular L_2 -mean on $\text{SO}(3)$

In fact, this algorithm is shown to be an instance of simple Riemannian gradient descent and it is shown that an implementation with arbitrary numerical accuracy would terminate only within a d_{\angle} -distance of $\delta^{-1} \tan(\delta) \varepsilon$ of the mean [10, Theorem 5]. For a Newton-type algorithm to compute this mean see [8].

The other distance measures can be discussed in a similar fashion and partly new results can be achieved this way. They will be reported elsewhere.

V. CAMERA RIG CALIBRATION

This section contains an excerpt of results published in [2]. Rotation averaging algorithms have many potential applications in the community of computer vision, especially to multi-camera systems. Here we show a potential application to the calibration of a non-overlapping multi-camera rig. Non-overlapping multi-camera rigs are of particular interest in practice. As the component cameras have little or no overlap in their fields of views, the effective overall field of view is wider, leading to efficient data acquisition.

A. Problem Formulation

We consider a camera rig consisting of two cameras, denoted left and right, fixed rigidly with respect to each other and individually calibrated. The camera rig undergoes rigid motion and captures several image pairs (cf. Figure 3).

We denote the coordinate frames of the cameras at time i by M_i^L and M_i^R .

$$M_i^L = \begin{pmatrix} L_i & t_i^L \\ 0^{\top} & 1 \end{pmatrix} \quad \text{and} \quad M_i^R = \begin{pmatrix} R_i & t_i^R \\ 0^{\top} & 1 \end{pmatrix}.$$

The first three rows of these matrices represent the projection matrices of the corresponding cameras, where image points are represented in coordinates normalized by the calibration matrix.

We denote the relative motion of M_0^R with respect to M_0^L by a transformation M^{LR} , so that $M^{LR} = M_0^R M_0^L^{-1}$. Since this relative motion remains fixed throughout the motion, we observe that $M^{LR} = M_i^R M_i^L^{-1}$ for all i . Next, the relative motion of M_j^L with respect to M_i^L is denoted by $M_{ij}^L = M_j^L M_i^L^{-1}$. Similarly, $M_{ij}^R = M_j^R M_i^R^{-1}$. Using the relation $M_i^R = M^{LR} M_i^L$, we find

$$M_{ij}^R = M^{LR} M_{ij}^L M^{LR^{-1}} \quad (1)$$

for all i, j . Now, we denote

$$M_{ij}^L = \begin{pmatrix} L_{ij} & t_{ij}^L \\ 0^{\top} & 1 \end{pmatrix} \quad \text{and} \quad M_{ij}^R = \begin{pmatrix} R_{ij} & t_{ij}^R \\ 0^{\top} & 1 \end{pmatrix}.$$

Observe that the relative rotations R_{ij} , L_{ij} and relative translations t_{ij}^R , t_{ij}^L may be computed via the essential matrix for the (i, j) image pairs.

Writing the transformation M^{LR} as

$$M^{LR} = \begin{pmatrix} S & s \\ 0^{\top} & 1 \end{pmatrix},$$

we deduce from (1) the equations

$$R_{ij} = S L_{ij} S^{-1} \quad (2)$$

$$t_{ij}^R = S t_{ij}^L + (I - L_{ij}) s \quad (3)$$

Calibration strategy Our prescribed task is to find the relative orientation of the right and left cameras, namely the transformation M^{LR} . Our method uses the following general framework.

- 1) Compute the relative rotations and translations (R_{ij}, t_{ij}^R) and (L_{ij}, t_{ij}^L) for many pairs (i, j) using the essential matrix.




Fig. 3. The calibration problem for a non-overlapping multi-camera rig.

- 2) Compute the relative rotation S from (2).
- 3) Solve linearly for s using (3).

Both these equations may be solved linearly. The rotation equation may be written as $SL_{ij} = R_{ij}S$, which is linear in the entries of S . In solving for the translation s , we note that the relative translations t_{ij}^L and t_{ij}^R are known only up to scale factors λ_{ij} and μ_{ij} . Then (3) may be written more exactly as $\lambda_{ij}t_{ij}^R = \mu_{ij}St_{ij}^L + (I - L_{ij})s$, where everything is known except for s and the scale factors λ_{ij} and μ_{ij} . Three image pairs are required to solve these equations and find s .

The strategy outlined here is workable, but relies on accurate measurements of the rotations L_{ij} and R_{ij} . In the following we will explain our strategies for rotation averaging that will lead to significantly improved results in practice. Although we have implemented the complete calibration algorithm, including estimation of the translation s , for the rest of this paper, we will consider only rotation estimation.

B. Averaging Rotations

The relative rotation estimates R_{ij} and L_{ij} obtained from individual estimates using the essential matrix will not be consistent. In particular, ideally, there should exist rotations L_i , R_i and S such that $L_{ij} = L_jL_i^{-1}$ and $R_{ij} = R_jR_i^{-1} = SL_{ij}S^{-1}$. If these two conditions are satisfied, then the relative rotation estimates R_{ij} and L_{ij} are consistent. In general they will not be, so we need to adjust them by a process of rotation averaging.

The cost function that we minimize is the residual error in all the rotation measurements R_{ij} and L_{ij} , defined by

$$\min_{S, L_i} \sum_{(i,j) \in \mathcal{N}} d(L_{ij}, L_jL_i^{-1})^p + d(R_{ij}, SL_jL_i^{-1}S^{-1})^p \quad (4)$$

There seems to be no direct method of minimizing this cost function under any of the distances we considered. Therefore, our strategy is to minimize the cost function by using rotation

averaging to update each L_i in turn, then conjugate rotation averaging to find S .

Conjugate rotation averaging here refers to a different rotation averaging problem than the one described above. Given n pairs of rotations (R_i, L_i) , the conjugate averaging problem is to find a rotation S such that $R_i \approx S^{-1}L_iS$ for all i . The appropriate minimization problem is then to find

$$\operatorname{argmin}_S \sum_{i=1}^n d(R_i, S^{-1}L_iS)^p.$$

An interpretation of this problem is that we are trying to find the fixed relative rotation S of two rigidly connected coordinate frames for which several corresponding rotation estimates L_i (for the “left” frame) and R_i (for the “right” frame) are available. This problem is sometimes referred to as the *hand-eye coordination problem*. See [2] for further details.

Initial values for each L_i are easily found by propagating from a given rotation L_0 assumed to be the identity, and then obtaining the initial S through conjugate averaging. Obviously, there is a gauge freedom in the solution, in that if S is any rotation, then the sets of rotations $R_i, i = 1, \dots, n$ and $R_iS, i = 1, \dots, n$ give the same result. This can be fixed by specifying one of the rotations to be the identity matrix. Our algorithm will not only obtain the camera rig rotation but also the rotations L_i and R_i for the left and right camera at each frame. The complete rotation estimation procedure follows (Algorithm 2). At each step of this algorithm, the total cost decreases, and hence converges to a limit. We do not at present claim a rigorous proof that the algorithm converges to even a local minimum, though that seems likely under most reasonable conditions. In particular, the sequence of estimates must contain a convergent subsequence, and the limit of this subsequence must be at least a local minimum with respect to each L_i and S individually.

- 1: Choose a tolerance $\epsilon > 0$.
 - 2: Let $L_0 = I$, determine L_i as the mean of all the rotations $L_{ij}^{-1}L_j$ where L_j has already been set.
 - 3: Estimate S from the equation $R_{ij}S = SL_{ij}$ as a conjugate rotation averaging problem.
 - 4: **loop**
 - 5: Update each L_j in turn by averaging all $L_{ij}L_i$, $S^{-1}R_{ij}SL_i$.
 - 6: Recompute and update S from the equation $R_{ij}S = SL_jL_i^{-1}$.
 - 7: **if** the cost function has decreased by less than ϵ since the last iteration **then**
 - 8: **return** S
 - 9: **end if**
 - 10: **end loop**
- Algorithm 2:** Iterative Rotation Averaging

C. Experiments

To evaluate the performance of the proposed algorithms, we conducted experiments on both synthetic data and real images. As a real example of a two-camera rig system, we have used a pair of wide-angle cameras to capture sequences of images. Images are captured at each camera illustrated in Figure 4. Feature points on the images are extracted using SIFT [9] and tracked through image sequences. These tracked features are transformed to image vectors on the unit sphere given the individual intrinsic calibrations. Outliers in the tracked features are removed using RANSAC [3] to fit the essential matrix using the normalized 8 point algorithm. Pairwise relative pose is obtained through decomposition of the essential matrix, and two frames bundle adjustment is utilized to refine the estimate, thus obtaining the relative rotations L_{ij} and R_{ij} . Finally, Algorithm 2 is applied to calibrate the camera rig, obtaining the relative rotation S and relative translation t .




Fig. 4. Images captured by camera rig with non-overlapping view

The image sequences captured by the left camera and the right camera contain 200 frames individually. As some pairs of image frames do not supply a successful relative motion estimation, we ultimately obtained 11199 pairs of relative motion estimates. Since both relative motion estimates L_{ij} and R_{ij} should have equal angle rotations, we use this criterion along with a minimum rotation angle requirement to select the best image pairs for further processing. Thus, we obtained 176 pairs of synchronized motions.

The convergence process is shown in Figure 5 where we illustrate the relation between $\log(\text{cost(it)} - \text{cost(e)})$ and iteration number it. Here, cost(e) is the convergence value.




Fig. 5. Convergence process on real camera rig image sequences. (a) Angular L^1 . (b) Angular L^2 . (c) Quaternion L^1 . (d) Quaternion L^2 . (e) Chordal L^1 . (f) Chordal L^2 .

VI. CONCLUSIONS

We have introduced a variant of the concept of geodesic convexity called weak convexity and applied it to the discussion of distance measures on $SO(3)$ and derived cost functions such as the one appearing in the definition of the Karcher mean. Our main contribution is a new and strong theorem on the location of global minima of such cost functions. This result significantly strengthens previous results that only considered local minima. We discussed an application of rotation averaging to camera rig calibration, an interesting problem in computer vision.

VII. ACKNOWLEDGMENTS

The authors would like to thank H. Li and N. Barnes for their help with the experiments described in this paper.

REFERENCES

- [1] Corcuera, J.M., Kendall, W.S.: Riemannian barycentres and geodesic convexity. *Math. Proc. Camb. Phil. Soc.* **127** (1999) 253–269
- [2] Dai, Y., Trunpf, J., Li, H., Barnes, N., Hartley, R.: Rotation averaging with application to camera-rig calibration. *Proceedings of the ninth Asian Conference on Computer Vision* (2009) to appear
- [3] Fischler, M. A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with application to image analysis and automated cartography. *Commun. Assoc. Comp. Mach.* **24** (1981) 381–395
- [4] Fletcher, P.T., Venkatasubramanian, S., Joshi, S.: The geometric median on riemannian manifolds with application to robust atlas estimation. *NeuroImage* **45**(1, Supplement 1) (2009) S143 – S152
- [5] Gramkow, C.: On averaging rotations. *Int. J. Comput. Vision* **42**(1-2) (2001) 7–16
- [6] Grove, K., Karcher, H., Ruh, E. A.: Jacobi fields and Finsler metrics on compact Lie groups with an application to differentiable pinching problems. *Math. Ann.* **211** (1974) 7–21
- [7] Karcher, H.: Riemannian center of mass and mollifier smoothing. *Comm. Pure Appl. Math.* **30**(5) (1977) 509–541
- [8] Krakowski, K., Hüper, K., Manton, J.: On the computation of the karcher mean on spheres and special orthogonal groups. In: *RoboMat 2007, Workshop on Robotics and Mathematics, Coimbra, Portugal* (2007)
- [9] Lowe, D. G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision* **60**(2) (2004) 91–110

- [10] Manton, J.H.: A globally convergent numerical algorithm for computing the centre of mass on compact Lie groups. In: Proceedings of the Eighth International Conference on Control, Automation, Robotics and Vision, Kunming, China (December 2004) 2211–2216
- [11] Moakher, M.: Means and averaging in the group of rotations. *SIAM J. Matrix Anal. Appl.* **24**(1) (2002) 1–16 (electronic)
- [12] Pennec, X.: Computing the mean of geometric features: application to the mean rotation. Technical Report INRIA RR-3371, INRIA (1998)
- [13] Sarlette, A., Sepulchre, R.: Consensus optimization on manifolds. *SIAM J. Control Optim.* **48**(1) (2009) 56–76