Int J Comput Vis (2013) 103:267–305
DOI 10.1007/s11263-012-0601-0

Rotation Averaging

Richard Hartley · Jochen Trumpf · Yuchao Dai · Hongdong Li

Received: 30 January 2012 / Accepted: 28 November 2012 / Published online: 30 January 2013

Abstract This paper is conceived as a tutorial on rotation averaging, summarizing the research that has been carried out in this area; it discusses methods for single-view and multiple-view rotation averaging, as well as providing proofs of convergence and convexity in many cases. However, at the same time it contains many new results, which were developed to fill gaps in knowledge, answering fundamental questions such as radius of convergence of the algorithms, and existence of local minima. These matters, or even proofs of correctness have in many cases not been considered in the Computer Vision literature. We consider three main problems: single rotation averaging, in which a single rotation is computed starting from several measurements; multiple-rotation averaging, in which absolute orientations are computed from several relative orientation measurements; and conjugate rotation averaging, which relates a pair of coordinate frames. This last is related to the hand-eye coordination problem and to multiple-camera calibration.

Keywords Geodesic distance · Angular distance · Chordal distance · Quaternion distance · L_1 mean · L_2 mean · conjugate rotation

1 Introduction

In this paper, we will be interested in three different rotation averaging problems. In the following description, $d(\mathtt{R}, \mathtt{S})$ denotes the distance between two rotations $\mathtt{R}$ and $\mathtt{S}$. Various different possible distance functions will be described later in the paper; for now, $d(\cdot, \cdot)$ is thought of as being any arbitrary metric on the space of rotations SO(3).

Single Rotation Averaging. In the single rotation averaging problem, several estimates are obtained of a single rotation, which are then averaged to give the best estimate. This may be thought of as finding a mean of several points $\mathtt{R}_i$ in the rotation space SO(3) (the group of all 3-dimensional rotations) and is an instance of finding a mean in a manifold.

Given an exponent $p \geq 1$ and a set of $n \geq 1$ rotations $\{\mathtt{R}_1, \ldots, \mathtt{R}_n\} \subset \mathrm{SO}(3)$ we wish to find the L^p-mean rotation with respect to d which is defined as

$$d^{\mathrm{p}} - \mathrm{mean}(\{\mathtt{R}_1, \ldots, \mathtt{R}_n\}) = \operatorname*{argmin}_{\mathtt{R}\in\mathrm{SO}(3)} \sum_{i=1}^{n} d(\mathtt{R}_i, \mathtt{R})^p .$$

Since SO(3) is compact, a minimum will exist as long as the distance function is continuous (which any sensible distance function is). This problem has been much studied in the literature, but there are still open problems, some of which are resolved here.

Conjugate Rotation Averaging. In the conjugate rotation averaging problem, $n \geq 1$ rotation pairs $(\mathtt{L}_i, \mathtt{R}_i)$ (the left and right rotations) are given, and we need to find a rotation $\mathtt{S}$ such that $\mathtt{R}_i = \mathtt{S}^{-1}\mathtt{L}_i\mathtt{S}$ for all i. This problem arises when the rotations $\mathtt{R}_i$ and $\mathtt{L}_i$ are measured in different coordinate frames, and the coordinate transformation $\mathtt{S}$ that relates these two frames is to be determined.

In the presence of noise, the appropriate minimization problem is then to find

$$\operatorname*{argmin}_{\mathtt{S}} \sum_{i=1}^{n} d(\mathtt{R}_i, \mathtt{S}^{-1}\mathtt{L}_i\mathtt{S})^p .$$

This problem is sometimes referred to as the *hand-eye coordination problem*, see for example Daniilidis (1998), Park and Martin (1994), and Zhang (1998).

R. Hartley (✉) · J. Trumpf · Y. Dai · H. Li
Australian National University, Canberra, ACT, Australia
e-mail: Richard.Hartley@anu.edu.au