

The International Journal of Robotics Research

<http://ijr.sagepub.com/>

Data Selection for Hand-eye Calibration: A Vector Quantization Approach

Jochen Schmidt and Heinrich Niemann

The International Journal of Robotics Research 2008 27: 1027

DOI: 10.1177/0278364908095172

The online version of this article can be found at:

<http://ijr.sagepub.com/content/27/9/1027>

Published by:



<http://www.sagepublications.com>

On behalf of:



Multimedia Archives

Additional services and information for *The International Journal of Robotics Research* can be found at:

Email Alerts: <http://ijr.sagepub.com/cgi/alerts>

Subscriptions: <http://ijr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://ijr.sagepub.com/content/27/9/1027.refs.html>

>> [Version of Record](#) - Aug 28, 2008

[What is This?](#)

Jochen Schmidt

Centre for Artificial Intelligence Research
Auckland University of Technology
Auckland, New Zealand
jochen.schmidt@aut.ac.nz

Heinrich Niemann

Lehrstuhl für Mustererkennung
Universität Erlangen-Nürnberg
91058 Erlangen, Germany

Data Selection for Hand-eye Calibration: A Vector Quantization Approach

Abstract

This paper presents new vector quantization based methods for selecting well-suited data for hand-eye calibration from a given sequence of hand and eye movements. Data selection can improve the accuracy of classic hand-eye calibration, and make it possible in the first place in situations where the standard approach of manually selecting positions is inconvenient or even impossible, especially when using continuously recorded data. A variety of methods is proposed, which differ from each other in the dimensionality of the vector quantization compared to the degrees of freedom of the rotation representation, and how the rotation angle is incorporated. The performance of the proposed vector quantization based data selection methods is evaluated using data obtained from a manually moved optical tracking system (hand) and an endoscopic camera (eye).

KEY WORDS—hand-eye calibration, vector quantization

1. Introduction

Hand-eye calibration, which has classically been used for calibrating the rigid transformation from the tip of a robot manipulator arm to a camera mounted on the arm, becomes more interesting for applications where similar problems arise, but which are not directly related to robotics. Examples include hand data provided by an optical tracking system instead of a robot, where the camera is mounted on an endoscope and moved manually as in Vogt (2006) or Schmidt et al. (2004)

or an inertial sensor mounted on a camera as in Aron et al. (2004). Self-calibration of a rigid stereo-camera system, where one camera can be treated as the ‘hand’ and the other as the ‘eye’, is described in Luong and Faugeras (2001) and Schmidt (2006).

A problem that is common to all hand-eye calibration algorithms is that the quality of the result is highly dependent upon the data used for computing the unknown transformation. The usual approach for solving this problem is to use robot movements that already take the restrictions on the data into account, which means that the movement has to be planned before recording. Suggestions of how this can be achieved were given in one of the original publications on hand-eye calibration by Tsai and Lenz (1989). In situations where planning such a well-suited movement is not possible (e.g. due to constraints on the available space) or cannot be controlled well (e.g. when using a hand-held camera), methods for data selection are required in order to obtain high-quality calibration results. In particular, in situations where the camera records images at frame rates of 25 images per second while the camera is moving continuously, using the movements in temporal order is a bad choice. Performing a data selection step before the actual hand-eye calibration is therefore essential in these cases.

We present a variety of new methods for automatic selection of well-suited data based on vector quantization. Using any of them before the actual hand-eye calibration makes calibration of continuously recorded data possible in the first place. As we will show in the experiments section, even when only a small number of planned poses are used as in the classic approach to acquiring data, running the proposed data selection algorithms can lead to an increase in accuracy. The algorithms presented here have been developed by Schmidt (2006). They have been partially revised, making the automatic threshold computation (Section 3.5) more effective. New experimental results will be presented in this article.

The problem of automatic data selection was already addressed in Schmidt et al. (2003), where it is applied in hand-eye calibration of an endoscopic surgery robot. This method first removes relative movements (defined by the rotation and translation between two given poses; see also Section 3.1) with small rotation angles using a fixed threshold. After this pre-selection step, *pairs* of relative movements are rated according to their suitability for hand-eye calibration. The goal is to use the best fraction of pairs for computing the hand-eye transformation.

As a rating criterion, it is proposed to use the scalar product between the rotation axes of two relative movements. A worst case estimate (if no movements are eliminated during pre-selection w.r.t. angle) of the time complexity of this approach is $O(N^4)$, N being the number of poses contained in the original input sequence. The method is more or less a brute-force approach; it cannot compete in computation time (or accuracy) with the vector quantization methods that we present in this article. Another problem is that well-matching *pairs* of relative movements are always selected, where one relative movement may be contained in multiple pairs. The pairs are afterwards used to form a linear system of equations for solving the hand-eye transformation. Since each relative movement results in one equation, it may happen that one movement is used more than once leading to two linearly dependent equations, one of them being redundant.

We have presented the first vector quantization based data selection method in Schmidt et al. (2004). This method is also described in Section 3.3. The pre-selection of movements using a fixed threshold has been adopted in Schmidt et al. (2004), while an algorithm for computing a suitable threshold automatically is now used. Schmidt et al. (2004) also contains an experimental comparison with the approach published in Schmidt et al. (2003); it concludes that the latter cannot compete with the vector quantization method at all and that the calibration results are mostly worse than for the vector quantization based methods.

Zhang et al. (2005) and Shi et al. (2005) have also presented an iterative method for data selection in hand-eye calibration. Their method considers pairs of movements and rates their suitability for hand-eye calibration by comparing the angle between the two rotation axes, the rotation angles and the norm of the translation vectors to thresholds. In contrast, our vector quantization based approach does not only rate pairs but also takes into account all available movements on a global scale.

The paper is structured as follows. A short introduction of hand-eye calibration will be given in Section 2, including a brief literature review and explaining the critical factors that influence calibration accuracy. Section 3 describes the new proposed data selection algorithms based on vector quantization. Experimental results, including a comparison with classic hand-eye calibration without data selection, are presented in Section 4. The paper concludes with Section 5. As the proposed algorithms make excessive use of different representa-

tions of three-dimensional (3D) rotations, we have included a short introduction for readers not familiar with these representations in Appendix A.

2. Hand-eye Calibration

2.1. Overview

The hand-eye calibration problem can be formulated as follows. Given a robot manipulator arm and a camera mounted on that arm, compute the rigid transformation from arm to camera, also called *hand-eye transformation*. Knowledge of this transformation is necessary for computing the camera pose from the pose of the robot arm, which is usually provided by the robot itself, while the pose of the camera is unknown. Once the hand-eye transformation is known, the camera pose is easily computed from the pose information provided by the robot.

The first hand-eye calibration methods were published by Shiu and Ahmad (1989) and Tsai and Lenz (1989); an early comparison of the methods available at that time was given in Wang (1992). The hand-eye calibration problem was formulated by Shiu and Ahmad (1989) as a matrix equation of the form

$$T_E T_{HE} = T_{HE} T_H, \quad (1)$$

where T_H is the robot arm (*hand*) movement, T_E the camera (*eye*) movement, and T_{HE} is the unknown hand-eye transformation, i.e. the transformation from gripper to camera. (Note that in some publications, T_{HE} is the transformation from camera to gripper.) Each matrix T_χ is a rigid transformation in homogeneous form, i.e.

$$T_\chi = \begin{pmatrix} R_\chi & t_\chi \\ \mathbf{0}_3^T & 1 \end{pmatrix}, \quad \chi \in \{H, E, HE\}. \quad (2)$$

The transformations T_χ consist of a 3×3 rotation matrix R_χ and a 3D translation vector t_χ ; $\mathbf{0}_3$ denotes the 3D null-vector.

The straightforward method of solving Equation (1) is to split it into two separate equations: one that contains only rotation and a second that contains rotation and translation i.e.

$$R_E R_{HE} = R_{HE} R_H \quad (3)$$

$$(R_E - I_{3 \times 3})t_{HE} = R_{HE}t_H - t_E, \quad (4)$$

where $I_{3 \times 3}$ denotes the 3×3 identity matrix. Thus, the rotational part R_{HE} of the hand-eye transformation can be determined first from Equation (3) and, after inserting it into Equation (4), the translational part t_{HE} is computed.

This is how hand-eye calibration is carried out by Shiu and Ahmad (1989); Tsai and Lenz (1989); Chou and Kamel (1991); Wang (1992). Various parameterizations of rotation were applied (see Appendix A for an introduction). The original works of Shiu and Ahmad (1989) and Tsai and Lenz (1989)

use the axis/angle representation; quaternions were used by Chou and Kamel (1991) and Horaud and Dornaika (1995). Dual quaternions were introduced by Daniilidis (1999, 2001).

In contrast to the former approaches, it was suggested by Chen (1991) that rotation and translation should be solved simultaneously, not separately. This approach is also followed by Horaud and Dornaika (1995), where a nonlinear optimization of rotation and translation is performed. An online calibration method based on Kalman filtering was proposed by Andreff (1997). Daniilidis (1999, 2001) introduced a hand-eye calibration algorithm based on dual quaternions that is also capable of handling rotation and translation simultaneously. However, in contrast to the former approaches, a linear solution is given. A method based on screw motions, very similar to Daniilidis' dual quaternion algorithm, can be found in Zhao and Liu (2006). Fassi and Legnani (2005) present a geometric interpretation of the original hand-eye calibration equation and analyze its properties.

Commonly, hand-eye calibration methods rely on the fact that the movement of the robot manipulator arm is provided by the robot, while the camera movement is computed using a calibration pattern and classic camera calibration methods (e.g. Tsai 1987; Zhang 1998, 2000). Andreff et al. (1999, 2001) and Schmidt et al. (2005) presented approaches that obtain the camera movement not by using a calibration pattern, but from point feature tracking and a structure-from-motion technique (see Faugeras and Luong (2001); Hartley and Zisserman (2003) for an introduction). In this case an additional scale factor has to be estimated, making the problem very similar to the self-calibration of a rigid stereo-camera system (Luong and Faugeras 1993, 2001; Faugeras and Luong 2001; Dornaika and Chung 2003; Schmidt 2006).

2.2. Critical Factors

Regardless of which algorithm is actually used, one important constraint is always valid for solving the general hand-eye calibration problem. At least two movements of the robot manipulator are necessary, where the axes of the rotations are non-parallel. This was shown by Tsai and Lenz (1989) and Chen (1991). If the movement is not general enough, the hand-eye parameters can be recovered only partially. For details, see e.g. Andreff et al. (2001).

The critical factors and criteria for improving hand-eye calibration accuracy have already been given by Tsai and Lenz (1989).

1. Maximize the angle between rotation axes of relative movements (influence on error in rotation, no translation recovery possible for parallel axes).
2. Maximize the rotation angle of relative movements (influence on error in rotation and translation).

3. Minimize the distance between the optical center of the camera and the calibration pattern (influence on error in translation).
4. Minimize the distance between the gripper coordinate system positions, i.e. small translational movement of the hand (influence on error in translation).

More details as well as an error analysis can be found in Tsai and Lenz (1989). If the movement of the robot gripper can be planned in advance, all items above may be controlled by the user. The usual way to fulfill the data requirements in robot hand-eye calibration is to use a calibration setup where the different positions of the gripper are chosen such that the data is well suited for calibration. Such a setup is described in e.g. Tsai and Lenz (1989).

When planning movements is not possible or gripper movement is confined to certain areas, not all of the above criteria can be controlled as desired. Particularly when hand-eye calibration is done for devices other than robot manipulator arms, controlled movements may be not possible at all or only with very low accuracy. An example of such an application that becomes more popular is the calibration of an optical tracking system (basically a camera) and an endoscope (optics and endoscopic camera). A so-called target, which consists of retroreflective markers, is mounted on the endoscope. Its 3D pose can then be determined by the optical tracking system. In this setup, the hand data is provided by the tracking system and the eye data by the endoscopic camera. With a known hand-eye transformation, the movement of the endoscopic camera can therefore be computed from the movement of the target. The endoscope is normally moved manually, and its exact pose cannot be controlled. For more details see Schmidt et al. (2004); Vogt (2006).

Another application is self-calibration of a stereo-camera system as discussed in Schmidt (2006). In this case, the rigid transformation between two cameras is estimated based on image information only, without using a calibration pattern. The cameras are either moved manually or implicitly by the head movement of a user in an Augmented Reality setting. Controlling criterion (3) in such a setup is usually not feasible. It would mean minimizing the distance of the cameras to the observed scene, which would restrict the freedom of the user's movements considerably, rendering one of the main advantages of a self-calibration approach useless.

However, the main difference between these new application areas and classic hand-eye calibration is the way data is recorded. In a traditional setup, images of a calibration pattern are acquired at a comparatively small number of positions, say 20. In an Augmented Reality setting, where the user wearing a stereo camera is allowed to move freely, images are recorded at a rate of 25 or 30 frames per second. This results in an amount of data higher by one to two orders of magnitude, recorded at more or less arbitrary positions. Also, having a continuous movement where rotation and translation change only slightly

between two consecutive frames, a trade-off between requirements (1) and (2) and requirement (4) has to be taken into account.

3. Data Selection

Various data selection algorithms are presented which select a globally consistent set of relative movements, optimizing the non-parallelism criterion (1) and the rotation angle criterion (2) from Section 2.2. Criterion (1) is considered to be the most important one here, since no recovery of the hand-eye translation is possible if the rotation axes used are parallel.

3.1. Pre-processing

The main purpose of the algorithms presented in this paper is to increase the hand-eye calibration accuracy when the hand positions are recorded continuously (i.e. when consecutive poses differ only minimally). Given the criteria from Section 2.2, it can be expected that processing the data in temporal ordering is suboptimal. This is supported by the experimental results presented later on; using these readily available movements between consecutive poses therefore cannot be recommended.

Deciding which data, i.e. relative movements, should be used as input for the data selection method instead is an important step that has considerable influence on the performance of the algorithm. It has been proposed by Schmidt et al. (2003) to consider *all possible* relative movements between the recorded hand poses that are contained in the data and use these as input for the data selection algorithm. Given two robot gripper poses at time i and j consisting of rotation $\mathbf{R}_i, \mathbf{R}_j$ and translation $\mathbf{t}_i, \mathbf{t}_j$, the relative movement given by $\mathbf{R}_{ij}, \mathbf{t}_{ij}$ can be computed as

$$\mathbf{R}_{ij} = \mathbf{R}_j^T \mathbf{R}_i \quad (5)$$

for rotation and

$$\mathbf{t}_{ij} = \mathbf{R}_j^T (\mathbf{t}_i - \mathbf{t}_j) = \mathbf{R}_{ij} \mathbf{R}_i^T (\mathbf{t}_i - \mathbf{t}_j) \quad (6)$$

for translation.

For N poses, the total number of all relative movements is $N(N-1)/2$, i.e. the time complexity of the pre-processing step equals $O(N^2)$.

3.2. Vector Quantization – Overview

Given a set of N_r relative movements represented by their rotation axes, the following algorithm aims to compute a new set of distinct axes consisting of N_s vectors where $N_s < N_r$. This is achieved by running a clustering algorithm on the vectors

representing axes, which computes a partitioning of the axes vectors.

A method which is suited very well to this task is *vector quantization* (Linde et al. 1980). Note that although we propose to use vector quantization (and the LBG algorithm – named after the authors Linde, Buzo and Gray (Linde et al. 1980) – mentioned later), this should be seen as an example rather than a strict condition for data selection. In fact, *any* clustering algorithm can be applied.

In general, vector quantization works as follows. An input vector $\mathbf{x} \in \mathbb{R}^n$ is mapped to a vector of the so-called *codebook* $\mathcal{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_{N_s}\}$, which is a set of N_s n -dimensional vectors that define a partitioning of \mathbb{R}^n . Given a distance measure $d(\cdot, \cdot)$ on vectors in \mathbb{R}^n (usually Euclidean distance), the input vectors are mapped:

$$\mathbf{x} \mapsto \mathbf{c}_\kappa, \quad \text{where } d(\mathbf{x}, \mathbf{c}_\kappa) < d(\mathbf{x}, \mathbf{c}_i)$$

$$\forall i = 1, \dots, N_s, i \neq \kappa. \quad (7)$$

Thus, the entries of the codebook \mathcal{C} are the cluster centers in \mathbb{R}^n . For finding the entries of the codebook, the well-known LBG algorithm is used. This is an iterative method that computes the codebook given the desired number of codebook entries. The complexity of the LBG algorithm for each iteration is $O(N_r N_s)$, which equals $O(N^2 N_s)$. Note that the number of movements N_r is reduced further considerably by a pre-selection step removing movements with small rotation angles as described in Section 3.3.

Various algorithms based on this idea are presented in Section 3.3. They differ in the input data used as well as in the dimensionality of vector quantization.

3.3. Three-dimensional Vector Quantization of Normalized Rotation Axes

The vector quantization based data selection algorithm presented in this section requires 3D rotation axes \mathbf{r}_i having norm one computed from relative movements as input data. This algorithm was first published in Schmidt et al. (2004).

Before selecting the movements according to their non-parallelism, a pre-selection is done according to their rotation angle since for angles close to zero the rotation axis is not well-defined (see Appendix A.2). For an angle of 180° , singularities in hand-eye calibration arise (Shiu and Ahmad 1989; Daniilidis 1999). These are exactly the cases where the rotation matrix has multiple real eigenvalues. This step removes these movements, thus optimizing criterion (2). Movements are discarded that have rotation angles greater than a given threshold θ_t and less than $180^\circ - \theta_t$, or higher than $180^\circ + \theta_t$ and less than $360^\circ - \theta_t$. The second interval is due to the fact that a rotation about an axis \mathbf{r} by an angle θ is the same as a rotation about the axis $-\mathbf{r}$ by the angle $360^\circ - \theta$. After the pre-selection step,

Input:	
N_r relative movements consisting of rotation and translation $\mathbf{R}_i, \mathbf{t}_i$ (Section 3.1),	
θ_t = threshold for pre-selection according to rotation angle,	
N_s = number of desired relative movements after data selection is complete (codebook size).	
Output:	
N_s relative movements consisting of rotation and translation $\mathbf{R}_\kappa, \mathbf{t}_\kappa$.	
FOR each relative movement i	
Compute axis \mathbf{r}_i (norm one) and angle θ_i from \mathbf{R}_i	
IF	$ \theta_i < \theta_t$ OR $(\theta_i > 180^\circ - \theta_t$ AND $ \theta_i < 180^\circ + \theta_t)$ OR $ \theta_i > 360^\circ - \theta_t$
THEN	Rotation angle too small: remove movement i from dataset
ELSE	Resolve ambiguities (restrict to a single hemisphere)
Compute codebook $\mathcal{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_{N_s}\}$ of size N_s using the remaining \mathbf{r}_i as training vectors	
FOR each remaining axis \mathbf{r}_i	
Classify \mathbf{r}_i to one of the partitions represented by codebook vector \mathbf{c}_κ : $\mathbf{r}_i \rightarrow \mathbf{r}_{i,\kappa}$	
Compute the distance $d(\mathbf{r}_{i,\kappa}, \mathbf{c}_\kappa)$	
FOR each codebook entry \mathbf{c}_κ	
Determine $\mathbf{r}_\kappa = \mathbf{r}_{j,\kappa}$, where $d(\mathbf{r}_{j,\kappa}, \mathbf{c}_\kappa) < d(\mathbf{r}_{i,\kappa}, \mathbf{c}_\kappa) \forall i, j$ of partition κ , $i \neq j$	
Select the relative movement $\mathbf{R}_\kappa, \mathbf{t}_\kappa$ that corresponds to \mathbf{r}_κ as one of the resulting movements	

Fig. 1. Data selection using a 3D vector quantization of normalized rotation axes.


only the rotation axes (normalized to one) are used for further processing. The complete algorithm is shown in Figure 1.

After pre-selection according to rotation angle and normalization of the rotation axes to one, the ambiguity in the axis/angle representation is resolved to ensure that similar rotation axes are actually close to each other in 3D. Since all normalized rotation axes $\mathbf{r}_i = (r_{ix} \ r_{iy} \ r_{iz})^T$ lie on a sphere in 3D space, this can be achieved by restricting the axes to one hemisphere. The hemisphere with non-negative r_{iz} -coordinate was chosen without loss of generality. If this coordinate of an axis \mathbf{r}_i is negative, the axis \mathbf{r}_i is substituted by $-\mathbf{r}_i$. Rotation axes having a zero r_{iz} -coordinate have to be handled separately by checking the r_{iy} - and r_{ix} -coordinates.


The next step is training the vector quantizer i.e. computation of the codebook vectors \mathcal{C} which results in a clustering of the rotation axes. Note that since all axes have norm one, the vectors are not uniformly distributed in space but lie on the surface of the unit sphere. An example obtained from real data is depicted in Figure 2. For the dataset shown in Figure 2(a), no pre-selection of the data with respect to small rotation angles was done, i.e. all relative movements were used as input for vector quantization. The resulting codebook is plotted in

Figure 2(b). Figure 2(c) shows the same dataset, where relative movements having a rotation angle smaller than 15° have been removed, and the generated codebook is depicted in Figure 2(d).


In many algorithms that apply vector quantization, the codebook vectors can be used directly for further processing; note that this is not the case for data selection as described here. Codebook vectors are computed as the center of gravity (i.e. mean values) of all input vectors belonging to a certain partition. Therefore, a codebook vector does not usually coincide with an element of the input vector set, which means that it cannot be related to an actual relative movement. Additional steps have to be taken in order to obtain a single rotation axis (and the associated relative movement) per partition. Firstly, each rotation axis \mathbf{r}_i has to be classified to one of the partitions defined by the codebook vectors. The classified axes are denoted by $\mathbf{r}_{i,\kappa}$. Secondly, for each rotation axis $\mathbf{r}_{i,\kappa}$ of a partition κ , the distance to the codebook vector \mathbf{c}_κ representing that partition is computed; the selected axis is the one where the distance to the codebook vector $d(\mathbf{r}_{i,\kappa}, \mathbf{c}_\kappa)$ is smallest. The relative movements belonging to the rotation axes selected this way can now be used for hand-eye calibration.




(a) input data, no threshold on rotation angle, i.e., $\theta_t = 0^\circ$



(b) codebook, no threshold on rotation angle, i.e., $\theta_t = 0^\circ$



(c) input data, threshold of $\theta_t = 15^\circ$



(d) codebook, threshold of $\theta_t = 15^\circ$

Fig. 2. Examples of vector quantization results in 3D based on real data: (a) input data, no threshold on rotation angle, i.e. $\theta_t = 0^\circ$; (b) codebook, no threshold on rotation angle, i.e. $\theta_t = 0^\circ$; (c) input data, threshold of $\theta_t = 15^\circ$; and (d) codebook, threshold of $\theta_t = 15^\circ$. The input data (rotation axes of norm one) in (a), (c) are shown as small dots and the codebook are plotted in (b), (d) as bold dots. Because of the normalization of the axes to one, all vectors lie on a sphere.

3.4. Two-dimensional Vector Quantization of Normalized Rotation Axes

The data selection algorithm presented before uses normalized rotation axes as input where the rotation angle is not encoded in the axis but handled separately, an axis being a 3D vector with only two degrees of freedom. Hence, the dimensionality of the vector quantization can be reduced from three to two by

using an appropriate parameterization of the axes. An obvious choice for this task is polar coordinates. Given a rotation axis \mathbf{r} , the polar coordinates λ, ρ of \mathbf{r} are computed as:

$$\lambda = \arctan \frac{r_y}{r_x}, \quad \rho = \arcsin r_z. \quad (8)$$

The data selection algorithm using polar coordinates is similar to that shown in Figure 1. The main difference is that an ad-




Fig. 3. Examples of vector quantization result using polar coordinates based on real data: (a) no threshold on rotation angle, i.e. $\theta_t = 0^\circ$ and (b) threshold of $\theta_t = 15^\circ$. The angles λ and ρ are given in rad.

ditional computation step, preceding the codebook generation, has to be introduced to convert the normalized rotation axes \mathbf{r}_i to polar coordinates using Equation (8). This reduces the dimensionality of the following vector quantization, as only 2D vectors that contain the angles λ_i and ρ_i are used.

An example of a vector quantization result using polar coordinates is shown in Figure 3. The same dataset as in Figure 2 was used.

3.5. Automatic Computation of Rotation Angle Thresholds

A drawback of the two data selection algorithms presented in Sections 3.3 and 3.4 is that a threshold θ_t for the rotation angle has to be set manually. There are mainly two ways to improve the algorithm.

1. Use a threshold, but compute it automatically from the available data.
2. Do not use a threshold at all, i.e. all data are used for vector quantization. The rotation angle is taken into account implicitly by an appropriate parameterization.

The former is discussed here; algorithms using the second option are shown in Sections 3.6 and 3.7.

In the data selection algorithms presented before a single threshold was used, which is applied to the lower (0° , 360°) and upper (180°) bounds of the rotation angle interval symmetrically. The best-suited movements have rotation angles located in the center of the two intervals at 90° and 270° . The algorithm for automatic threshold computation is different in two ways. Two separate thresholds for the upper and lower

bound are calculated, and these are not necessarily symmetric. The structure chart for automatic threshold computation is shown in Figure 4.

The algorithm requires that the desired remaining fraction of movements be specified as an input parameter. This is a clear advantage over using a threshold for the angle directly. When an explicit angle threshold is provided, more often than not it will be chosen either too high or too low, as we have no *a priori* knowledge of the data used for calibration. If the threshold is too high, there may be no movement left after applying the threshold which makes calibration impossible. If chosen too low, lots of movements are possibly processed further that are actually not very good and thus distort the calibration result. Since the rotation angles are highly dependent on the recorded image sequence, a general recommendation for choosing a threshold on the angle is not possible. In contrast, specifying a percentage of movements that are to remain after pre-selection, allows the best movements to be selected while it is guaranteed that a sufficient number of movements is left for calibration.

First, the rotation angle θ_i is computed for each relative movement; the rotation axes are irrelevant for threshold determination. Since only the amount of rotation is of interest we take the absolute values of θ_i . Then, all angles are normalized to the interval 0° to 180° and stored in a list \mathcal{L} which can be accessed by an index ranging from 0 to $N_r - 1$. After sorting the list in ascending order, the smallest and largest rotation angles α_1 and α_u that are contained in the recorded image sequence can be found in the entries $\mathcal{L}(0)$ and $\mathcal{L}(N_r - 1)$.

Recall that after normalization, the best-suited movements have a rotation angle of $\sim 90^\circ$. Three cases have to be distinguished. If the smallest rotation angle α_1 is greater than 90° ,

Input:	
N_r relative movements consisting of rotation and translation $\mathbf{R}_i, \mathbf{t}_i$ (Section 3.1),	
δ = desired remaining fraction (0 to 1) of relative movements after pre-selection.	
Output:	
θ_u, θ_l = upper and lower threshold for pre-selection according to rotation angle.	
FOR each relative movement i	
Compute rotation angle θ_i from \mathbf{R}_i	
Take absolute angle, i.e. $\theta_i := \theta_i $	
IF	$\theta_i > 180^\circ$
THEN	Normalize angle to interval 0° to 180° : $\theta_i := 360^\circ - \theta_i$
Store angle θ_i in a list \mathcal{L} (indexed from $0, \dots, N_r - 1$)	
Sort list \mathcal{L} (ascending)	
$\alpha_l := \mathcal{L}(0), \alpha_u := \mathcal{L}(N_r - 1)$	
IF	$\alpha_l > 90^\circ$
THEN	$\theta_l := 0^\circ, \theta_u := 180^\circ - \mathcal{L}((N_r - 1) - \text{Round}((1 - \delta) \cdot (N_r + 1)))$
ELSE	IF $\alpha_u < 90^\circ$
THEN	$\theta_l := \mathcal{L}(\text{Round}((1 - \delta) \cdot (N_r - 1))), \theta_u := 0^\circ$
ELSE	IF $\alpha_l \leq 90^\circ$ AND $\alpha_u \geq 90^\circ$
THEN	Find index i , where $\mathcal{L}(i) \geq 90^\circ$ and $\mathcal{L}(j) < 90^\circ \forall j = 0, \dots, i - 1$
	Compute fractions ξ_l, ξ_u of list elements left and right of i (cf. (9))
	Compute fractions ρ_s, ρ_b to be removed on left and right (cf. (10))
	IF $\xi_l \geq \xi_u$
	THEN $\theta_l := \mathcal{L}(\text{Round}(\rho_b(N_r - 1)))$
	$\theta_u := 180^\circ - \mathcal{L}((N_r - 1) - \text{Round}(\rho_s(N_r + 1)))$
	ELSE $\theta_l := \mathcal{L}(\text{Round}(\rho_s(N_r - 1)))$
	$\theta_u := 180^\circ - \mathcal{L}((N_r - 1) - \text{Round}(\rho_b(N_r + 1)))$

Fig. 4. An algorithm for computing upper and lower rotation angle thresholds θ_u, θ_l automatically.

no movements are contained in the sequence that have to be removed at the lower bound. Therefore, the lower threshold θ_l is set to 0° . All movements have to be removed at the upper bound, and the upper threshold is defined by the list entry having index $(N_r - 1) - \text{Round}((1 - \delta) \cdot (N_r + 1))$.

A similar situation with reversed roles of thresholds arises if the largest angle α_u is smaller than 90° . In this case, all movements have to be removed at the lower bound. The upper threshold θ_u is set to 0° , while the lower is given by the list entry at index $\text{Round}((1 - \delta) \cdot (N_r - 1))$.

The case where angles both below and above 90° exist is slightly more complex, as both thresholds have to be chosen asymmetrically depending on the number of movements with rotation angles greater and smaller than 90° . Therefore, the first step is to identify the position i of the first angle in the list that is equal to or larger than 90° , i.e. the ‘center’ of the list

with respect to the best angle contained in the sequence. The fraction of angles left and right of the identified position i can now be computed:

$$\xi_l = \frac{i}{N_r - 1}, \quad \xi_u = 1 - \xi_l, \quad (9)$$

where ξ_l, ξ_u are the fraction of movements in the lower and upper part with respect to 90° , respectively. When deciding which movements are to be removed, it is important to keep as many close to 90° as possible. In the majority of cases, most movements will be deleted in the part of the list containing the larger amount of movements and only a small fraction (if any) from the part containing the smaller amount of movements. We compute two factors ρ_s and ρ_b that determine the fraction of movements to be deleted in the smaller (ρ_s) and the larger (ρ_b) parts of the list. These factors are given by:

Input:	
N_r relative movements consisting of rotation and translation $\mathbf{R}_i, \mathbf{t}_i$ (Section 3.1),	
N_s = number of desired relative movements after data selection is complete.	
Output:	
N_s relative movements consisting of rotation and translation $\mathbf{R}_\kappa, \mathbf{t}_\kappa$.	
FOR each relative movement i	
Compute axis/angle ω_i from \mathbf{R}_i ; angle $\theta_i = \omega_i $	
IF	$\theta_i > 180^\circ$
THEN substitute ω_i by a rotation about the negative axis and angle $360^\circ - \theta_i$	
Compute codebook $\mathcal{C} = \{c_1, \dots, c_{N_s}\}$ of size N_s using the ω_i as training vectors	
FOR each vector ω_i	
Classify ω_i to one of the partitions represented by codebook vector c_κ :	
$\omega_i \rightarrow \omega_{i,\kappa}$	
Compute the distance $d(\omega_{i,\kappa}, c_\kappa)$	
FOR each codebook entry c_κ	
Determine $\omega_\kappa = \omega_{j,\kappa}$, where $d(\omega_{j,\kappa}, c_\kappa) < d(\omega_{i,\kappa}, c_\kappa)$	
$\forall i, j$ of partition $\kappa, i \neq j$	
Select the relative movement $\mathbf{R}_\kappa, \mathbf{t}_\kappa$ that corresponds to ω_κ as one of the resulting movements	

Fig. 5. Data selection using axis/angle representation.

$$\rho_s = \max\left(\frac{1}{2}((1 - \delta) - |\xi_l - \xi_u|), 0\right)$$

$$\rho_b = \min((1 - \delta), \max(\xi_l, \xi_u) - \min(\xi_l, \xi_u)) + \rho_s. \quad (10)$$

The first addend in the equation for ρ_b removes as many movements as possible from the larger part, i.e. either all that have been required or an amount that makes it the same size as the originally smaller part. The remaining movements can then be removed symmetrically from both parts, the amount being defined by ρ_s .

If most movements are located in the lower part of the list, the thresholds θ_l, θ_u are then defined by the list elements having index $\text{Round}(\rho_b(N_r - 1))$ and $(N_r - 1) - \text{Round}(\rho_s(N_r + 1))$, respectively. In the case that most movements are located in the upper part of the list, ρ_s and ρ_b change roles when computing the indices of the list elements.

Strictly speaking, checking the first two cases where all movements are located on one side of 90° is not necessary as both are covered by Equation (10). Nevertheless, they have been included in the structure chart as we believe the concept is easier to understand this way.

3.6. Vector Quantization using Axis/Angle Representation

This section presents a data selection algorithm that does not need thresholds on rotation angles as no pre-selection of rela-

tive movements is done, i.e. all available movements are used. A structure chart is shown in Figure 5.

Instead of treating rotation axis and angle separately, the axis/angle representation is used as described in Appendix A.2 where the angle θ is encoded as the norm of the axis vector \mathbf{r} in a 3D vector ω having three degrees of freedom.

We start with a set of N_r relative movements, now represented by their rotation axes with angles encoded in ω_i . The result is a set of N_s vectors, $N_s < N_r$, where the corresponding selected movements are a trade-off between criteria (1) and (2) as defined in Section 2.2. Movements having small rotation angles will be found in the resulting dataset if their rotation axes fit the remaining data well.

Initially, all ω_i are normalized such that angles are in the range 0 – 180° . This is different from the normalization in the previous data selection methods, where the sign of the axes elements is used. There are always two options: either the sign of an axis or the rotation angle is controlled, never both. Depending on the application, either option may have advantages and disadvantages.

In contrast to the methods presented before, where movements with rotation angles close to zero are scattered, they will now be concentrated near the origin of the coordinate system. The pre-selection using a rotation angle threshold can therefore be omitted, as the vector quantization step will only select a few movements having unsuitable rotation angles while the majority will be alright. This feature is the main advantage of

the algorithm. Also, 3D vector quantization can be used in a straightforward fashion, as the 3D axis/angle vectors ω_i have three degrees of freedom.

The basic principle of the rest of the algorithm remains the same as previously, except that vector quantization is now carried out on the vector where rotation axis and angle are encoded.

3.7. Vector Quantization using Quaternions

As in Section 3.6, a data selection algorithm is presented here that does not remove movements having small rotation angles and therefore has no need for thresholds. The quaternion representation of 3D rotations is now used (see Appendix A.3).

After the quaternion representation has been computed for each relative movement, the ambiguity in the quaternion representation has to be resolved. Since the quaternions q_i and $-q_i$ represent the same rotation, we restrict the quaternions to the hyper-hemisphere with positive real part. This can be done similarly to the axis ambiguity resolution discussed in Section 3.3, with the difference that a quaternion consists of four elements instead of three. The remaining part of the algorithm is similar to the data selection methods discussed before.

The main advantage of quaternions compared to using the rotation axis is that the quaternion is well defined for arbitrary rotation angles. While the rotation axis is undefined for a rotation angle of zero, i.e. for movements where the rotation matrix equals $I_{3 \times 3}$, the corresponding quaternion is defined and equals 1. The main disadvantage of using quaternions is that these consist of four elements with only three degrees of freedom; therefore, a 4D vector quantization has to be used instead of a 3D vector quantization. As before when rotation axes were discussed, polar angles representing quaternions could of course be used, as these have norm one and thus lie on a hypersphere.

4. Experimental Results

This section presents an experimental evaluation of the vector quantization based data selection algorithms. It starts with an introduction to the metrics used for residual error computation in Section 4.1. The datasets used are described in Section 4.2, followed by the experimental results in Sections 4.3–4.6. We will look into the following topics in particular: how the codebook size used for vector quantization influences the calibration error (Section 4.3); how the pre-selection threshold on the rotation angle affects the result (Section 4.4); and which of the presented data selection methods perform best (Section 4.5)? Finally, Section 4.6 shows how the data selection performs compared to manually selecting well-suited poses during data acquisition (although this is not the main application area of the proposed methods). The actual hand-eye calibration was done using the linear dual quaternion algorithm by Daniilidis (1999, 2001).

4.1. Residual Error Metrics

For experimental evaluation, we need error metrics for rotation and translation that measure the accuracy of hand-eye calibration. Commonly, the error in translation is given as a relative error, while for rotation an absolute error metric is used (e.g. Horaud and Dornaika 1995; Daniilidis 1999; Andreff et al. 2001). In this paper, absolute and relative errors will be shown for both rotation and translation. The absolute residual error for translation is given by

$$\epsilon_{\text{tabs}} = \frac{1}{N} \sum_{i=1}^N \|\hat{t}_i - t_i\| \quad (11)$$

and the relative residual error by

$$\epsilon_{\text{trrel}} = \frac{1}{N} \sum_{i=1}^N \frac{\|\hat{t}_i - t_i\|}{\|t_i\|}, \quad (12)$$

where N is the number of translation vectors used for error computation, t_i is the true translation vector, and \hat{t}_i is the vector estimated by hand-eye calibration.

Different metrics for errors in rotation are used in literature. While the norm of the difference between two rotation matrices is given in Horaud and Dornaika (1995), this work follows Daniilidis (1999) and Andreff et al. (2001) and uses instead the norm of quaternion differences for relative residual errors, given by

$$\epsilon_{\text{Rrel}} = \frac{1}{N} \sum_{i=1}^N \frac{\|\hat{q}_i - q_i\|}{\|1 - q_i\|}. \quad (13)$$

The norm of quaternion differences is obviously connected to the rotation angle as well as to the angle between the two rotation axes:

$$\|\hat{q}_i - q_i\|^2 = 2 - 2(\cos \hat{\theta}_i \cos \theta_i + \hat{r}_i^T r_i \sin \hat{\theta}_i \sin \theta_i), \quad (14)$$

where $\hat{\theta}_i, \theta_i$ are the rotation angles and \hat{r}_i, r_i the rotation axes corresponding to the quaternions \hat{q}_i and q_i , respectively.

For rotations about the same axis but by different angles, this metric has the property that it is directly connected to the residual rotation angle as Equation (14) can be simplified to:

$$\|\hat{q}_i - q_i\| = \sqrt{2 - 2 \cos \frac{\hat{\theta}_i - \theta_i}{2}}. \quad (15)$$

The absolute residuals can either be given using quaternions as well, or in degrees based on the axis/angle representation of the residual rotation matrix R_{resi} , which is given by:

$$R_{\text{resi}} = \hat{R}_i^T R_i. \quad (16)$$

A simplified absolute rotational residual error can now be defined by the rotation angle θ_{resi} , which can be computed from one of the complex eigenvalues of R_{resi} :



Fig. 6. Left: optical tracking system smARTtrack1 in the laboratory and right: smARTtrack1 in the operating room during an endoscopic surgery (images courtesy of F. Vogt).

$$\epsilon'_{Rabs} = \frac{1}{N} \sum_{i=1}^N |\theta_{resi}|. \quad (17)$$

The advantage of using this metric instead of the quaternion based metric is simply that an absolute residual given in degrees makes it easier for the reader to judge whether the error is high or low.

The metrics presented above are used for computing a prediction error, which is the residual between the predicted eye position computed from hand data and the estimated hand-eye transformation, and the real (calibrated) eye pose. In order to give an overall residual error, a set of relative movements (we use 100) is selected randomly from the complete set of all possible relative movements (Section 3.1). Note that again it is of disadvantage to use relative movements between subsequent positions because the movements will usually be small. This results in large relative errors and thus does not reflect the actual quality of the estimated hand-eye transformation. The results shown in Tables 2–5 have been obtained by iterating the above process 100 times and averaging the resulting residual errors.

The reason why the randomly chosen movements are selected from the complete dataset rather than the set obtained after vector quantization is that the latter would not result in a valid residual error that can be used to describe the actual calibration accuracy, because these data have been used for parameter estimation. As the estimation process optimizes the hand-eye parameters on the provided data, the residual error would always be minimal.

To summarize, the residual errors in translation shown in the following were computed using Equations (11) and (12).

For relative residual errors in rotation, Equation (13) was used. The absolute rotational errors show the average rotation angle in degrees, which was computed using Equation (17). All results are given with an accuracy of three valid digits.

4.2. Description of Datasets

Instead of a robot as in classic hand-eye calibration, we have used an optical tracking system in our experiments. The infrared optical tracking system smARTtrack1 by Advanced Realtime GmbH (shown in Figure 6) provides pose data of a so-called target (the *hand*) that is fixed to an endoscope. It is a typical optical tracking system consisting of two infrared cameras and the target, which is built from markers that can easily be identified in the images captured by the cameras. Spheres with a retro-reflective surface are used, and marker identification is simplified by active illumination with infrared light. The 3D position of each visible marker is calculated by the tracking system; knowledge of the geometry of the target then allows us to calculate its pose. The accuracy of the pose is 0.19 mm in the x and y direction, 0.36 mm in the z direction and 0.14° for rotation.

A CCD camera is mounted rigidly on the endoscope, which is moved manually. The objective of hand-eye calibration is to determine the unknown transformation from the target pose provided by the optical tracking system to 3D camera coordinates.

The camera (eye) poses are computed using a calibration pattern and standard camera calibration techniques (Zhang 1998, 2000). We use an asymmetric pattern as shown in Figure 7 (left), which consists of 49 circular calibration points




Fig. 7. Left: asymmetric 7 × 7 calibration pattern with marked corners and right: processed image of calibration pattern after ellipse-fitting.

Table 1. Properties of the datasets used in the experiments, including data selection information.

Sequence	ART1	ART2	ART3	ART4
Number of poses	270	190	200	200
Total number of relative movements	36 315	17 955	19 900	19 900
Number of movements after applying threshold	7261	8976	8954	6964
Minimum angle in dataset	0.000 001 71°	0.0388°	0.135°	0.0406°
Maximum angle in dataset	80.3°	80.9°	126°	73.8°
Minimum angle after applying threshold	38.5°	25.0°	48.4°	34.5°
Maximum angle after applying threshold	80.3°	80.9°	126°	73.8°
Median angle in dataset	21.8°	25.2°	42.8°	25.2°
Median angle after threshold	47.0°	41.2°	73.4°	47.6°

arranged in a 7 × 7 pattern. The corners of the pattern are coded using smaller dots to allow for resolving ambiguities when the camera is moved. We extract contours in the captured image of the calibration pattern, and perform ellipse-fitting on each contour (Figure 7, right). The center of each ellipse is used as a 2D calibration point which provides sub-pixel accuracy; the measured back-projection error is smaller than 0.2 pixels.

Results from eight representative datasets acquired this way are shown here. These have been selected carefully out of many more experiments that we have conducted, to describe effects that can be observed when using the proposed data selection methods. Four sets contain a large number of data continuously recorded while the endoscope was moved manually. These are denoted by ART1, ART2, ART3 and ART4 in the following. The remaining four contain only a small number of poses, which were recorded at manually selected positions as in classic hand-eye calibration.

The continuously recorded datasets differ mainly in the number of poses contained in each sequence (270 for ART1, 190 for ART2, 200 for ART3 and ART4) and in the type of movement made while they were recorded. The camera-endoscope configuration, i.e. the hand-eye transformation, is different for all datasets as the camera was re-mounted on the endoscope every time. An overview of the main properties of the datasets, including the results of automatic threshold computation, are shown in Table 1.

The the accuracy of the hand-eye calibration is heavily dependent on the amount of rotational movement in terms of both distinct rotation axes as well as large rotation angles. The quality of a dataset can therefore be assessed by plotting the distribution of normalized rotation axes and histograms of rotation angles, as shown in Figures 8 and 9.

As can be seen in the plots of normalized rotation axes after pre-selection in Figure 8, the rotational movement in ART1 and ART3 is smaller than in ART4, and considerably smaller than




Fig. 8. Comparison of normalized rotation axes contained in the continuously recorded datasets after pre-selection: (a) *ART1*; (b) *ART2*; (c) *ART3*; and (d) *ART4*. The variation in the rotation axes is much larger for *ART2* and *ART4* compared to *ART1* or *ART3*, making the former better suited for hand-eye calibration.

in *ART2*. It is important to note that a dataset has to contain a certain amount of rotational movement to achieve good results. Therefore, based on rotation axes, the datasets *ART2* and *ART4* are much better suited for hand-eye calibration than *ART1* and *ART3*, due to the small coverage of the latter.

Figure 9 shows histograms of rotation angles contained in the datasets before (left column) and after (right column)

pre-selection of movements using an automatically computed rotation angle threshold. It can be observed that the majority of movements before pre-selection have small rotation angles, and only *ART3* contains angles greater than 90° . Recall that for obtaining a very accurate calibration we would need movements with rotation angles around 90° . Based on the histograms we can conclude that *ART3* is the best-suited dataset.

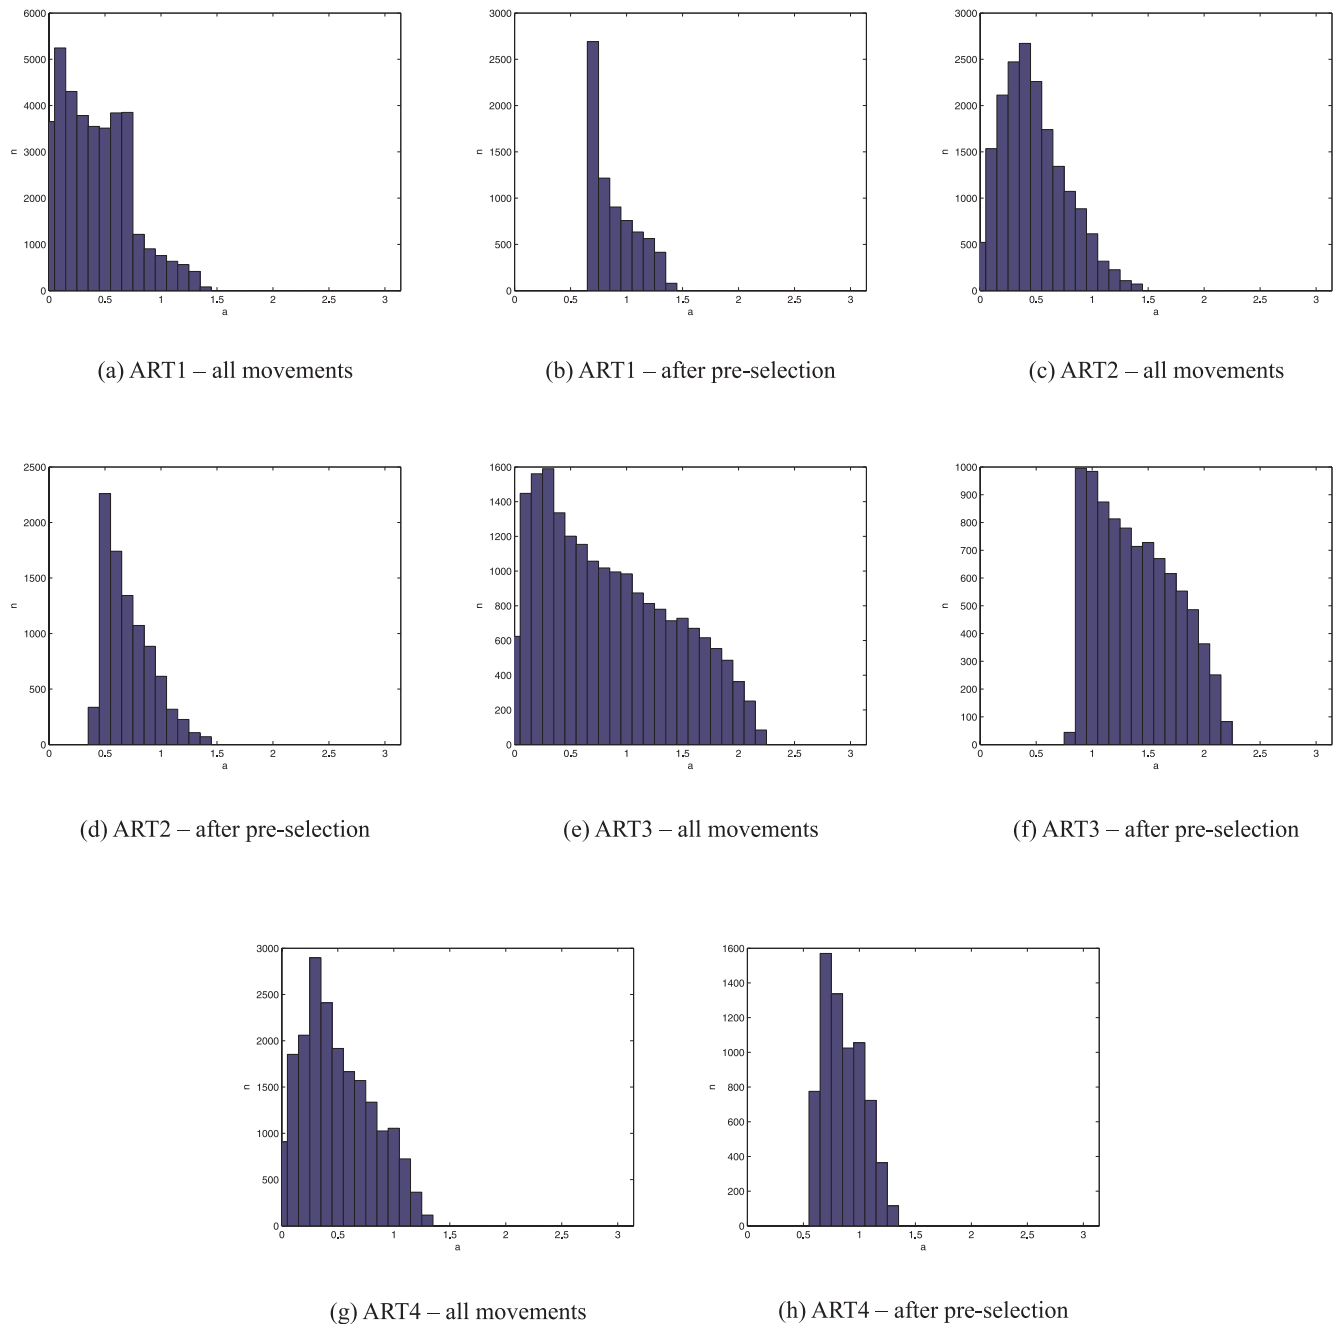


Fig. 9. Histograms of rotation angles (rad) in the datasets: (a) *ART1* – all movements; (b) *ART1* – after pre-selection; (c) *ART2* – all movements; (d) *ART2* – after pre-selection; (e) *ART3* – all movements; (f) *ART3* – after pre-selection; (g) *ART4* – all movements; and (h) *ART4* – after pre-selection.

However, combined with the distribution of axes (Figure 8(c)) which has low coverage, judging the calibration outcome for *ART3* in advance from the plots is inconclusive.

Based on these figures it can already be predicted that *ART1* is probably a dataset that results in a low-accuracy calibration, as neither rotation axes nor angles are well suited for hand-eye

calibration. It is likely that *ART2* and *ART4* will yield better results, as the coverage of axes is good and the histograms of angles are acceptable, although not optimal.

In order to get an idea of the quality of the hand-eye calibration obtained from continuously recorded sequences, the results were compared to the accuracy of performing hand-

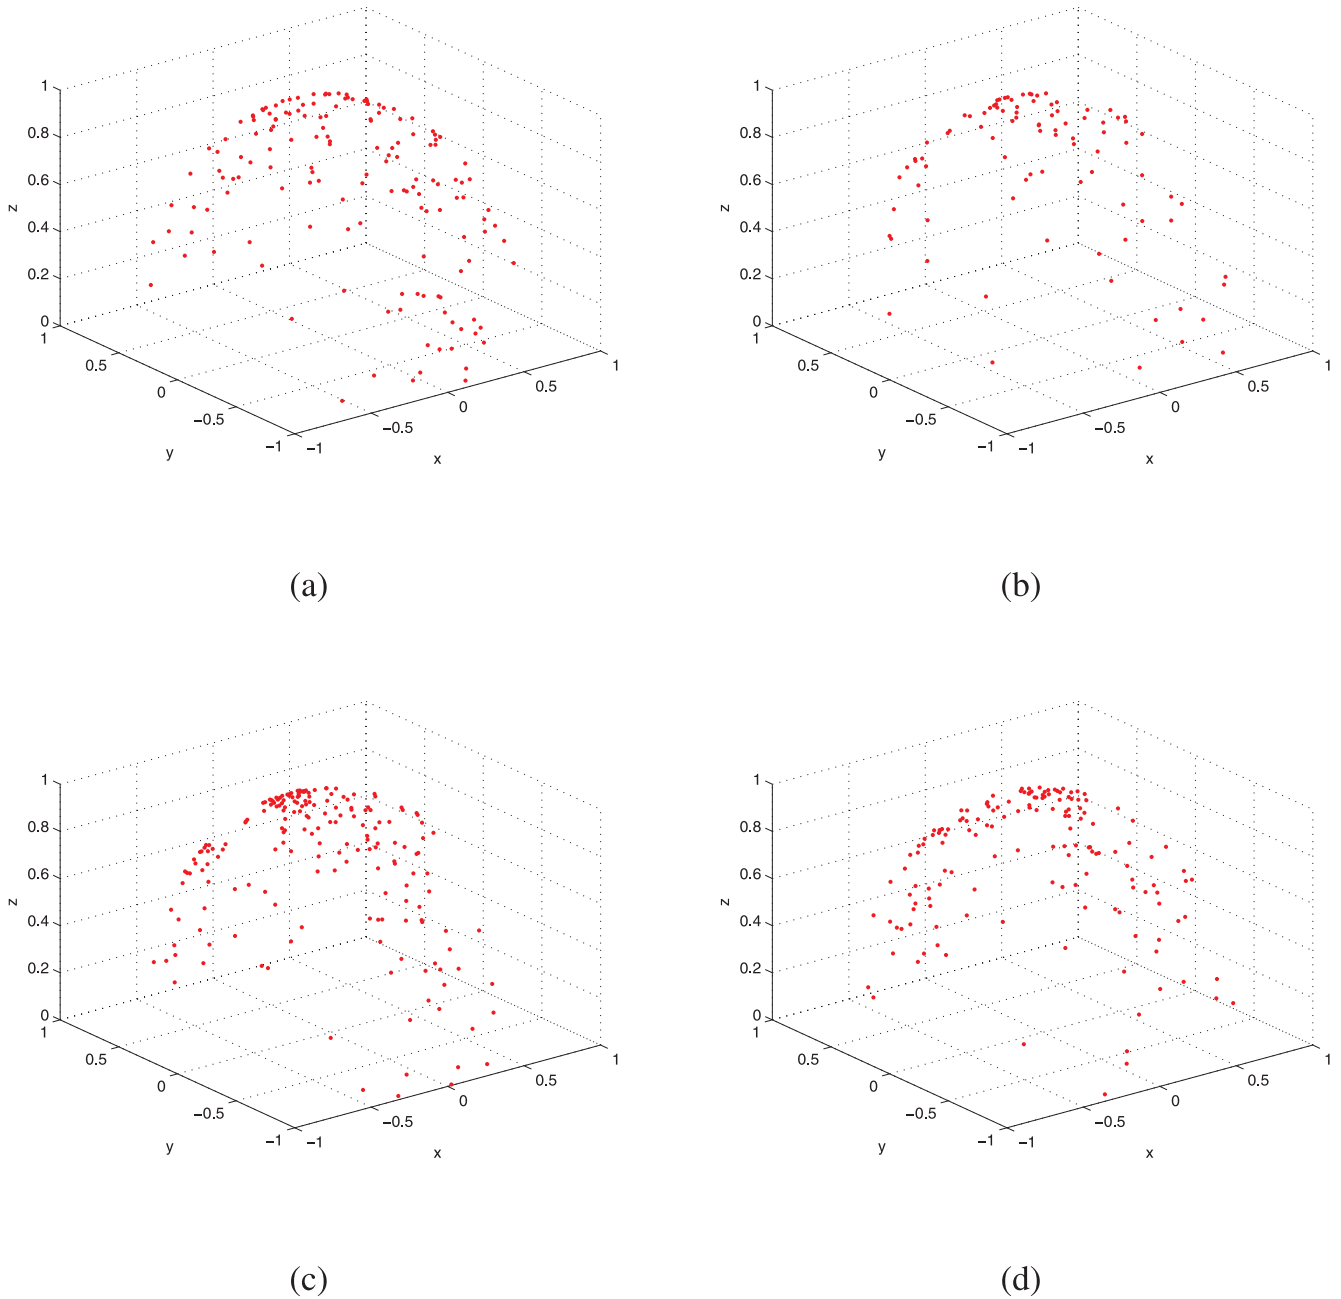


Fig. 10. Comparison of normalized rotation axes contained in the datasets recorded at manually selected positions: (a) *ART5*; (b) *ART6*; (c) *ART7*; and (d) *ART8*. These datasets contain less poses than the continuously recorded ones, but the distribution of rotation axes is much better suited for hand-eye calibration.

eye calibration in the classic way i.e. using a small number of poses recorded at manually selected distinct positions that are well-suited for hand-eye calibration. These were acquired using the same camera-endoscope configuration as continuously recorded data, and are denoted by *ART5* (18 poses), *ART6* (14 poses) (both corresponding to *ART2*), *ART7* (20 poses, cor-

responding to *ART3*) and *ART8* (18 poses, corresponding to *ART4*).

Figure 10 shows plots of the distribution of rotation axes of all relative movements contained in these four datasets. Comparing these figures to the plots in Figure 8, it can be observed that the coverage is much better than for the continuous se-

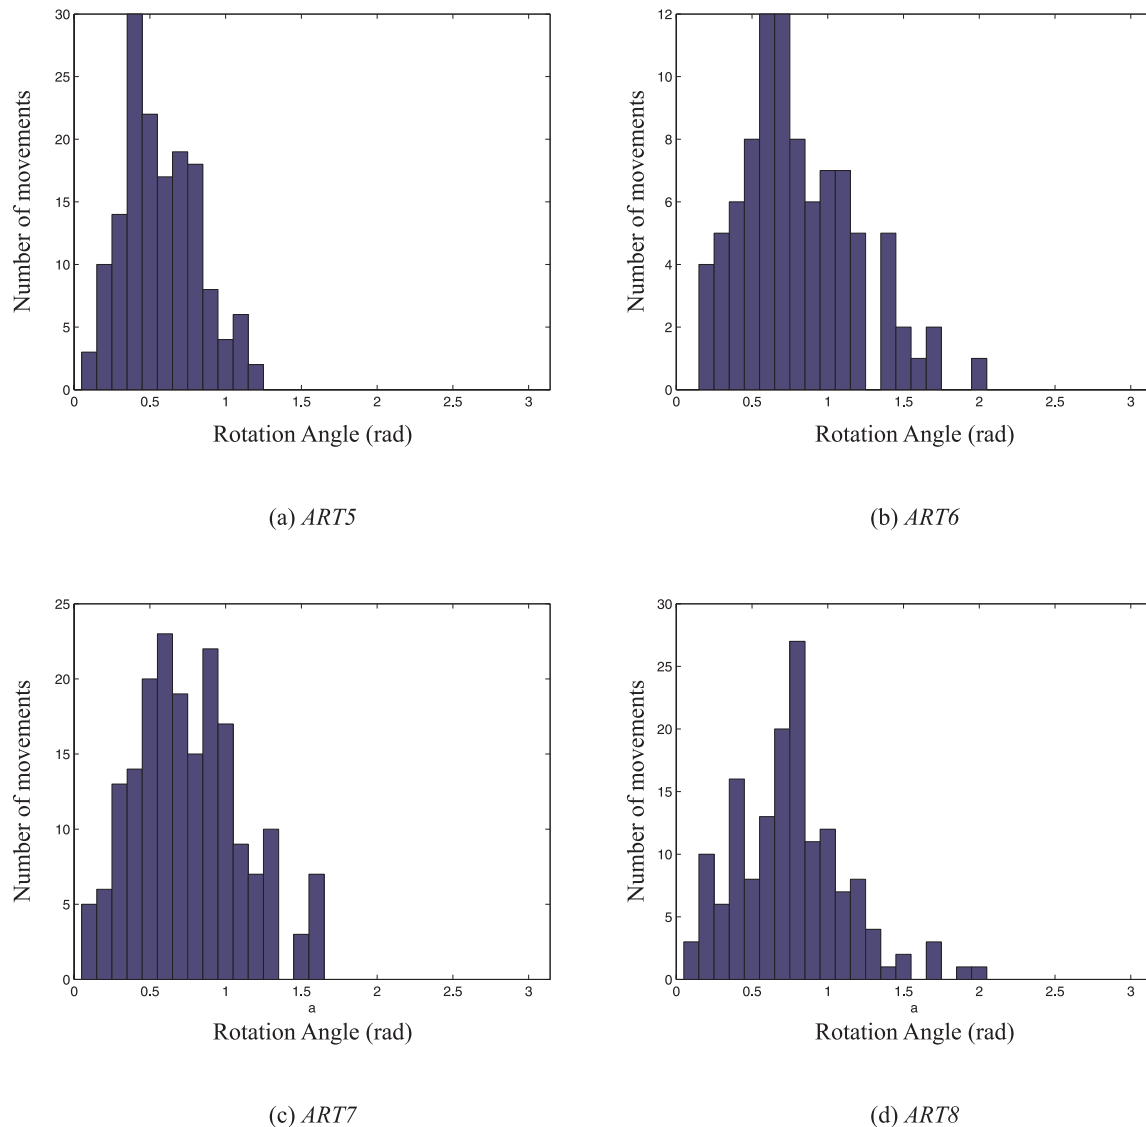


Fig. 11. Histograms of rotation angles (rad) in the datasets recorded at manually selected positions: (a) *ART5*; (b) *ART6*; (c) *ART7*; and (d) *ART8*.

quences. They obviously contain much less data, but it is not quantity that counts. Histograms of the rotation angles (without any pre-selection) are shown in Figure 11. As opposed to the histograms in Figure 9, the manually recorded datasets contain less very small angles close to 0° . Although they too are far from being optimal, with the exception for *ART5* they all contain angles greater than (but still close to) 90° .

To summarize, even without performing an actual calibration, it can be concluded from the plots alone that in most cases continuously recorded data will lead to a less accurate calibration compared to data containing manually selected positions. However, in many applications it may be very inconvenient or even impossible (e.g. due to time constraints during a surgical

intervention as in the case of the endoscope) to use manually selected positions.

In the following sections we present an evaluation of the data selection algorithms, compare continuously recorded data with data acquired at manually selected positions calibrated the classic way and show that even for the manually selected datasets, performing a data selection can result in a more accurate calibration.

4.3. Codebook Size

In this section the influence of one of the two data selection parameters is analyzed, namely the codebook size used for

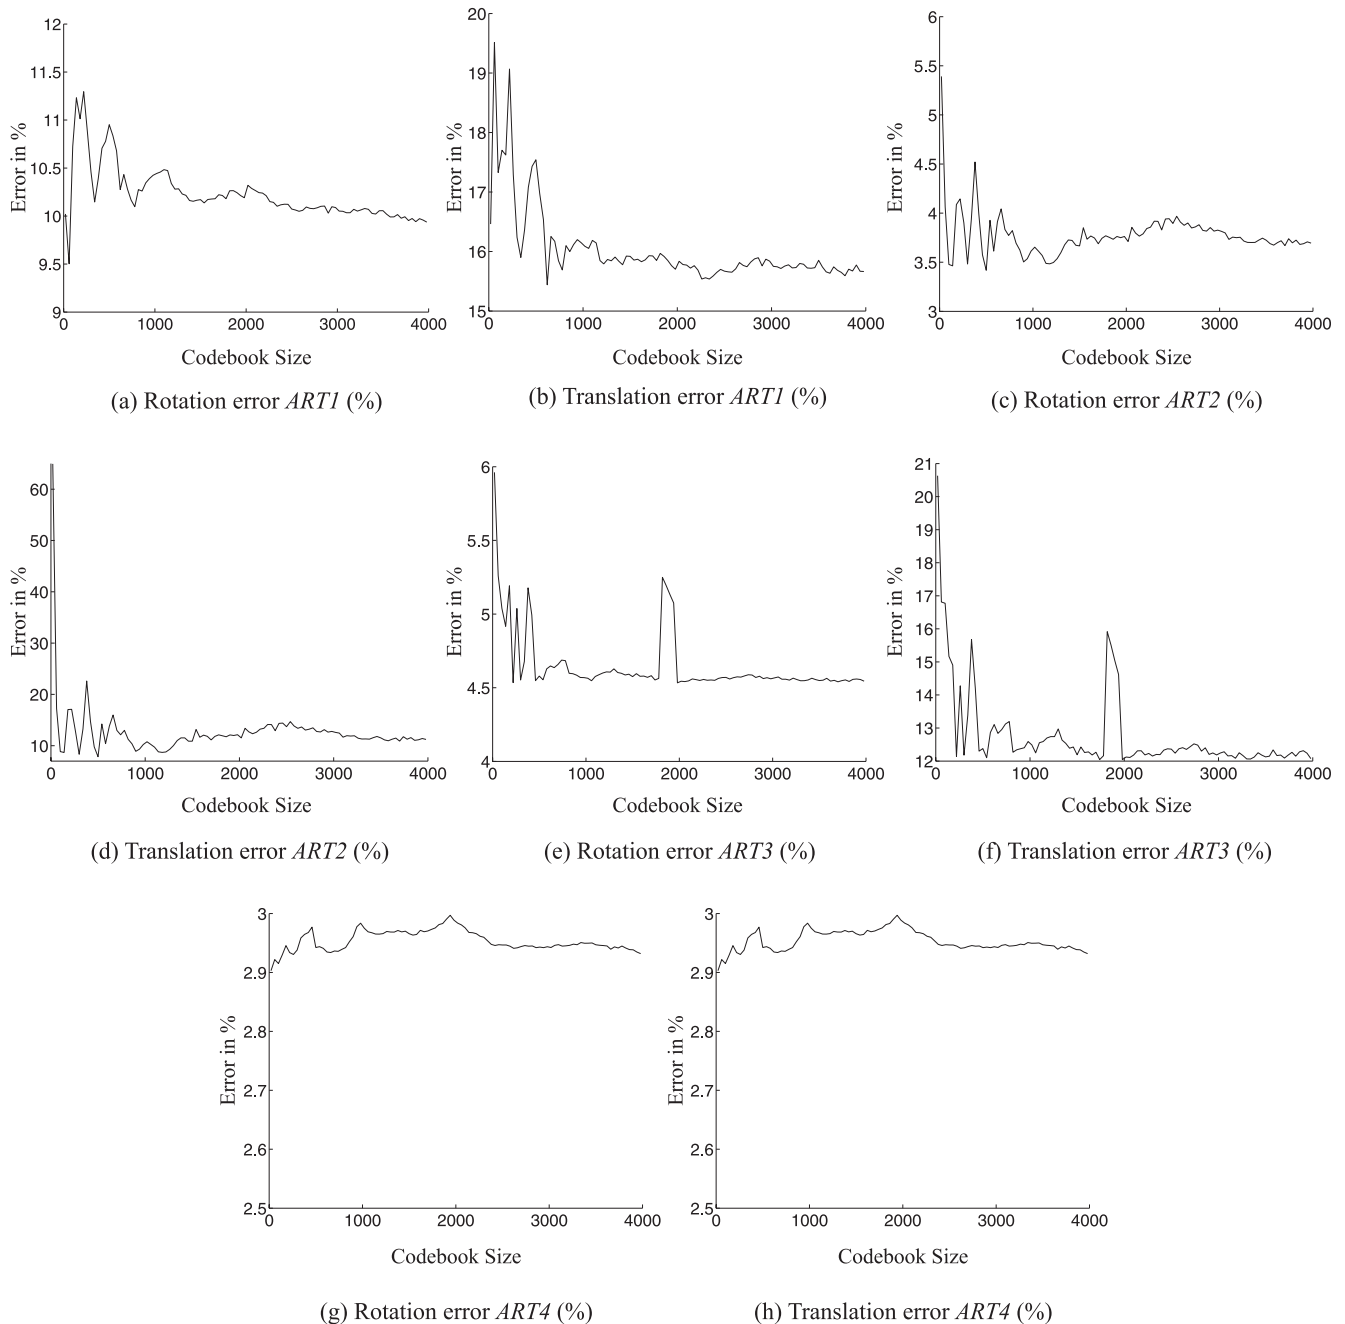


Fig. 12. Mean relative residual errors in rotation and translation (%) dependent on the codebook size used for vector quantization for the datasets: (a) rotation error *ART1*; (b) translation error *ART1*; (c) rotation error *ART2*; (d) translation error *ART2*; (e) rotation error *ART3*; (f) translation error *ART3*; (g) rotation error *ART4*; and (h) translation error *ART4*.

vector quantization. The hand-eye calibration residual errors were computed varying the codebook size, while the other parameter was left constant. For data selection, the 3D vector quantization algorithm on normalized rotation axes (Section 3.3) was chosen and the pre-selection step that removes

movements having small rotation angles was done using the automatic threshold computation (Section 3.5).

Plots of the relative residual errors in rotation and translation for the continuously recorded datasets *ART1*, *ART2*, *ART3* and *ART4* are shown in Figure 12; the graphs for absolute

residuals have been omitted, as they look very much like the relative ones and therefore do not provide much more information.

The codebook size was varied from 20 to 4000 relative movements, which is 0.22–57% (depending on the dataset) of the number of relative movements remaining after pre-selection with respect to the rotation angle, and only 11–22% of the total number of relative movements contained in the sequence. Typically, what can be observed is the behavior of the datasets *ART1* and *ART2* and, with the exception of some outliers in the middle, *ART3* as shown in Figures 12(a)–12(f). Fluctuations are quite high for small codebook sizes, the residual error becoming more stable as a larger amount of data is used for calibration. In contrast to the other datasets, *ART4* is stable from the beginning (note the scale of the vertical axis). This is probably due to the fact that *ART4* is obviously the best-suited dataset for hand-eye calibration, judging by the residual errors.

For practical purposes, low values for codebook size is preferable to high values, as the number of vectors in the codebook has a twofold impact on computation time. Vector quantization itself will take longer to compute as will the actual hand-eye calibration as more data are used. All datasets shown can arguably be considered to be relatively stable from a codebook size of ≥ 2000 for all sequences, which corresponds to about 5.5% (*ART1*), 10% (*ART3*, *ART4*) and 11% (*ART2*) of the total number of relative movements before pre-selection or about 22–29% of the relative movements after pre-selection.

Note that the residuals for *ART1* are considerably higher than those of the other sequences. This is not due to the data selection or hand-eye calibration algorithms used but is inherent in the data, which does not contain sufficient information for an accurate calibration due to only small rotational movement. In other words, rotation axes are more concentrated and the amount of rotation is smaller than in the other datasets.

4.4. Rotation Angle Threshold

The second data selection parameter that influences hand-eye calibration accuracy determines how much data is discarded during the pre-selection stage due to small rotation angles. For the experimental evaluation in this section, the automatic threshold computation algorithm described in Section 3.5 was used.

During the experiments codebook size was fixed for each dataset, while the percentage of relative movements that are removed by automatic threshold computation was variable. The codebook sizes used are 1200 (*ART1*), 1000 (*ART2*), 900 (*ART3*) and 700 (*ART4*). Residual error plots dependent on the fraction of data left after pre-selection according to the rotation angle are shown in Figure 13. A value of 1 on the horizontal axis is equivalent to 100% of the data used, i.e. no pre-selection at all, while 0 would indicate that no data was left

after pre-selection. As the latter case is pointless, the evaluation was only done up to a fraction where enough data was left to allow for hand-eye calibration, namely 7%.

The problem when removing relative movements based on the rotation angle alone is that an unknown amount of these movements may be suited quite well for hand-eye calibration in terms of non-parallel rotation axes, so there will always be a trade-off. This issue has been addressed before in Sections 3.6 and 3.7 where data selection algorithms have been presented that do not require a threshold for pre-selection, as none are carried out. Whether this results in a higher calibration accuracy than applying a rotation angle threshold is evaluated in Section 4.5.

As shown in the histograms in Figure 9, the datasets contain a considerable number of movements with relatively small rotation angle. Therefore, judging solely on rotation angle, what would be expected is that the plots in Figure 13 are U-shaped. When no movements are removed from the dataset, the residual error should be high because unsuitable data are used for calibration and skew the result. On the other hand, when most movements have been removed, the calibration residual should increase due to the very small amount of noisy data that is used, making hand-eye calibration more sensitive to erroneous movements.

While the increasing residual error at the lower end can be observed in all plots in Figure 13, the increase at the right end is visible only for the *ART3* and *ART4* datasets.

In general, a value of 20–40% for the data to be left after pre-selection with respect to the rotation angle is a relatively good choice for automatic threshold computation in most cases. If possible, smaller values are preferable to higher values, because the computation time of vector quantization depends on the size of the datasets after pre-selection.

4.5. Comparison of Data Selection Methods

In this section, a comparison of the performance of the various vector quantization based data selection methods is presented. Data selection is essential if continuous movements are used and the results when using relative movements between consecutive poses (i.e. without any data selection) are also shown.

The following methods were compared.

1. Consecutive movements: relative movements between consecutive poses are used. No data selection is done. These results are presented in order to show how much can be gained by data selection when continuously recorded sequences are used.
2. 3D VQ, normalized axes: data selection using normalized 3D rotation axes with two DOF as presented in Section 3.3.

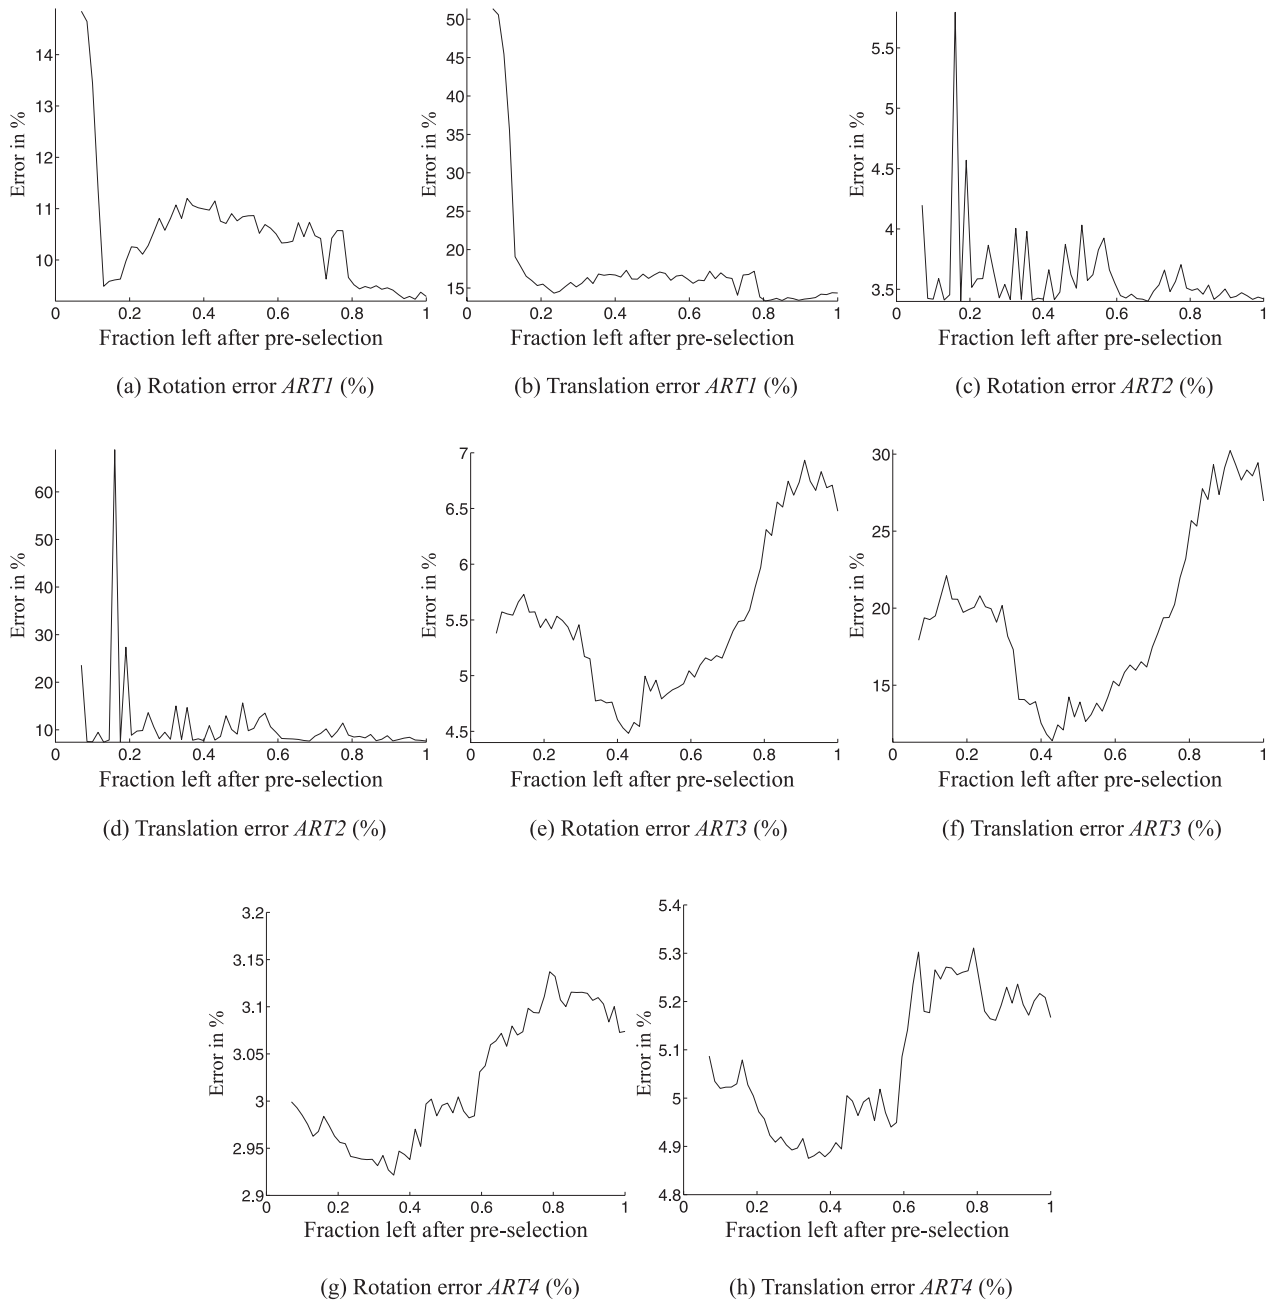


Fig. 13. Mean relative and absolute errors (%) in rotation and translation dependent on the fraction of relative movements left after pre-selection with respect to the rotation angle: (a) rotation error *ART1*; (b) translation error *ART1*; (c) rotation error *ART2*; (d) translation error *ART2*; (e) rotation error *ART3*; (f) translation error *ART3*; (g) rotation error *ART4*; and (h) translation error *ART4*.

3. 2D VQ, polar coordinates: data selection based on the polar coordinate representation of normalized rotation axes as presented in Section 3.4.
4. 3D VQ, axis/angle: data selection with three DOF based on the axis/angle representation of rotation axes, where

the rotation angle is encoded as the norm of the axis. This method was presented in Section 3.6.

5. 4D VQ, quaternions: data selection with three DOF based on the quaternion representation of rotations as shown in Section 3.7.

Table 2. Comparison of different data selection methods, sequence ART1 (parameters: codebook size 1200, $\delta = 0.2$).

Method	Translation		Rotation	
(1)	19.9 mm	49.8%	2.89°	14.7%
(2), automatic threshold	5.67 mm	15.8%	1.53°	10.3%
(3), automatic threshold	5.43 mm	15.3%	1.51°	10.3%
(4), no threshold	5.75 mm	15.8%	1.49°	10.1%
(4), automatic threshold	5.75 mm	16.0%	1.44°	10.1%
(5), no threshold	5.98 mm	16.4%	1.52°	10.2%
(5), automatic threshold	5.61 mm	15.7%	1.43°	10.0%

In case of (2) and (3), rotation angle thresholds were computed automatically using the method presented in Section 3.5. Obviously, this is not applicable for method (1).

Note that for methods (4) and (5) no pre-selection of the data with respect to small rotation angles is necessary as the angle is handled implicitly during vector quantization. The residuals for this case are described by 'no threshold'. However, the removal of relative movements with small rotation angles can nevertheless be done, and the results using automatically computed thresholds are also shown.

The tables show absolute and relative residual errors for rotation and translation. In the case of translation, the residuals were computed using Equation (11) for absolute errors and Equation (12) for relative errors.

The relative rotation error was computed based on quaternions using Equation (13). In order to obtain an impression of the order of magnitude of the absolute rotational error, it has been decided to show these in degrees rather than to give the absolute quaternionic residual. For this purpose, the absolute residuals have been computed from the rotation angle (in axis/angle representation) of the residual rotation matrix given by Equation (17). This value is highly correlated to the quaternionic residual. Note, however, that due to the different rotation representations, there will be slight deviations in some cases where absolute and relative residuals from different experiments are compared.

The calibration results are shown in Tables 2–5.

As predicted, using consecutive movements is always the worst case, with residual errors that render the calibration result totally useless in most cases. It can be observed that the data selection method used has virtually no influence on the residual errors in rotation. This result is as expected, as the rotation matrix can always be computed, even for movements which are not general enough, while this is not true for translation.

Recommending a single data selection method based on the results presented here is not easy, as there is no single method that consistently gives the best results. However, the follow-

Table 3. Comparison of different data selection methods, sequence ART2 (parameters: codebook size 1000, $\delta = 0.5$).

Method	Translation		Rotation	
(1)	4.98 mm	23.9%	0.854°	4.26%
(2), automatic threshold	1.98 mm	10.4%	0.626°	3.63%
(3), automatic threshold	3.12 mm	15.6%	0.797°	4.12%
(4), no threshold	1.52 mm	8.09%	0.563°	3.46%
(4), automatic threshold	3.37 mm	16.6%	0.777°	4.07%
(5), no threshold	2.11 mm	11.0%	0.662°	3.73%
(5), automatic threshold	2.21 mm	11.5%	0.652°	3.70%

Table 4. Comparison of different data selection methods, sequence ART3 (parameters: codebook size 1000, $\delta = 0.45$).

Method	Translation		Rotation	
(1)	9.94 mm	37.3%	2.61°	7.61%
(2), automatic threshold	2.80 mm	12.3%	1.05°	4.57%
(3), automatic threshold	4.01 mm	16.7%	1.28°	4.95%
(4), no threshold	3.96 mm	16.5%	1.48°	5.28%
(4), automatic threshold	3.69 mm	15.6%	1.36°	5.09%
(5), no threshold	4.12 mm	17.1%	1.55°	5.40%
(5), automatic threshold	3.94 mm	16.5%	1.42°	5.20%

Table 5. Comparison of different data selection methods, sequence ART4 (parameters: codebook size 700, $\delta = 0.35$).

Method	Translation		Rotation	
(1)	8.34 mm	17.7%	1.10°	4.59%
(2), automatic threshold	1.67 mm	4.87%	0.555°	2.93%
(3), automatic threshold	1.68 mm	4.89%	0.561°	2.95%
(4), no threshold	1.68 mm	4.85%	0.578°	2.99%
(4), automatic threshold	1.65 mm	4.85%	0.544°	2.90%
(5), no threshold	1.75 mm	4.96%	0.587°	3.01%
(6), automatic threshold	1.65 mm	4.86%	0.541°	2.90%

ing insights are to be gained. Using any data selection is better than doing nothing, as all methods perform much better than using consecutive movements directly. We believe that the method of choice is using normalized rotation axes with a pre-selection according to the rotation angle with an automatically computed threshold (case (2)), because it is consistently good and ranked in the top three in all experiments.

The results using rotation representations where the rotation angle is implicitly coded, such as quaternions and axis/angle ((4), no threshold and (5), no threshold), did not perform as well as the others, with the exception of the *ART2* experiment. Results improved when the automatic threshold computation was done, removing movements with small rotation angles.

With the exception of *ART1*, using polar coordinates was inferior to most other methods. Table 5 lists the results for the *ART4* experiment. If the dataset is suited well for calibration, i.e. the coverage of rotation axes is high and there is a sufficient amount of rotation angles close to 90° , it makes virtually no difference which data selection method is used. The difference between the best and worst results (not taking into account using consecutive movements for obvious reasons) is only 0.1 mm, and the difference between the best result and the recommended method using normalized rotation axes is only 0.02 mm.

The experiments were run on a computer with an Intel P4 2.6 GHz CPU. While the implementation was not particularly optimized for speed, we want to give at least a rough guideline regarding the computation times that can be expected for data selection and hand-eye calibration.

Using consecutive movements is obviously the fastest way to calibrate, as no data selection is performed, and takes about 50 ms. For the remaining methods, computation time depends mainly on the amount of data during vector quantization and its dimensionality. Not surprisingly, the fastest algorithm is that based on polar coordinates as only a 2D quantization has to be done. Computation times varied between 5.8 s (*ART4*) and 11 s (*ART1*). Accordingly, using quaternions and no threshold (i.e. 4D quantization and all data) is always the slowest option, with computation times between 10 s (*ART4*) and 31 s (*ART1*). This is similar for the axis/angle method using no threshold, where the time required is only slightly less, namely 9.5 s (*ART2*) to 26 s (*ART1*). When using a pre-selection with an automatically computed threshold, the computation times of these two algorithms are very similar to that for the normalized rotation axes method, which varies between 6.1 s (*ART4*) and 12 s (*ART1*).

4.6. Comparison with Manual Selection

There are obviously situations when acquiring well-suited data at manually selected positions is less convenient than in others. In a classic setup with a robot manipulator arm, the robot can easily be programmed to move to defined positions that are used for calibration. However, in setups such as hand-eye calibration of a camera mounted on an endoscope or a stereo-camera system used in Augmented Reality, the device is manually controlled by the user and obtaining good data is no longer straightforward. Two topics are therefore addressed in this section.

Table 6. Comparison of the continuously recorded sequences *ART2*, *ART3* and *ART4* to datasets using the same camera-endoscope configuration recorded at manually selected positions with a small number of poses. It can be observed that a manual selection of positions is not always superior to the conveniently recorded continuous sequences.

Dataset	Translation	Rotation
<i>ART5</i>	2.31 mm 10.2%	0.556° 3.46%
<i>ART6</i>	1.65 mm 8.02%	0.556° 3.47%
<i>ART2</i>	1.98 mm 10.4%	0.626° 3.63%
<i>ART7</i>	1.85 mm 8.41%	0.892° 4.34%
<i>ART3</i>	2.80 mm 12.3%	1.05° 4.57%
<i>ART8</i>	2.08 mm 5.49%	0.672° 3.21%
<i>ART4</i>	1.67 mm 4.87%	0.555° 2.93%

1. How well does hand-eye calibration based on continuously recorded sequences perform compared to calibration by the classic method using a small number of poses recorded at manually selected distinct positions that are well-suited for hand-eye calibration?
2. Can calibration accuracy be increased by running a data selection on datasets with manually selected positions, which is already good when consecutive poses are used?

Table 6 shows calibration results for datasets consisting of about 20 poses recorded at manually selected positions (calibrated the classic way), described in Section 4.2, and compares them to the corresponding continuously recorded data (calibrated using data selection with normalized rotation axes and automatic threshold computation). The camera-endoscope configuration (i.e. actual hand-eye transformation) of *ART5* and *ART6* is the same as for the continuously recorded dataset *ART2*. Likewise, *ART7* has the same configuration as *ART3*, and *ART8* the same as *ART4*.

Comparing *ART5* and *ART6* on one side to *ART2* on the other shows that the manual selection of positions while recording data is not always better than the conveniently recorded continuous sequence. In particular, the translation estimate of *ART2* is considerably better than that of *ART5*. The result for the manually selected positions (*ART7*) is much better than that of the continuously recorded *ART3*. The situation for the last set *ART8* and *ART4* shows again that a continuously recorded sequence can lead to better results than a manual selection of the positions for recording.

It is obvious that much depends on the dataset itself and the poses contained in it. When there is insufficient information available, the calibration results are also poor. Of course, in practice this may be more often the case for continuously recorded sequences, as the user usually does not choose the

Table 7. Comparison between the classic hand-eye calibration method using consecutive poses (left) and the method using the data selection proposed in this work (right). Clearly, an additional data selection improves the calibration accuracy on these datasets, which contain only a small number of poses recorded at manually selected distinct positions.

Dataset	Consecutive frames				Data selection			
	Translation		Rotation		Translation		Rotation	
<i>ART5</i>	2.31 mm	10.2%	0.556°	3.46%	2.13 mm	9.76%	0.552°	3.45%
<i>ART6</i>	1.65 mm	8.02%	0.556°	3.47%	1.32 mm	7.11%	0.538°	3.41%
<i>ART7</i>	1.85 mm	8.41%	0.892°	4.34%	1.73 mm	8.14%	0.865°	4.30%
<i>ART8</i>	2.08 mm	5.49%	0.672°	3.21%	1.86 mm	5.15%	0.518°	2.80%

positions as carefully as in the case where the data are acquired at distinct positions.

Even although the data selection methods have been developed for continuously recorded sequences where they are essential, the question remains whether applying them to the datasets containing only a small number of poses at manually selected positions would lead to even better calibration results. A comparison between the classic method using consecutive poses and the method using the data selection proposed in this paper is therefore presented in Table 7.

Again, the data selection method chosen for this comparison was vector quantization of normalized rotation axes with automatic threshold computation. In all cases, the calibration with data selection was superior to that using consecutive poses as is usually done for hand-eye calibration. Except for *ART8* (where the absolute residual error in rotation changed by 0.154°, which is a 23% decrease compared to calibration without data selection), the rotation residuals did not change by much. The translational residual error, however, improved considerably for all datasets, which is a result of using well-suited data for hand-eye calibration.

The results highlight that even when the user tries to perform a small amount of well-described calibration movements as accurately as possible, using data selection will still increase calibration accuracy.

5. Conclusion

We have proposed new data selection methods that can improve the accuracy of hand-eye calibration in many cases, and make it possible in the first place in situations where a manual selection of positions is inconvenient or even impossible. Examples include areas other than the classic setup with a robot manipulator arm and a camera. These include calibration of a camera mounted on an endoscope where the ‘hand’ data are provided by an optical tracking system or applications such as self-calibration of a rigid stereo-camera system. These appear to be unrelated to hand-eye calibration initially, but nevertheless can be solved using algorithms adopted from hand-eye calibration.

The result of the data selection algorithm is a dataset that is well suited for hand-eye calibration as it removes relative movements with small rotation angles and selects those movements where the rotation axes are different. Data selection as presented in this paper is based on applying a clustering algorithm on the data. For this purpose, we proposed to use vector quantization, but in general any clustering algorithm is suitable.

A variety of methods are proposed, which differ in the dimensionality (2D, 3D and 4D) of the vector quantization compared to the degrees of freedom (two or three) and whether an automatically computed threshold or no threshold at all is used for incorporation of the rotation angle. The methods using no threshold are based on 3D and 4D vector quantization using the axis/angle or quaternion representation of rotations, respectively. They are a trade-off between the non-parallelism criterion for the rotation axes and the fact that for movements with small rotation angles the axis is not well defined. The former methods remove movements with small angles in a pre-processing step and use only the differences in the rotation axes as a selection criterion.

The performance of the proposed vector quantization based data selection methods was evaluated using data obtained from an optical tracking system (hand) and an endoscopic camera (eye) that was calibrated using a calibration pattern.

Firstly, the parameters that influence hand-eye calibration accuracy were evaluated: the codebook size of the quantizer and the thresholds used for pre-selection of movements with respect to their rotation angle. It has been found that the fluctuations of the residual error for small codebook sizes are relatively high and become less for increasing codebook size. According to experimental results, using about 10% of the total number of relative movements as the codebook size (corresponding to about 20–30% of the movements left after pre-selection), is recommended.

Before vector quantization, a pre-selection step is performed that discards movements with small rotation angles from the data, as these are not suitable for hand-eye calibration. The data remaining after this pre-selection has a high influence on the computation time required during vector quantization.

These experiments were performed for fixed codebook sizes (but different sizes for each dataset), while the percentage of movements that are removed by automatic threshold computation was varied. High residual errors can be expected when either most movements have been removed or when all movements have been left in the data used for calibration. In most cases, a value of 20–40% for the data remaining after pre-selection is a good choice for automatic threshold computation.

Next, the different data selection methods proposed in this paper were compared. The new methods are based on vector quantization and differ in the parameterization used for representing 3D rotation (normalized rotation axes, polar coordinates, axis/angle, quaternions) and in the way the data pre-selection with respect to the rotation angle is done (automatically computed thresholds or no thresholds). These methods were compared to using consecutive movements, i.e. using the movements as they are, without any data selection. As predicted, using consecutive movements is always the worst case, with residual errors that render the calibration results totally useless in most cases.

As expected, the data selection method used has virtually no influence on the residual errors in rotation, only the translational residuals differ. To obtain the best results regarding accuracy, it is recommended to apply the vector quantization based data selection using normalized rotation axes. It should be noted, however, that much depends on the dataset itself. In other words, when the information contained in the movements is not general enough in terms of different rotation axes and high rotation angles, there is no way to obtain good calibration results, no matter which data selection or hand-eye calibration algorithm is used.

It has also been found that even in situations where calibration data are acquired using a small number of manually selected positions, applying an additional data selection can improve the accuracy of hand-eye calibration over the classic approach using consecutive movements. Even when the user is asked to perform a small number of well-described calibration movements, applying data selection on the acquired data will still result in improved accuracy of the hand-eye transformation.

Acknowledgements

The ART datasets used in the experiments were provided by F. Vogt, SFB 603, TP B6.

Appendix: 3D Rotation Parameterization

As the data selection algorithms presented in Section 3 make extensive use of various representations of 3D rotations, a brief

overview of the commonly used rotations is given here. A rotation in 3D is usually given by a rotation matrix $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ with the property that the column (and row) vectors are orthonormal, i.e.

$$\mathbf{R}\mathbf{R}^T = \mathbf{I}_{3 \times 3}, \quad \det(\mathbf{R}) = 1. \quad (18)$$

A rotation matrix \mathbf{R} therefore has nine elements but only three degrees of freedom. The set of all these matrices forms the rotation group $\text{SO}(3)$.

A.1. Cardan and Euler Angles

We discuss the Cardan and Euler angle representations only briefly as they are not employed for data selection due to inherent problems. However, as both representations are widely used, they should not be completely omitted.

To obtain the Cardan angle representation, an arbitrary rotation matrix \mathbf{R} is decomposed into a product of three rotations by the angles α , β and γ about the x , y and z axes of the coordinate system. In contrast to the Euler angle representation, the Cardan angles α , β and γ are defined with respect to the axes of the original coordinate system.

In the Euler angle representation, an arbitrary rotation matrix \mathbf{R} is also decomposed into a product of three rotations by the angles ϕ , ψ and φ , where

- ϕ defines a rotation about the z axis of the original coordinate system;
- ψ defines a rotation about the x' axis, which is the image of the x axis of the original coordinate system after the first rotation; and
- φ defines a rotation about the z'' axis, which is the image of the z axis of the original coordinate system after the previous two rotations have been computed.

Cardan and Euler angles are probably the most well-known parameterizations for rotations in 3D. These two representations sometimes become confused in literature, but usually the conclusions drawn for Cardan and Euler angles stay the same. One of the main drawbacks is that since matrix multiplication is not commutative, the Cardan/Euler angle representation is not unique meaning that a permutation of the order of the rotations about the coordinate-system axes yields different Cardan/Euler angles. Probably the most important drawback of these parameterizations is the existence of so-called *gimbal lock* singularities, where one degree of freedom is lost, i.e. two of the three angles belong to the same degree of freedom. For a more detailed discussion see Watt and Watt (1992).

A.2. Axis/Angle

An arbitrary rotation \mathbf{R} can be represented as a rotation about *one* axis $\mathbf{r} \in \mathbb{R}^3$ by the angle θ . This will be denoted here as *axis/angle representation*. Since only the direction of the rotation axis \mathbf{r} is of importance, \mathbf{r} has only two degrees of freedom and thus can be normalized to one. Hence, axis and angle can be combined into a single vector $\boldsymbol{\omega}$ with three degrees of freedom, its direction giving the rotation axis and its length the rotation angle, thus making it a minimal parameterization:

$$\boldsymbol{\omega} = \theta \mathbf{r}, \quad \text{and for } \theta \neq 0: \theta = |\boldsymbol{\omega}|, \quad \mathbf{r} = \frac{\boldsymbol{\omega}}{|\boldsymbol{\omega}|}. \quad (19)$$

Computing a rotation matrix \mathbf{R} from $\boldsymbol{\omega}$ can be done using Rodrigues' formula (Faugeras 1993; Hartley and Zisserman 2003)

$$\begin{aligned} \mathbf{R} &= \mathbf{I}_{3 \times 3} + \frac{\sin \theta}{\theta} [\boldsymbol{\omega}]_{\times} + \frac{1 - \cos \theta}{\theta^2} [\boldsymbol{\omega}]_{\times}^2 \\ &= \mathbf{I}_{3 \times 3} + \sin \theta [\mathbf{r}]_{\times} + (1 - \cos \theta) [\mathbf{r}]_{\times}^2, \end{aligned} \quad (20)$$

where $[\mathbf{x}]_{\times}$ denotes the skew-symmetric matrix that represents the outer vector product of \mathbf{x} and \mathbf{y} as a matrix multiplication:

$$\mathbf{x} \times \mathbf{y} = [\mathbf{x}]_{\times} \mathbf{y} = \begin{pmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{pmatrix} \mathbf{y}. \quad (21)$$

The computation of axis and angle from a rotation matrix \mathbf{R} is carried out as follows. Eigen-decomposition of \mathbf{R} yields the three eigenvalues 1 and $\cos \theta \pm i \sin \theta$. The axis \mathbf{r} is the eigenvector corresponding to the eigenvalue 1. The angle θ is calculated from one of the remaining eigenvalues. Note that the axis/angle representation is not unique: a rotation about an axis \mathbf{r} by an angle θ is the same as a rotation about the axis $-\mathbf{r}$ by the angle $2\pi - \theta$. Therefore, one has to check the consistency of the direction of the axis and the angle, which can be done by inserting both into Equation (20). Another problem arises for a rotation angle of 0° , i.e. if $\mathbf{R} = \mathbf{I}_{3 \times 3}$. In this case, all three eigenvalues are equal to one which results in a non-unique rotation axis. This is obvious, since for an angle of 0° no rotation is done at all, meaning that the axis can obviously be chosen arbitrarily.

A.3. Quaternions

Quaternions are numbers i.e. they form one of the four existing normed division algebras. (The others are the real and complex

numbers and the Octonions; Conway and Smith 2003; Baez 2001.) They are in a certain sense similar to complex numbers: instead of only one imaginary part, quaternions have three of them. The concept of quaternions was introduced by Sir William Rowan Hamilton and presented to the Royal Irish Academy in 1843 (Hamilton 1844, 1847, 1848). The set of quaternions is usually denoted as \mathbb{H} . Unit quaternions form the special unitary group $\text{SU}(2)$, which can be represented as all complex unitary 2×2 matrices having determinant one. Since $\text{SU}(2)$ is a double cover of the special orthogonal group $\text{SO}(3)$, there exist two quaternions for each rotation matrix. More details on quaternions can be found in Faugeras (1993); Kuipers (1999); Conway and Smith (2003).

A quaternion \mathbf{q} is defined as follows:

$$\mathbf{q} = q_r + q_1 i + q_2 j + q_3 k, \quad q_r, q_1, q_2, q_3 \in \mathbb{R}, \quad (22)$$

where q_r is the real part and q_1, q_2, q_3 are the imaginary parts. Multiplication and summation are done component-wise, with

$$i^2 = j^2 = k^2 = ijk = -1 \quad (23)$$

which is equivalent to

$$i^2 = j^2 = k^2 = -1,$$

$$ij = -ji = k, \quad jk = -kj = i, \quad ki = -ik = j. \quad (24)$$

A quaternion is often written as a 4-tuple

$$\mathbf{q} = (q_r, q_1, q_2, q_3) \quad \text{or} \quad \mathbf{q} = (q_r, \mathbf{q}_{\text{im}}), \quad (25)$$

where \mathbf{q}_{im} is a 3-vector containing the imaginary parts. In contrast to complex numbers, the commutative law of multiplication is *not* valid (Equation (24)), i.e.

$$\exists \mathbf{q}_1, \mathbf{q}_2 \in \mathbb{H}, \quad \text{where } \mathbf{q}_1 \mathbf{q}_2 \neq \mathbf{q}_2 \mathbf{q}_1. \quad (26)$$

Similar to complex numbers, a conjugate quaternion is defined as

$$\mathbf{q}^* = q_r - q_1 i - q_2 j - q_3 k. \quad (27)$$

The norm of a quaternion \mathbf{q} is given by

$$|\mathbf{q}| = \sqrt{\mathbf{q} \mathbf{q}^*} = \sqrt{\mathbf{q}^* \mathbf{q}} = \sqrt{q_r^2 + q_1^2 + q_2^2 + q_3^2}. \quad (28)$$

The multiplicative inverse of \mathbf{q} is

$$\mathbf{q}^{-1} = \frac{1}{\mathbf{q} \mathbf{q}^*} \mathbf{q}^*. \quad (29)$$

Hence, for a unit quaternion ($|\mathbf{q}| = 1$), the inverse of multiplication equals the conjugate, i.e. $\mathbf{q}^{-1} = \mathbf{q}^*$.

Just as the multiplication of two unit complex numbers defines a rotation in two dimensions, a multiplication of two unit quaternions yields a rotation in 3D. Let \mathbf{p} be a 3D point

to be rotated, \mathbf{r} a rotation axis with $|\mathbf{r}| = 1$ and θ the angle of rotation about this axis. Define the quaternions

$$\mathbf{q} = \left(\cos \frac{\theta}{2}, \sin \frac{\theta}{2} \cdot \mathbf{r} \right),$$

$$\mathbf{p}' = (0, \mathbf{p}). \quad (30)$$

Then

$$\mathbf{p}'_{\text{rot}} = \mathbf{q}\mathbf{p}'\mathbf{q}^{-1} = \mathbf{q}\mathbf{p}'\mathbf{q}^* \quad (31)$$

since \mathbf{q} is a unit quaternion. \mathbf{p}'_{rot} is the quaternion corresponding to the rotated point.

Since a quaternion representing a rotation is computed from axis and angle, it is not unique as the two quaternions

$$\mathbf{q}_1 = \left(\cos \frac{\theta}{2}, \sin \frac{\theta}{2} \cdot \mathbf{r} \right)$$

and

$$\mathbf{q}_2 = \left(\cos \frac{2\pi - \theta}{2}, \sin \frac{2\pi - \theta}{2} \cdot (-\mathbf{r}) \right)$$

$$= \left(-\cos \frac{\theta}{2}, -\sin \frac{\theta}{2} \cdot \mathbf{r} \right)$$

define the same rotation. Which of the two quaternions is used does not matter, but one has to be careful when measuring the distance of two rotations (e.g. for describing rotation residual errors) by the distance between quaternions. However, in contrast to the axis/angle representation where $\mathbf{R} = \mathbf{I}_{3 \times 3}$ results in an undefined rotation axis \mathbf{r} , the corresponding quaternion is defined and equals 1 (i.e. (1, 0, 0, 0)).

The computation of a quaternion from a rotation matrix is carried out using the axis/angle representation as shown in Equation (30). The computation of a rotation matrix \mathbf{R} from a quaternion can be done as follows (Faugeras 1993):

$$\mathbf{R} = \begin{pmatrix} r_1 & r_2 & r_3 \end{pmatrix}, \quad (32)$$

where

$$\mathbf{r}_1 = \begin{pmatrix} q_r^2 + q_1^2 - q_2^2 - q_3^2 \\ 2(q_1 q_2 + q_r q_3) \\ 2(q_1 q_3 - q_r q_2) \end{pmatrix},$$

$$\mathbf{r}_2 = \begin{pmatrix} 2(q_1 q_2 - q_r q_3) \\ q_r^2 - q_1^2 + q_2^2 - q_3^2 \\ 2(q_2 q_3 + q_r q_1) \end{pmatrix},$$

$$\mathbf{r}_3 = \begin{pmatrix} 2(q_1 q_3 + q_r q_2) \\ 2(q_2 q_3 - q_r q_1) \\ q_r^2 - q_1^2 - q_2^2 + q_3^2 \end{pmatrix}. \quad (33)$$

A.4. Discussion

The different representations for rotation matrices introduced in this section model 3×3 rotation matrices, which have nine elements but only three degrees of freedom with less than nine parameters.

Due to the many disadvantages of Euler and Cardan angle representations, using these should be avoided whenever possible. Axis/angle is a minimal parameterization having the drawback that, for rotations with small angles, the rotation axis is not well defined. This problem does not occur when using unit quaternions. However, quaternions are a non-minimal parameterization as they have four elements with three degrees of freedom. Both axis/angle and quaternions are non-unique, i.e. there are always two different representations for the same rotation. However, this causes only slight problems in practice which are not comparable to the non-uniqueness of Cardan and Euler angle representations. Therefore, due to their advantages, the rotation representation of choice for estimating rotation is either axis/angle or quaternions, depending on the application.

An interesting result supporting this conclusion is based on the fact that quaternions as well as axis/angle representation are a so-called *fair parameterization* of 3D rotations, while Cardan and Euler angles are not (Hornegger and Tomasi 1999). A parameterization is fair if it does not introduce more numerical sensitivity than is inherent to the problem itself, which is guaranteed if any rigid transformation of the space to be parameterized results in an orthogonal transformation of the parameters.

References

- Andreff, N. (1997). Towards the Embedding of On-Line Hand-Eye Calibration into Visual Servoing. In *IEEE/RSJ/INRIA Workshop On New Trends in Imagebased Robot Servoing*, pp. 64–70.
- Andreff, N., Horaud, R., and Espiau, B. (1999). On-line Hand-Eye Calibration. In *Proceedings of Second International Conference on 3-D Digital Imaging and Modeling (3DIM'99)*, Ottawa, pp. 430–436.
- Andreff, N., Horaud, R., and Espiau, B. (2001). Robot Hand-Eye Calibration Using Structure from Motion. *International Journal of Robotics Research*, **20**: 228–248.
- Aron, M., Simon, G., and Berger, M.-O. (2004). Handling Uncertain Sensor Data in Vision-Based Camera Tracking. In *International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 58–67. IEEE Computer Society.
- Baez, J. C. (2001). The Octonions. *Bulletin of the American Mathematical Society*, **39**(2): 145–205.
- Chen, H. (1991). A Screw Motion Approach to Uniqueness Analysis of Head-Eye Geometry. In *Proceedings of*

- Computer Vision and Pattern Recognition (CVPR)*, Maui, Hawaii, pp. 145–151. IEEE Computer Society Press.
- Chou, J. C. K. and Kamel, M. (1991). Finding the Position and Orientation of a Sensor on a Robot Manipulator Using Quaternions. *International Journal of Robotics Research*, **10**(3): 240–254.
- Conway, J. H. and Smith, D. A. (2003). *On Quaternions and Octonions: Their Geometry, Arithmetic, and Symmetry*. A K Peters, Ltd.
- Daniilidis, K. (1999). Hand-Eye Calibration Using Dual Quaternions. *International Journal of Robotics Research*, **18**(3): 286–298.
- Daniilidis, K. (2001). Using the Algebra of Dual Quaternions for Motion Alignment. In G. Sommer (Ed.), *Geometric Computing with Clifford Algebras*, Chapter 20, pp. 489–500. Springer-Verlag.
- Dornaika, F. and Chung, R. (2003). Stereo Geometry from 3D Ego-Motion Streams. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, **33**(2): 308–323.
- Fassi, I. and Legnani, G. (2005). Hand to sensor calibration: A geometrical interpretation of the matrix equation $AX=XB$. *Journal of Robotic Systems*, **22**(9): 497–506.
- Faugeras, O. (1993). *Three-Dimensional Computer Vision: A Geometric Viewpoint*. Cambridge, MA: MIT Press.
- Faugeras, O. and Luong, Q.-T. (2001). *The Geometry of Multiple Images, The Laws that Govern the Formation of Multiple Images of a Scene and Some of Their Applications*. Cambridge, Massachusetts: MIT Press.
- Hamilton, W. R. (1844). On a New Species of Imaginary Quantities Connected with a Theory of Quaternions. *Proceedings of the Royal Irish Academy*, **2**: 424–434.
- Hamilton, W. R. (1847). On Quaternions. *Proceedings of the Royal Irish Academy*, **3**: 1–16.
- Hamilton, W. R. (1848). Researches Respecting Quaternions: First Series. *Transactions of the Royal Irish Academy*, **21**: 199–296.
- Hartley, R. I. and Zisserman, A. (2003). *Multiple View Geometry in Computer Vision* (2nd ed.). Cambridge: Cambridge University Press.
- Horaud, R. and Dornaika, F. (1995). Hand-Eye Calibration. *International Journal of Robotics Research*, **14**(3): 195–210.
- Hornegger, J. and Tomasi, C. (1999). Representation Issues in the ML Estimation of Camera Motion. In *Proceedings of the International Conference on Computer Vision (ICCV)*, Corfu, Greece, pp. 640–647. IEEE Computer Society Press.
- Kuipers, J. B. (1999). *Quaternions and Rotation Sequences: A Primer with Applications to Orbits, Aerospace, and Virtual Reality*. Princeton University Press.
- Linde, Y., Buzo, A., and Gray, R. (1980). An Algorithm for Vector Quantizer Design. *IEEE Transactions on Communications*, **28**(1): 84–95.
- Luong, Q.-T. and Faugeras, O. (1993). Self-Calibration of a Stereo Rig from Unknown Camera Motions and Point Correspondences. Technical Report 2014, INRIA.
- Luong, Q.-T. and Faugeras, O. (2001). Self-Calibration of a Stereo Rig from Unknown Camera Motions and Point Correspondences. In A. Gruen and T. Huang (Eds.), *Calibration and Orientation of Cameras in Computer Vision*, pp. 195–229. Berlin, Heidelberg: Springer-Verlag.
- Schmidt, J. (2006). *3-D Reconstruction and Stereo Self-Calibration for Augmented Reality*. Berlin: Logos Verlag.
- Schmidt, J., Vogt, F., and Niemann, H. (2003). Robust Hand-Eye Calibration of an Endoscopic Surgery Robot Using Dual Quaternions. In B. Michaelis and G. Krel (Eds.), *Pattern Recognition, Proceedings of the 25th DAGM Symposium*, Volume 2781 of *Lecture Notes in Computer Science*, Berlin, Heidelberg, pp. 548–556. Springer-Verlag.
- Schmidt, J., Vogt, F., and Niemann, H. (2004). Vector Quantization Based Data Selection for Hand-Eye Calibration. In B. Girod, M. Magnor, and H.-P. Seidel (Eds.), *Vision, Modeling, and Visualization 2004*, Stanford, USA, pp. 21–28. Aka/IOS Press, Berlin.
- Schmidt, J., Vogt, F., and Niemann, H. (2005). Calibration-Free Hand-Eye Calibration: A Structure-from-Motion Approach. In W. G. Kropatsch, R. Sablatnig, and A. Hanbury (Eds.), *Pattern Recognition, 27th DAGM Symposium*, Volume 3663 of *Lecture Notes in Computer Science*, pp. 67–74. Springer-Verlag, Berlin, Heidelberg, New York.
- Shi, F., Wang, J., and Liu, Y. (2005). An Approach to Improve Online Hand-Eye Calibration. In J. S. Marques, N. P. de la Blanca, and P. Pina (Eds.), *IbPRIA (1)*, Volume 3522 of *Lecture Notes in Computer Science*, pp. 647–655. Springer.
- Shiu, Y. C. and Ahmad, S. (1989). Calibration of Wrist-Mounted Robotic Sensors by Solving Homogeneous Transform Equations of the Form $AX = XB$. *IEEE Transactions on Robotics and Automation*, **5**(3): 16–29.
- Tsai, R. Y. (1987). A Versatile Camera Calibration Technique for High-Accuracy 3D Machine Vision Metrology Using Off-the-Shelf TV Cameras and Lenses. *IEEE Journal of Robotics and Automation*, **RA-3**(4): 323–344.
- Tsai, R. Y. and Lenz, R. K. (1989). A New Technique for Fully Autonomous and Efficient 3D Robotics Hand/Eye Calibration. *IEEE Transactions on Robotics and Automation*, **5**(3): 345–358.
- Vogt, F. (2006). *Augmented Light Field Visualization and Real-Time Image Enhancement for Computer Assisted Endoscopic Surgery*. Tönning: Der Andere Verlag.
- Wang, C. (1992). Extrinsic Calibration of a Vision Sensor Mounted on a Robot. *IEEE Transactions on Robotics and Automation*, **8**(2): 161–175.
- Watt, A. and Watt, M. (1992). *Advanced Animation and Rendering Techniques*. Addison-Wesley.
- Zhang, J., Shi, F., and Liu, Y. (2005). An Adaptive Selection of Motion for Online Hand-Eye Calibration. In S. Zhang and R. Jarvis (Eds.), *Australian Conference on Artificial Intelligence*, Volume 3809 of *Lecture Notes in Computer Science*, pp. 520–529. Springer.

- Zhang, Z. (1998). A Flexible New Technique for Camera Calibration. Technical Report MSR-TR-98-71, Microsoft Research.
- Zhang, Z. (2000). A Flexible New Technique for Camera Calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **22**(11): 1330–1334.
- Zhao, Z. and Liu, Y. (2006). Hand-Eye Calibration Based on Screw Motions. In *Proceedings of International Conference on Pattern Recognition (ICPR)*, Volume 3, pp. 1022–1026. IEEE Computer Society.