A portable measuring system for robot calibration

G.D. van Albada, A. Visser, J.M. Lagerberg and L.O. Hertzberger
Faculty of Mathematics and Computer Science
University of Amsterdam

95ME073

Abstract

Currently, much production time is lost in teaching and re-teaching robot programs. New techniques that avoid this, such as off-line programming, require an accurate kinematic model of the robot. To construct this model, measurements of the robot are required.

At the University of Amsterdam a self-calibrating measuring system has been developed, using a camera mounted in the robot hand and a reference object. Analysis of a single image of the reference object allows us to obtain an estimate of the position and orientation of the viewing point. A much better estimate is obtained by analysing a series of images, as the essential geometrical and optical camera parameters can be derived from the redundancy in the measurements.

Our experimental results show that we can compute the pose of the camera for the different images with an accuracy in the order of 0.15 mm and 1.5 arc minute. This is usually adequate for robot calibration.

Introduction

The positioning accuracy of robot systems can be characterised in a number of different ways. The repeatability defines the accuracy to which the end-effector re-attains the same position (and orientation) in its working volume, when the robot joint values are repeated. The repeatability tends to be in the order of a few tenths of a millimetre and a few minutes of arc. It defines the limit to the accuracy attainable with a particular robot using only its internal position sensors. The kinematic positioning accuracy is the accuracy to which the robot attains a given position in its working volume when the robot controller is commanded to go to that position. The robot controller requires an inverse kinematic model of the robot to compute the corresponding joint values. It is found in practice that these models often lack in accuracy, the more so as the actual parameters of the robot may change due to wear, minor accidents and repairs. The kinematic positioning accuracy often is not much better than several millimetres and can be as bad as a few cm. A third accuracy measure is the dynamic positioning accuracy, which defines the accuracy at which the robot end-effector can be made to follow a specified path (at a given speed). The dynamic positioning accuracy depends on a large number of system properties and will not be considered further in this paper.

Currently, most robots are programmed through teaching, i.e. the robot end-effector is moved to the required positions for the programme, and the joint-encoder values are recorded. In this way the programmed points will be reached with only the repeatability error, and no (inverse) kinematic model is needed. Teaching is, however, a time-consuming procedure that needs to be executed with the robot removed from the production-line. Every robot in the production-line will be slightly different, and therefore will have to be taught its programme separately. After repair or minor accidents, the robots properties will be somewhat changed, necessitating re-teaching.

When sufficiently accurate kinematic models are available, robot programmes can be created off-line, and taught programmes can be adapted to execute correctly on another robot.

The CAR project (ESPRIT 5220)[1] was concerned with all aspects of the generation and maintenance of such kinematic models. Constructing a kinematic model is a multi-step process. First, a model describing all the links, actuators and their parameters must be constructed. Next, those parameters that need to be identified must be specified, and a measuring programme that allows identification of these parameters must be constructed. In this programme, the robot will be moved through a sequence of poses where the position (and sometimes the orientation) of the end-effector will be measured with high accuracy. Finally, the measurements will be used in a parameter-identification procedure. In this paper, we will only be concerned with the measuring procedure itself. Procedures for the construction of the kinematic model and the parameter identification can be found in papers by Albright[3] and Schröer[4].

Depending on the purpose for which the robot is being measured – a full identification of all robot parameters, or a partial identification of only some parameters – the range of robot poses that needs to be measured may vary. For a partial identification measurements obtained in only a part of the robot operating volume will generally suffice, especially if those poses can be attained with different robot configurations. Currently, the most commonly used measuring system for robot calibration consists of two or more theodolites, with which end-effector positions can be measured in most of the robot operating volume. These systems are quite accurate and are suitable for full system identification, but they are expensive and require significant expertise of the operator. Within the CAR project a need for a low-cost, portable and easy to use measuring system suitable for partial identifications was felt. In close collaboration with the other project partners, we have developed a prototype for such a system at the University of Amsterdam. In this paper we will first give a short resume of the requirements on the system and their implications for the design. Papers describing various aspects of this system can be found in Van Albada et al. [1, 2].

Requirements and design

As explained in the introduction, we strove to develop a portable, inexpensive and easy-to-use measurement system for partial recalibration of robots. Some of the more specific requirements for this system were:

- An operating volume of at least $(0.5 \text{ m})^3$.
- A measuring accuracy of about 0.1 mm and 1 minute of arc
- Insensitivity to initial positioning errors of at least a few cm.
- Non-contact
- Minimal interaction with robot control program (e.g. no sensor-based position adjustment)
- Robustness and low maintenance.

The low maintenance requirement almost automatically implies that the system should contain no (high precision) mechanical components and that it must be self-calibrating. The other requirements can best be met by using an optical system, preferably based on the use of a camera and a reference object. For a camera-based system, there is still a choice between a monocular, or a multi-camera system, and a choice between placing the camera on the robot, or the reference object. Each option has its specific advantages and draw-backs. Placing the camera on the robot has the disadvantage that the power and the signal cables must somehow be led from the robot-hand to the recording device (frame-grabber). As the robot has to go through a large range of movements, this is a problem of significant complexity.

---

[1] In CAR the following companies and institutes co-operated: Fraunhofer-Institut für Produktionsanlagen und Konstruktionstechnik (IPK Berlin, prime contractor), Leica (UK) Ltd., University of Amsterdam, Dept. of Computer Systems, TGT (Ireland), KUKA Schweißanlagen und Roboter GmbH, Volkswagen AG. ESPRIT projects are 50% funded by the EEC.

On the other hand, placing the camera at a fixed position, and the reference object in the robot hand, creates problems in the measurement of the orientation of the robot hand. Orientation measurements are essentially measurements of the difference in position between two ends of the reference object. Given a position measuring accuracy of 0.1 mm, a reference object of about 50 cm by 50 cm is required to attain an orientation accuracy of 1 minute of arc. Attaching such a large object to the robot hand can be a problem and it can limit the range of attainable poses of the robot. Deformation of the reference object during motion of the robot is a significant risk that is difficult to evaluate or prevent.

Obtaining accurate measurements with a monocular system requires a wide angle camera working at close range. The attainable accuracy with a multi-camera system should be significantly better than with a monocular system, but a multi-camera system is inherently more complicated and requires a more complex set-up and self-calibration procedure than a monocular system. Also, placing two or more sufficiently separated cameras on the robot can lead to a variety of problems.

For the reasons stated above we have opted for a single wide angle camera on the robot and a reference object fixed in the workspace.

Principle of operation

The image of a known object, obtained with a camera with known properties, can contain enough information to compute the position of the camera relative to the object. When a sufficiently varied set of images is obtained, most of the properties of the camera and the reference object can also be determined. This technique of measuring spatial relationships using images is known as photogrammetry.

The following paragraphs show what information is involved and are intended to make plausible that such a procedure is indeed possible. Let's assume that we have a left-handed, orthonormal co-ordinate system[2] (X, Y, Z) tied to the reference object and a 2-dimensional orthonormal co-ordinate system (U, V) tied to the image. We can then describe the imaging process by a pair of non-linear functions u and v, so that a point at $\vec{\mathbf{x}}_o = (x_o, y_o, z_o)$ will be imaged at (u_o, v_o), with:

$$u_o = u\,(\vec{\mathbf{x}}_o, \mathbf{P}, \mathbf{M}), \qquad (1a)$$

$$v_o = v\,(\vec{\mathbf{x}}_o, \mathbf{P}, \mathbf{M}), \qquad (1b)$$

where $\mathbf{P} = (x_c, y_c, z_c, \alpha_c, \beta_c, \gamma_c)$, is the six-tuple giving the position and orientation of the camera (the precise definition of the angle parameters α, β, and γ will not concern us here), and $\mathbf{M}$ is a tuple of all N_{cam} parameters in the model used for the camera. $\mathbf{M}$ is assumed to be the same for all images. When we obtain N_{pos} images with camera poses $\mathbf{P}_j$ (j = 1, .. N_{pos}), using a reference object with N_{mark} measurable markers at positions $\vec{\mathbf{x}}_{o,i}$ (i = 1, .. N_{mark}), we have a total of N_{total} unknown quantities:

$$N_{total} = 6\,N_{pos} + 3\,N_{mark} + N_{cam}$$

Of these N_{total} values, seven cannot be determined in principle, but must be defined by additional equations. These quantities relate to the absolute position and orientation of the reference object, which cannot be measured by any known device and to its scale, which can only be measured by comparison to some other standard.

---

[2] This is the usual everyday type of co-ordinate system.

Every measured marker position in an image provides two equations. If we assume the **M** and all the $\vec{\mathbf{x}}_{o,i}$ are known, measuring three markers per image would appear to suffice to derive the position the camera. However, due to the non-linearity of the system of equations, at least four measured markers per image are needed to remove ambiguities in the solution. Each image yields two equations for each marker in the image; each marker should, therefore, appear in at least two images taken from two sufficiently different directions to compute its position vector $\vec{\mathbf{x}}_{o,i}$. In practice the above constraints on the number of images and the numbers of marker per image can easily be met. Therefore, all N_{total} - 7 unknowns can be solved, and the system becomes essentially self-calibrating.

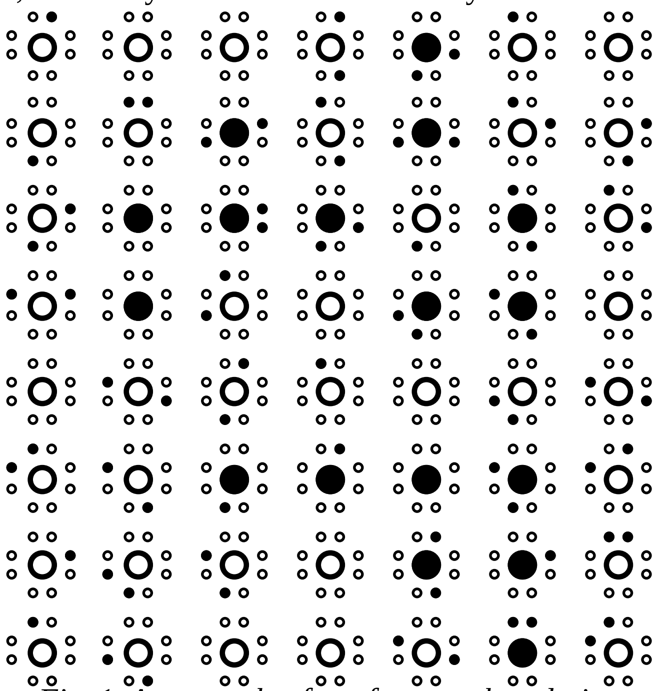

Fig. 1. An example of a reference plate design

In our measuring procedure, we use a reference plate on which we have printed 504 circular markers (Figure 1). Images taken from very close to the reference plate will sometimes show as few as 20 marker images; images taken from a larger distance will show nearly all. The plate has been designed so that even small groups of markers can nearly always be recognised from their appearance. As the calibration of a robot requires a large number of positions to be measured, obtaining a sufficiently large set of images to allow self-calibration of the system is not a problem. Our measurement series generally contain between 40 and 100 images, with an average of some 200 markers measured per image.

We solve the non-linear system of equations by taking initial estimates of the $\vec{\mathbf{x}}_{o,i}$ and the **M**. Using these and the positions of the markers identified in the image, an initial estimate of the $\mathbf{P}_j$ can be obtained by correcting the measured positions for the distortions of the camera and computing the homogeneous transformation from the $\vec{\mathbf{x}}_{o,i}$ to the corrected image positions. In this way an initial solution sufficiently close to the final solution is obtained that we can iterate to the final solution in a limited number of steps, using the linearised equations. The large number of markers per image improve the attainable accuracy, allow us to judge the quality of the model fit, and leave room to reject outliers.

The camera model

The choice of an appropriate camera model is essential for the measuring procedure to produce accurate and reliable results. Apparently small, but systematic discrepancies between the model and the real camera can lead to unacceptable errors in the measured position.

In first approximation, almost any camera can be described by a pin-hole, or camera-obscura model. This model describes in a straightforward, geometric manner where a point on the object will be projected in the image plane. Homogeneous transformations provide a popular way to describe this projection. In most cases, the actual image formation process can be described as a series of perturbations of this idealised process. These perturbations describe the purely geometric differences with the ideal camera, differences affecting the brightness of the image, and an effect influencing both brightness and geometry. The camera model is used in different ways in various stages of our data reduction procedure.

1. Geometric distortion. This describes where the actual image point lies relative to its idealised position. For wide-angle lenses this is quite a large effect near the edges of the image. Geometric distortion depends primarily on the location of the image point in the image. There is also a small effect depending on the parallax (the inverse of the distance) of the object. The geometric distortion effect is taken into account in the image recognition and the model fitting phases of the reduction process. For one of our lenses, we found that the distortion can be described quite well as a radial displacement relative to a point near the centre of the image. The magnitude of the displacement depends on the third and fifth order of the radial distance to that point. The proportionality constants depend slightly on the parallax of the object. For another lens this description did not suffice. The remaining distortion was modelled as a “displacement field,” with additional small displacements depending on the position in the image only.
2. Vignetting and non-uniform illumination. These effects influence the brightness of the image. In combination with an extended PSF (see 3, below) they can lead to an apparent shift in the position of an object in the image. In our procedure, where we measure the position of the markers by tracing the edges, vignetting and illumination effects must be measured and corrected for. This is accomplished using standard image-processing techniques in the first steps of the image processing procedure.
3. The point-spread function (PSF). Light from a single point on the object will be distributed over an area in the image. The PSF depends on the position in the image, on the lens-aperture (which should be kept fixed), on the parallax of the object, and on the colour of the light. For a CCD camera there are a number of additional contributions due to the sampling of the image, the shift-out procedure and the transmission to the frame-grabber/digitiser. Modelling and fitting the PSF is a complex and expensive procedure. It can be avoided as long as the PSF is sufficiently small and symmetric. In our set-up, reducing the effect of the PSF is accomplished by using a small aperture (*f*/11) and a careful design of the measuring procedure.

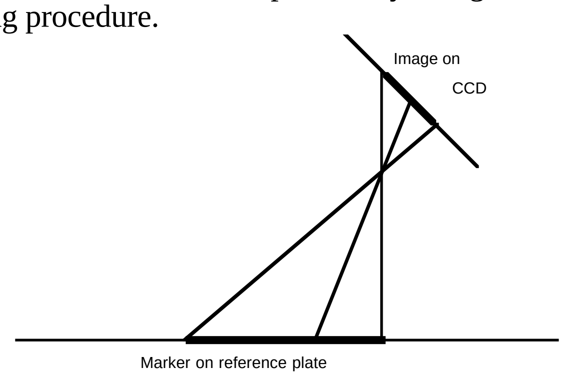

Fig. 2 - The centre of the marker image is not the image of the marker centre

Using images to measure the position of objects implies deriving a model that fits the images as accurately as possible. The most comprehensive way of doing this, is to make a model reproducing the observed intensities in the image. This requires modelling of the illumination, the vignetting, the PSF and the camera distortion.

For reasons of simplicity and efficiency, we have opted for an approach in which we measure the contours of the markers in the image, and derive the marker positions from these contours. This requires an initial model of the perspective and camera distortion, as projection effects will significantly displace the centre of the marker image relative to the marker (Fig. 2). Care must be exercised to obtain a symmetric and reasonably small PSF, i.e. to obtain sharp images.

In modelling the image formation process, we thus consider only the vignetting and illumination effects (in the image processing phase) and the perspective and geometric distortion effects (in the model fitting phase). It is assumed that the effects of the PSF on the measured positions will mostly cancel out for the circular markers that we use. By using mostly open circles (annuli), we add an extra symmetry: the outer edge appears as a black circle on a bright background, the inner edge as a white circle on a dark background. If the PSF is strongly asymmetric, the two circles would be shifted in opposite directions and appear to lie at a different position. We can verify that this effect is small (when the difference is large, both measurements are rejected) and average out any small residual difference.

Data processing

Obtaining camera positions and orientations from a set of images involves a number of processing steps, some of which are schematically shown in Fig. 3.

Before any image can be obtained, suitable robot poses must be computed for the eventual calibration procedure. These poses must be constrained so that workable images of the reference plate can be obtained. A procedure to generate such poses was created by Schröer et al. at IPK Berlin [2, 13].

After setting up the robot, the image acquisition system and the reference plate, the robot can be sent through the required sequence of poses. As thermal expansion of the various components can have an effect that is easily as large as the desired tolerances, the system must be allowed to attain a stable operating temperature before the measurements are taken. In order to correct for gain and dark-current effects in the CCD, it is desirable to obtain some dark images (lens cap on) and white images (uniform white surface at sufficiently close range to be quite unsharp) at the beginning and end of the measuring sequence. It is important to ensure that sufficiently different poses are obtained that the camera can also be calibrated. Both the orientation of the camera and it position relative to the reference plate should be varied. If the robot has too few degrees of freedom to accomplish this, additional images obtained with the camera off the robot can be used for camera calibration.

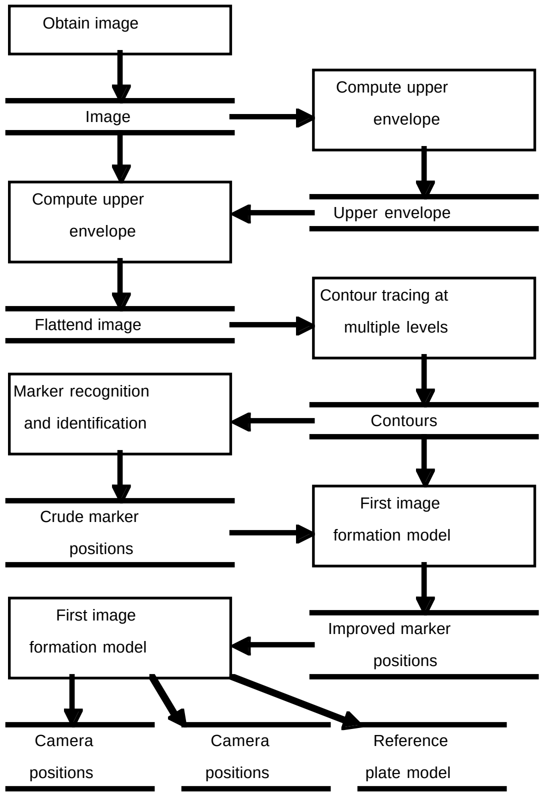

Fig. 3 - The photogrammetric procedure - from image to camera positions and system parameters.

Once all the images have been obtained, the data processing can be performed away from the robot. This data processing involves image processing, image recognition, and model fitting. In the image processing step, transformations are applied to the measured intensities to reduce various effects in the image that could hamper the image recognition procedure. The most important of these is the non-uniform brightness of the image of the reference plate, due to vignetting and illumination effects. This is done by computing an upper envelope to the measured intensities in the image (see e.g. Verbeek et al. [5]), and dividing the image by this upper envelope. In the resulting image, the body of the reference plate appears uniformly bright, and the black markers will be much darker, though not quite uniformly black. By carefully choosing a number of contour levels and tracing these contours in the grey-scale image, the positions of the markers can be determined with an accuracy of 0.03 to 0.1 pixels. In the image recognition step, various tests are applied to retain only good quality marker contours. Information on the relative size of the markers and the pattern of open and filled markers (see Fig. 1) is used to recognise and identify the markers. Using default values for the camera distortion, the projection functions $u()$ and $v()$ (equations 1a,b) can be estimated and the marker positions can be corrected for projection effects. These improved marker positions are used in the final steps of the model fitting phase, where all images are processed together to solve the complete set of unknowns and apply the self-calibrating capability of the system. In this step the equations are linearised by computing the derivatives of $u()$ and $v()$ with respect to their parameters for all measured markers. In this way a sparse Jacobian matrix of up to some 60 000 by 800 elements is obtained used in a least-squares solution for the camera positions and -parameters. The corrections to the marker positions on the reference plate are solved for simultaneously, but in a slightly different way. This final model fitting step is executed iteratively until the system has converged. The residuals in the fits to the observations are used to estimate the accuracy in the computed pose and parameter values.

Performance and possible enhancements

More comprehensive error estimates for the measured poses than those using only the residuals in the fits indicate that accuracies of about 0.15 mm and 1.5 minute of arc rms. can routinely be obtained. Ongoing research indicates that improvements in the image acquisition phase can significantly improve these figures.

When the measurements are to be used for further processing, as is the case in the calibration of robots, a shortcut appears promising in further improving the measurement results. Our analysis shows that the errors in the positions and angles are very strongly correlated. This means that the volume of the error ellipsoid is much smaller than the product of the errors in each of the components, i.e. the measurements are in some sense better than the individual error estimates indicate.

The measured camera positions (and orientations) actually are only used as an intermediate result. The end result that we try to obtain is a - much smaller - set of parameters for the robot model. By concatenating the photogrammetric procedure to the robot calibration procedure and eliminating the camera poses as explicit intermediate results, a better accuracy and an even more stable solution should be obtainable. The disadvantage is that robot calibration program and measuring procedure cannot be used independently anymore.

Conclusions

Robot calibration is rapidly becoming more and more necessary for the successful employment of robots in an increasingly wide range of applications. The widespread use of robot calibrations requires that affordable, easily used equipment becomes available. Equipment based on photogrammetric techniques promises to be a strong contender in this area.

Our prototype measuring system shows that a suitable system, providing an acceptable measuring accuracy, can be built even with very modest investments. It also demonstrates that the measuring process can be highly automated, so that little expert knowledge is required by the user.

There is still room for further improvements in the equipment (cameras with more than twice the resolution of our camera are now coming on the market) and data processing, promising quite good results at a very reasonable cost.

Acknowledgements

The research described in this paper was partly funded by the EU through ESPRIT II project 5220 “CAR”. Stephen Kyle of Leica, Steve Albright, Klaus Schröer and Michael Grethlein of IPK, Rudolf Le Poole of the Leiden Observatory and Arnold Smeulders of the University of Amsterdam have contributed significantly to the development of the system through their expert advice and support.

References

[1] G.D. van Albada, J.M. Lagerberg and A. Visser, Eye in Hand Calibration, Industrial Robot, Vol. 21, no 6 (MCB University Press, 1994) pp. 14-17

[2] G.D. van Albada, J.M. Lagerberg, A. Visser and L.O. Hertzberger, A low-cost pose-measuring system for robot calibration, Robotics and Autonomous Systems (in press)

[3] S.L. Albright, Calibration system for robot production control and accuracy, in R. Bernhardt and S. Albright, Robot Calibration, Eds. (Chapman & Hall, London, 1993), pp. 37-56.

[4] K. Schröer, Theory of kinematic modelling and numerical procedures for robot calibration, in R. Bernhardt and S. Albright, Robot Calibration Eds. (Chapman & Hall, London, 1993), p.157-193.

[5] P.W. Verbeek, H.A. Vrooman, L.J. van Vliet, Low-level Image Processing by Max-Min Filters, Signal Processing 15 (1988) p. 249-258.