Eye in Hand Robot Calibration

van Albada, G.D.; Lagerberg, J.M.; Visser, A.

DOI
10.1108/EUM0000000004168

Publication date
1994

Document Version
Final published version

Published in
Industrial Robot

Eye in Hand Robot Calibration

G.D. van Albada, J.M. Lagerberg and A. Visser

Introduction

Robot calibration can serve various purposes. The static and dynamic positioning accuracy of robots have become the bottle-neck for the introduction of off-line programming techniques. These techniques require the robot's position to be predicted with sufficient accuracy. Robot calibration will improve the positioning accuracy. Another important application of robot calibration is its use as a diagnostic tool in robot production and maintenance. Inaccuracies and wear in specific components of the robot may be identified using accurate measurements and a suitable kinematic model.

A large number of robot measurement systems are now available commercially, each with its own range of applicability and its own requirements. Yet, there is a dearth of systems that are portable, accurate and low cost.

In this article we present a simple measuring system that may fill this gap. It is based on a camera in the robot hand plus a known reference object in the robot workspace. Our prototype allows us to measure a robot's position and orientation in a volume of $1m^3$, with an accuracy of 0.20mm and 2.0 minutes of arc. The work described in this article was performed for CAR, ESPRIT project nr. 5220[1].

The research described in this article was partly funded by the EEC through ESPRIT II project 5220 "CAR". Stephen Kyle of Leica, and Steve Albright, Klaus Schröer and Michael Grethlein of IPK have contributed significantly to the development of the system through their expert advice and support.

Design and Implementation of the Measuring System

The measuring system discussed in this article was designed on the basis that the system:

- should be able to provide (static) position and orientation data compatible with the repeatability of current robot systems;
- should be low cost, portable, easy to operate by non-expert personnel and sufficiently robust to be used in an average industrial environment;
- need not be able to work in the full workspace of the robot, but should be able to measure a large number of poses in a limited volume.

On the basis of these criteria, various measuring techniques have been examined, as described in an initial CAR report[2, pp. 101-23]. This study indicated that the system should also be self calibrating; that it should work on the basis of optical sensors, that it should contain no moving parts and a minimum number of specially manufactured components.

The most obvious solution that promised to approach the required accuracy proved to be a system based on a single camera in the robot hand, plus a specially designed, passive, flat reference object (reference plate) positioned in the robot workspace (see photograph).

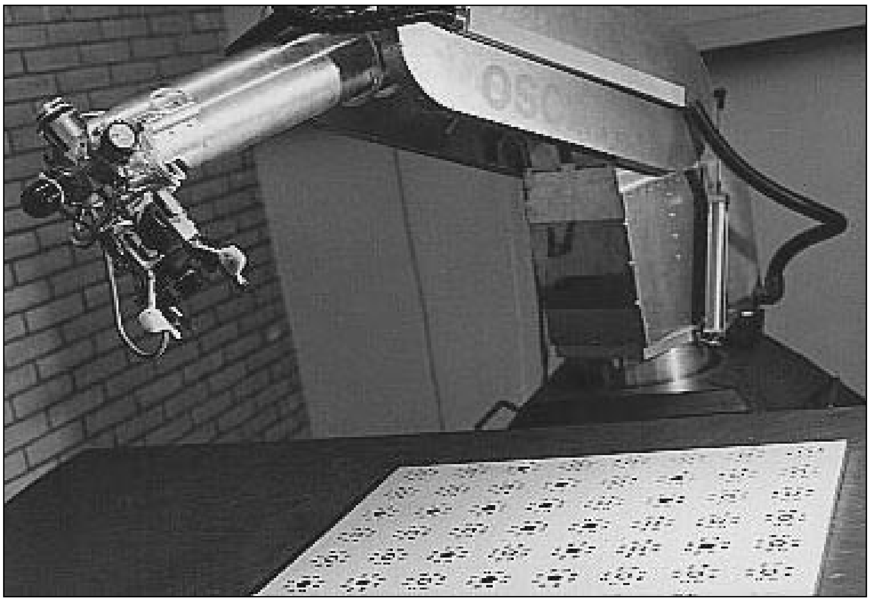

The reference plate as seen by the camera

To be able to self-calibrate the camera at least a few images are required that contain a large number of measurable positions covering the entire image. The larger the reference plate is, the better the camera can be calibrated. Also, position measurements with an accuracy of 0.1mm imply a reference object size exceeding 30cm if an orientation accuracy of 1 minute is to be attained. Attaching such a large object to the robot flange and placing cameras in the workspace may be a problem. The camera parameters were found to vary slightly from run to ron. It is therefore necessary to recalibrate these every time.

We have implemented a prototype version of the measuring system using a simple off-the-shelf camera. The reference plate consists of a blank aluminium plate with a black pattern of rings printed on it. Results, presented in this article, show that the accuracy of the camera system, due to its self-calibration capacity, is generally sufficient for robot calibration.

Measuring Procedure

The measuring procedure begins with the selection of the model parameters of the robot that need to be redetermined. Using this set of model parameters, the pose generation program[2] will generate a set of measurable poses that will allow computation of these parameters.

Using these poses, a robot program is generated which directs the robot along a path containing these positions. At each measuring pose the robot stops and one or more images of the reference plate are obtained. The actual joint parameters at the measuring poses may be recorded, if possible, but may otherwise be assumed to be equal to the commanded values.

Next, the images obtained are processed off-line, to obtain the positions of the camera relative to the reference plate, plus the parameters of the camera. This “photogrammetric procedure” is the innovative part of our system. This procedure can be made self-calibrating when a collection of sufficiently different images is available. The poses are obtained by iterating two tasks: the *image-processing* procedure and the *image-reconstruction* procedure. The former tries to recognize and identify the markers on the reference plate and to determine their positions in the image, the latter fits a model that can predict the position of the markers in every image by the computation of the camera positions for each image and the camera parameters. The predictions are fed back to the identification part of the *image-processing* procedure.

Using the calibration procedure developed at IPK Berlin (Schröer)[2, pp. 157-93] the unknown robot parameters, plus the position of the reference plate relative to the robot base, can be derived from these measurements.

Experimental Results

For the CAR project we developed and tested a prototype system based on the principles described in the preceding section. The prototype demonstrates the viability of the approach. In the initial stage we verified our design of the image-reconstruction model with a least-square fit using the singular value decomposition method. This method is helpful for ill-conditioned systems, but involves a lot of computations. After a small refinement of our model very good condition numbers were obtained, and large sets of images could be processed with the classical and far more efficient Gaussian elimination method.

Refinement of the Camera Model

In the first experiment 16 images were taken from different viewpoints. Originally only six camera parameters were estimated: five parameters that describe the geometric transformation between the optical centre and the image plane and the third order distortion coefficient *k*.

The experiment showed that in less than ten steps the system converged to a solution. The vector field of the residuals (vector differences between the measured and the computed image points) showed that the distortion was not modelled in an optimal way. To give a clear picture of the vector field, we divided the area of the image in 10 by 10 parts, in which we added the residuals. The result is given in Figure 1. At the corners of the image the residuals are quite large compared to the inner residuals. Also, they point in the same direction, indicating the presence of residual distortion effects. Therefore, we added the co-ordinates of the centre of the distortion as parameters to the system, and the fifth order term of the radial lens distortion. The vector field of the residuals shows that the effect of the distortion at the corners is diminished. See Figure 2 in which the new sum vector of the residuals in each part is plotted (in order to also show the smaller residuals, the relative length of the vectors has been increased by a factor 10 relative to those in Figure 1). To suppress the effects of outliers, an automatic rejection procedure for points that resulted in an unusually bad fit was incorporated.

The following camera parameters were obtained. The third order term was negative, so we can speak of pin-cushion distortion. The centre of the distortion and the projection of the optical centre on the image were very

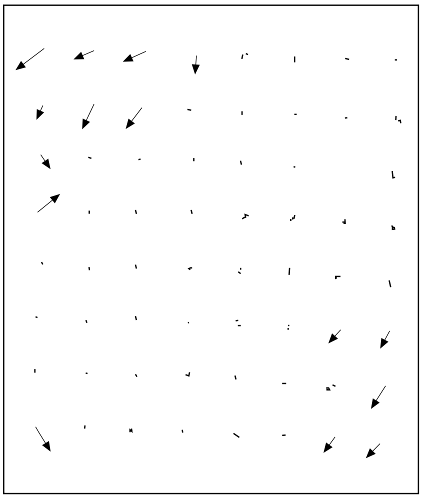

Figure 1. Residuals between Calculated and Measured Image Points, with a Six-parameter Camera Model (20x)

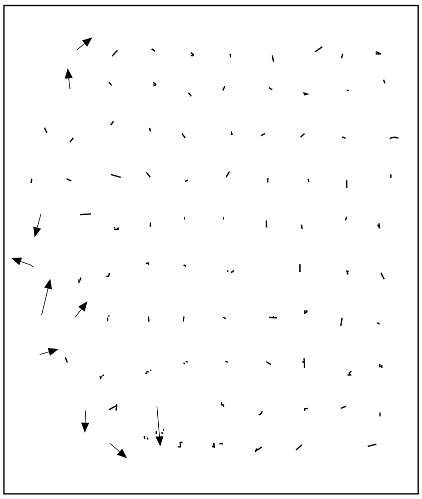

Figure 2. Residuals between Calculated and Measured Image Points, with a Nine-parameter Camera Model (200x)

close to each other. This clearly shows that the lens of the camera is symmetric with respect to its centre, and that the chip is almost perpendicular to the optical axis.

Measurements on the OSCAR Robot

In the second experiment 124 images were obtained using the "OSCAR" robot at the University of Amsterdam. For this experiment we computed a number of error quantifiers. One significant quantity is the accuracy to which the measured points can be fitted by the model. In our case a rms fitting error per measured point of 0.11 pixel was found. From the rms fitting error per image, plus the assumption that the remaining errors per measured point are uncorrelated, we computed the expected error covariance matrices for the position and orientation and for the camera parameters and from those the expected errors in the measurements.

The results until now show a formal rms accuracy of 0.10mm and 1.0 minutes of arc, with a number of significantly worse points (Figures 3 and 4). However, the calculated formal accuracy of the measurements depends on a number of assumptions, such as the constancy of the camera properties throughout a sequence of measurements and the independence of the residual errors, i.e. it does not take into account systematic and correlated errors.

Additional experiments indicate that the formal error estimates do not take into account some very real error sources and therefore tend to be overly optimistic. In the following section we will give a short overview of possible sources of these systematic errors. Further analysis of the principal error sources has resulted in an extension of the number of camera parameters to 11. We now also solve for distortions in the reference plate.

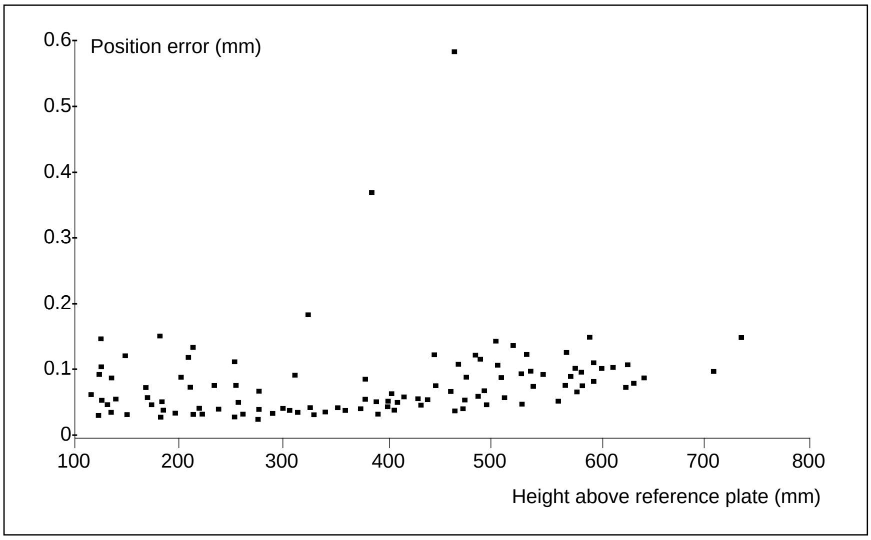

Figure 3. The Estimated rms Position Errors $(\sigma_x^2 + \sigma_y^2 + \sigma_z^2)$ for a Collection of 124 Images, as a Function of the *z*-height of the Camera. Notice that the Position Error Increases with the Height above the Reference Plate

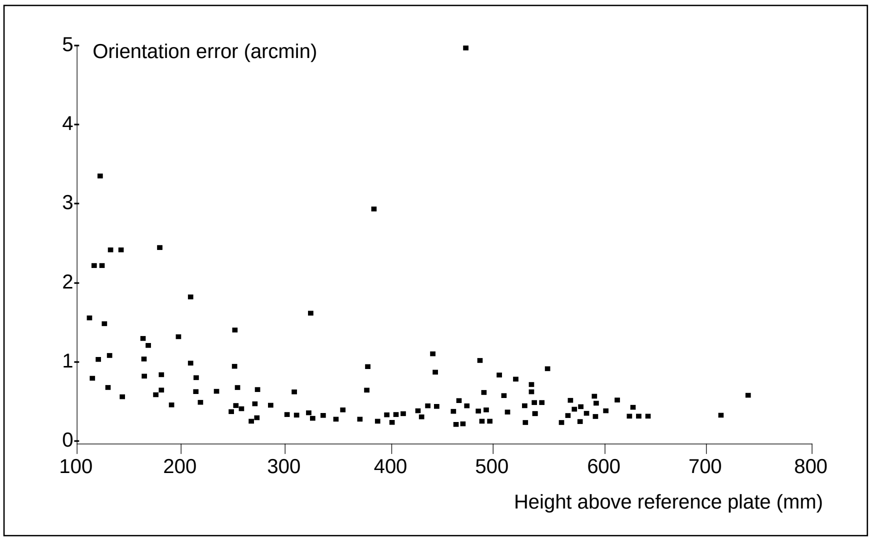

Figure 4. The Estimated rms Angle Errors $(\sigma_\alpha^2 + \sigma_\beta^2 + \sigma_\gamma^2)$ for a Series of 124 Images, as a Function of the *z*-height of the Camera. For the Range of *z*-heights Shown, the Orientation Error is Mostly Independent of *z*. It Tends to Increase for Small Values of *z*, as Fewer Reference Points Will be Visible

Discussion

In this section we will discuss various possible sources of measurement errors, the most important of which we have taken into account. How further improvements in the measuring procedure can be obtained is indicated. Images of the reference plate are obtained with a camera consisting of a lens, a mounting, a CCD and a frame grabber. Each of the components in this procedure contributes its own set of errors. These errors either have to be minimized by adopting a suitable measuring procedure, or have to be modelled in order to remove their contribution.

Here the contributions of each component will be discussed in term.

The Reference Plate

The measurement reference plate consists of a white flat plate with a large number of black, ring-shaped markings in a regular, grid-shaped pattern, as illustrated in the photo. A flat reference object was selected because it can be more easily constructed and maintained than a three-dimensional object, although the latter is, in principle, better for photogrammetric applications. The accuracy of the reference plate must well exceed the desired measuring accuracy.

It is not too difficult to manufacture a sufficiently accurate plate. However, some problems should be taken into account, specifically the effect of temperature changes. Typical expansion coefficients of solids are in the order of 10^{-5} C^{-1}. So, our reference plate (0.6m by 0.5m) will expand about 0.005 mm per degree C; i.e. changes in the ambient temperature in the order of a few degrees C will result in scale changes comparable to the desired measurement accuracy.

If a sufficiently large number of images is obtained, errors of the positions of the markings can be recognized by our measuring procedure, and can in principle also be measured.

The Mounting

The mounting is important as it fixes the position of the lens relative to the detector, in our case a CCD. The mounting can allow the distance of the lens to the detector to be changed (focusing), the aperture to be changed and the lens to be removed from the camera.

Each of these options implies a mechanical change to the optical system, leading to non-reproducible variations in its properties. For that reason, a fixed focus, fixed aperture lens is preferred.

The Lens

The camera lens will be used to produce images of the reference plate over a range of object distances from 0.2m and 1.4m. The images should be as sharp as possible to obtain good measurements.

Lenses are known to display a large variety of imaging errors, affecting the quality of the image. Note that the achieved measuring accuracy of the system is significantly better than one could expect from the major imaging effects.

This is achieved by a combination of sub-pixel interpolation in the grey-scale image in the computation of the position for each marker, the use of information from a large number of pixels for each marker, and the use of information derived from up to about 700 markers in each image.

The principal error types are the following:

- Defocusing and depth-of-field. With a fixed focus a sharp image is produced only for an object at a precise distance. The further the object is removed from that distance, the more the image is smeared.

  The distance range over which this effect stays within acceptable bounds is referred to as the depth-of-field. The effect is always present, but its effect on the image can be reduced by using wide-angle lenses and stopping down the aperture. For our system the effect maximally is in the order of 15′.
- Diffraction. Due to the wave nature of light, there is a limit on the degree to which the lens can be stopped down. For small aperture diameters a diffraction pattern becomes visible: the Airy disk.

  The radius of the Airy disk for our system is about 3′. Optimally, one should choose the aperture of the lens so that the loss of sharpness due to the depth of field and diffraction are approximately equal.
- Distortion. The pin-cushion distortion in our system gives a significant contribution. We modelled it with third and fifth order radial terms in the least-square fit of our image-reconstruction procedure.
- Astigmatism, image plane curvature and coma will affect the sharpness of the image in the corners; these effects are reduced by choosing a small aperture.
- Chromatic aberration can be significant and is best reduced by using an optically flat colour filter (about 10nm band pass).

“The measuring system is self-calibrating”

Off-the-shelf lenses, such as the lenses used in our experiments, are usually optimized to yield an image that is pleasing to the eye. For ultimate performance, a specially designed lens should be used, making use of the specific trade-offs allowed for photogrammetry, but this will increase the cost of the system.

The CCD and the Frame Grabber

The image produced by the lens must be detected using a CCD or a similar (rectangular) array of detector elements.

The output of the detector elements is usually converted to a standard video signal, which is digitized using a frame grabber. A problem with this procedure is that the outputs of adjacent detector elements can be mixed in an unpredictable fashion in the output signal. Synchronization errors between the camera and the frame grabber can also lead to geometric distortions that vary from line to line in the image. For these reasons, the use of a “pixel synchronous” detection system is preferred. The imperfect synchronization between camera and frame-grabber appears to be the main source of errors in our current system.

Conclusions

In this article we have presented a low-cost method, based on photogrammetry, to obtain measurements for the calibration of robot systems. The measuring system is self-calibrating.

The method has been implemented and tested and provides promising results for practical application. The components used are relatively inexpensive, and can easily be combined to yield a portable system.

As most of the data processing has been highly automated, such a system will be usable by non-expert personnel.

By combining the video camera with a fast frame grabber and recording system, or alternatively with a video recorder, dynamic measurements should be obtainable.

The relative locations and orientations of two robots in a work cell can be found by placing the reference plate between the robots and calibrating both robots with that common reference.

Note and Reference

1. In CAR the following companies and institutes co-operated: Fraunhofer-Institut für Produktionsanlagen and Konstruktionstechnik (IPK Berlin, prime contractor), Leica (UK) Ltd., University of Amsterdam, Department of Computer Systems, TGT (Ireland), KUKA Schweißanlagen and Roboter GmbH, Volkswagen AG. ESPRIT projects are 50 per cent funded by the EU.
2. Bernhardt, R. and Albright, S., *Robot Calibration*, Chapman & Hall, London, 1993.

G.D. van Albada, J.M. Lagerberg and A. Visser are with the Faculty of Computer Science, University of Amsterdam, Amsterdam, The Netherlands. Tel: +31 20 525 7463; Fax: +31 20 525 7490.