3-D Reconstruction and Camera Calibration from Images with known Objects

Gudrun Socher   Torsten Merz   Stefan Posch

Universität Bielefeld, Technische Fakultät,

AG Angewandte Informatik,

Postfach 100131, 33501 Bielefeld, Germany

gudrun@techfak.uni-bielefeld.de

Abstract

We present a method for camera calibration and metric reconstruction of the three-dimensional structure of scenes with several, possibly small and nearly planar objects from one or more images. We formulate the projection of object models explicitly according to the pin-hole camera model in order to be able to estimate the pose parameters for all objects as well as relative poses and the focal lengths of the cameras. This is accomplished by minimising a multivariate non-linear cost function. The only information needed is simple geometric object models, the correspondence between model and image features, and the correspondence of objects in the images if more than one view of the scene is used. Additionally, we present a new method for the projection of circles using projective invariants. Results using both simulated and real images are presented.

keywords: Least-squares model fitting, model-based vision, 3-D reconstruction, camera calibration, projective invariants.

1   Introduction

We present a method for camera calibration and metric reconstruction of the three-dimensional structure of scenes with several, possibly small and nearly planar objects in one process. One or more images from uncalibrated cameras are used. The only information needed is simple geometric models containing descriptions of objects as a set of vertices, edges, and circles, the correspondence between model and image features, and – for more than one view of the scene – the correspondence of objects in the images. Calibration is obtained by direct observation of objects in the scene. Therefore, the method lends itself for on-line calibration of an active vision system. No preceding calibration with a special calibration pattern

---

is necessary and no mechanic or thermic influences resulting from different times of exposure distort the reconstruction.

Our approach is inspired by model-based methods and approaches applying projective geometry. Model-based approaches as presented by [Lowe 91, Goldberg 93] use three-dimensional models of single objects in one image to estimate the object's pose relative to the camera. Special simplified partial derivatives which are difficult to extend for more objects or additional images are used. Preceding camera calibration is necessary if not enough significant model features are detectable in the image or to stabilise the solution. [Mohr *et al.* 93, Boufama *et al.* 93, Faugeras 92] use images from uncalibrated cameras and estimate a projective reconstruction only from known point correspondences. Additional metric information is incorporated in a second step to derive a reconstruction in Euclidean space. Inaccuracies due to noise or false matches introduced in the first step are difficult to correct in the second step where information available from the 3-D scene is taken into account.

[Crowley *et al.* 93] achieve robust results using the objects in the scene for camera calibration. Their approach holds for single objects and shows how to calibrate with a minimum number of model points.

In contrast to these approaches, we formulate the projection of the object models to one or more images explicitly according to the pin-hole camera model in order to be able to estimate the pose parameters for all objects and for all cameras and the focal lengths of the cameras. Seemingly complicated projection functions with complex partial derivatives have the advantage that a minimum number of parameter values are to be estimated. E.g., three parameters determine a rotation rather than the nine entries of a rotation matrix. Furthermore, geometric information is explicitly captured without additional constraints. Our method holds for any model feature, not only for points, and for any number of known objects and images. Moreover, we present a new method for the projection of circles using projective invariants. A multivariate non-linear cost function measuring the deviation of projected model features from detected image features is minimised simultaneously for all detected image features in all images using the Levenberg-Marquardt method. Constraints such as planarity or location of features other than those encoded in the object models can be incorporated easily.

Experimental results from both simulated and real images are presented and show the robustness of our approach for nearly planar scenes with small objects.

2 Model-based 3-D Reconstruction and Camera Parameter Estimation

Model-based 3-D reconstruction is a quantitative method to estimate simultaneously the best viewpoint of all cameras and the object pose parameters by fitting the projection of a three-dimensional model to given two-dimensional features. The model-fitting is accomplished by minimising a cost function measuring all differences between projected model features and detected image features as a function of the objects' pose and the camera parameters. Common features in the scenes we are dealing with are points and circles. The projection of circles results in ellipses.

The projection of one object model depends on $7 \cdot n$ parameters where n is the number of used images. The seven parameters per view are the focal length, three rotational, and three translational parameters.

Projection of model points

The projection of a model point is the transformation of the point[1] $\boldsymbol{x}_o$ from model coordinates o to the camera coordinate system l and the subsequent projection onto the image plane b_l. This can be expressed in homogeneous coordinates[2] as

$$\tilde{\boldsymbol{x}}_{b_l} = \mathcal{P}^p_{b_l o}(\boldsymbol{x}_o) = \Phi^{-1}\left(\mathcal{T}_{b_l l} \cdot \mathcal{T}_{lo} \cdot \Phi(\boldsymbol{x}_o)\right) \tag{1}$$

$$= \Phi^{-1}\left(\left(\begin{pmatrix} \frac{b}{d_x} & 0 & C_x & 0 \\ 0 & \frac{b}{d_y} & C_y & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \cdot \right.\right.$$

$$\left.\left.\begin{pmatrix} \cos\theta\cos\phi & \cos\theta\sin\phi & -\sin\theta & t_x \\ \sin\psi\sin\theta\cos\phi - \cos\psi\sin\phi & \sin\psi\sin\theta\sin\phi + \cos\psi\cos\phi & \sin\psi\cos\theta & t_y \\ \cos\psi\sin\theta\cos\phi + \sin\psi\sin\phi & \cos\psi\sin\theta\sin\phi - \sin\psi\cos\phi & \cos\psi\cos\theta & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \cdot \Phi(\boldsymbol{x}_o)\right)\right).$$

Φ is a function for the transformation from affine to homogeneous coordinates. The projection of a model point in a second image plane b_r needs one additional transformation $\mathcal{T}_{rl}$ from the reference coordinate system which we place in the first camera coordinate system l to the second camera coordinate system r,

$$\tilde{\boldsymbol{x}}_{b_r} = \mathcal{P}^p_{b_r o}(\boldsymbol{x}_o) = \Phi^{-1}\left(\mathcal{T}_{b_r r} \cdot \mathcal{T}_{rl} \cdot \mathcal{T}_{lo} \cdot \Phi(\boldsymbol{x}_o)\right). \tag{2}$$

Projection of model circles

The perspective projection of circles which are planar figures can be understood as a collineation in the projective plane $\mathbb{IP}^2$. The cross ratio is invariant under every collineation (see [Semple & Kneebone 52]). Let $A, B, C,$ and D be four points in a projective plane, no three of them being collinear and a pencil of lines passing through these four points. P is the centre of the pencil. Its cross ratio k is

$$k = [PA, PB; PC, PD] = \frac{\overline{A'C'}}{\overline{A'D'}} \cdot \frac{\overline{B'D'}}{\overline{B'C'}}, \tag{3}$$

with A', B', C', D' being the intersections of this pencil with some line not passing through P.

The locus of the centres P of a pencil of lines passing through A, B, C, D having a given cross ratio is a conic through A, B, C, D (Theorem of Chasles, see [Semple & Kneebone 52, Mohr 93] and Fig. 1). Thus, a conic is uniquely defined by four points and a cross ratio. Its quadratic form is

$$ax^2 + 2bxy + cy^2 + 2dx + 2ey + f = 0 \tag{4}$$

and the coefficients are determined using

$$k L_{AB} L_{CD} + (1-k) L_{AC} L_{BD} = 0 \tag{5}$$

---

[1] Vectors are written in small bold type characters.

[2] Homogeneous transformations are denoted by $\mathcal{T}$ with subscripts indicating destination and source coordinate frame of the transformation.

with $L_{IJ} = (x_I - x)(y_I - y_J) - (y_I - y)(x_I - x_J);$ $\quad I, J \in \{A, B, C, D\}$ (ref. [Mohr 93]). To determine the cross ratio of a circle with radius r and the four points $(r, 0)$, $(0, r)$, $(-r, 0)$, and $(0, -r)$ we apply eq. (5) to the equation of a circle $-kr^2x^2 + (2kr^2 - 4r^2) xy - kr^2y^2 + kr^4 = 0$. This results in $2kr^2 - 4r^2 = 0 \Rightarrow k = 2$.
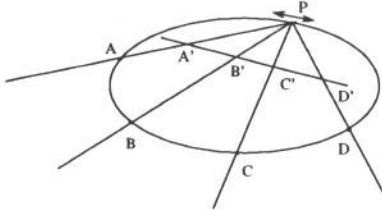

Thus, the quadratic form of a projected model circle is easily computed using four projected points on the circle and the corresponding cross ratio in eq. (5). This method holds with little extension for the projection of general ellipses, too.

Fig. 1: Cross ratio of a pencil of lines on a conic. The representation of an ellipse e as centre point m, radii l_1 and l_2, and orientation θ is more convenient as the quadratic form and enables the component-wise comparison with a detected image ellipse. This representation is obtained from the quadratic form with

$$m = \frac{1}{\delta} \left(\begin{vmatrix} b & c \\ d & e \end{vmatrix}, - \begin{vmatrix} a & b \\ d & e \end{vmatrix} \right)^T \text{ and } \theta = \frac{1}{2} \arctan \left(\frac{2b}{a-c} \right) \text{ with } \delta = \begin{vmatrix} a & b \\ b & c \end{vmatrix}. \quad (6)$$

Let λ_1 and λ_2 be the real solutions of the polynomial $\lambda^2 - (a+c) \cdot \lambda + \delta = 0$, then the radii are

$$l_1 = \sqrt{\left| \frac{\Delta}{\delta \cdot \lambda_1} \right|} \quad \text{and} \quad l_2 = \sqrt{\left| \frac{\Delta}{\delta \cdot \lambda_2} \right|} \quad \text{with} \quad \Delta = \begin{vmatrix} a & b & d \\ b & c & e \\ d & e & f \end{vmatrix}. \quad (7)$$

The projection of a model circle to the first and to the second image plane are denoted by

$$\tilde{x}_{b_l} = \mathcal{P}_{b_l o}^e(x_o) = \Gamma_b \left(\mathcal{T}_{b_l l} \cdot \mathcal{T}_{lo} \cdot \Gamma_c(x_o) \right), \quad (8)$$

$$\tilde{x}_{b_r} = \mathcal{P}_{b_r o}^e(x_o) = \Gamma_b \left(\mathcal{T}_{b_r r} \cdot \mathcal{T}_{rl} \cdot \mathcal{T}_{lo} \cdot \Gamma_c(x_o) \right). \quad (9)$$

An image ellipse $\tilde{x}_b$ is described by its centre point m, the radii l_1 and l_2 and its orientation θ. The function realising the transformation (eq. (6) and eq. (7)) of the projected model circle in homogeneous coordinates to the ellipse representation is Γ_b. A model circle x_o is characterised by its centre point, the radius and a normal vector in model coordinates o. The function Γ_c calculates the four points that are projected and their cross ratio in homogeneous coordinates. This formulation of the perspective projection of a model circle allows us to measure easily the deviation of projected and detected ellipses comparing five parameters.

Model-fitting

The pose of an object is well estimated from the image data if the value of the non-linear multivariate cost function

$$C(a) = \sum_{i=1}^{N} \sum_{j \in B} \left(x_{b_{j_i}} - \mathcal{P}_{b_j o}^i(a, x_{o_i}) \right)^T \cdot \Sigma^{-1} \cdot \left(x_{b_{j_i}} - \mathcal{P}_{b_j o}^i(a, x_{o_i}) \right). \quad (10)$$

is minimal. The cost function C measures the deviation of projected model features $\boldsymbol{x}_{o_i}$ – these can be points or circles – from the corresponding image features. The vector $\boldsymbol{a}$ contains all unknown parameters. B is the set of images of a scene. N is the number of corresponding model and image feature pairs. Depending on the feature, the vectors $\boldsymbol{x}_{b_{j_i}}$ and $\boldsymbol{x}_{o_i}$ contain different representations and the projection functions $\mathcal{P}^i_{b_{j}o}$ are the respective transformations. Σ is a covariance matrix which is used to model the admissible tolerance with respect to deviation from projected model to detected image features.

Camera Parameter Estimation

Classical camera calibration methods (e.g. [Tsai 85]) can not be performed on-line as they demand a special calibration pattern. Depth estimation is then a two-step process and it may lead to suboptimal solutions. We have explicitly modelled the camera parameters in our projection functions and thus they are estimated using the knowledge of the 3-D structure of the objects in the scene. We estimate the external camera parameters and the focal length. The results show that principal point and scale factors are stable enough for our off-the-shelf CCD cameras to assume fixed values. The influence of lens distortion to the results of our approach is quite small. We demonstrate this with simulated data in section 4. Nevertheless, it is possible to model the estimation of lens distortion in a manner similar to that of [Li 94].

[Tsai 85] shows that full camera calibration is possible with five coplanar reference points. A solution for calibration derived with four coplanar points is unique because four coplanar points determine a collineation in a plane and any further imaginary points in that plane as intersections of lines between lines through the four points can be derived. Six non coplanar points determine a unique solution as well (see [Yuan 89]).

Calibration is possible with one camera view. Taking a stereo image leads to much more robust results. Furthermore, the pose of a circle with known radius can not be computed uniquely from one view (see [Ma 93, Safaee-Rad et al. 92]). Taking at least two images for reconstruction, the pose of a circle in space is, if the focal lengths are known, uniquely defined up to the direction of its normal vector (ref. [Dhome et al. 90]). The sign of the normal can be determined due to the visibility of the projected ellipse. The focal lengths are determined with any non-circular object in the scene with at least four visible coplanar model points.

3 Minimisation

The main problem of non-linear parameter estimation is to find a method which guarantees convergence of the cost function (eq. 10) to a global minimum. The minimisation using the Levenberg-Marquardt method (see [Press et al. 88]), which is a combination of Newton's method and a gradient descent, converges to the nearest local minimum. The global minimum is found with good initial parameter values. However, we do not have initial parameter estimates. Thus, we divide the global model fitting problem into three stages to step-wise enhance and monitor the estimates of the parameters.

The Jacobian matrix $\left(\partial \mathcal{P}^i_{b_j o}(\boldsymbol{a}, \boldsymbol{x}_{o_i})/\partial \boldsymbol{a}\right)_{ij}$ is an essential prerequisite for the minimisation and it is obtained by applying the generalised chain rule for the concatenation of projection and feature representation extraction. The partial derivatives are computed using MAPLE[3].

Stage I: In the first stage, the poses of all objects are reconstructed individually, and separately for each camera view. As few parameters are to be estimated, the individual reconstructions are performed very quickly; however the minimisations have to be monitored in order not to let them converge to false local minima because of inappropriate initial values. The initial value for the focal length is chosen to be a commonly used length (e.g. 15mm). For the translation in z-direction we take a typical object distance (e.g. 2m). The initial x- and y-translation parameter values are calculated from the assumed focal length and z-translation, tracing the view ray through an image point and a model point. Rotation parameters can be set to any values.

During minimisation the focal length is monitored. If it leaves an admissible range (10-100mm in our case), the object is rotated by negating two rotational parameters and the minimisation is restarted with the other parameters reset to their original initial values. The cost function is also monitored during minimisation. If the process converges to a local minimum with inadmissibly high costs, the z-translation parameter is modified according to a predefined scheme. To improve the speed of convergence, it is useful to additionally adjust the x- and y-translation parameters which should be consistent with the current z-translation and focal length.

This monitored Levenberg-Marquardt iteration is stopped if either the change of the parameter estimates from one iteration step to the next is less than a given threshold, or if the model fitting does not succeed, i.e. if a maximum number of iterations is reached or if the same local minimum is found despite modified parameter values.
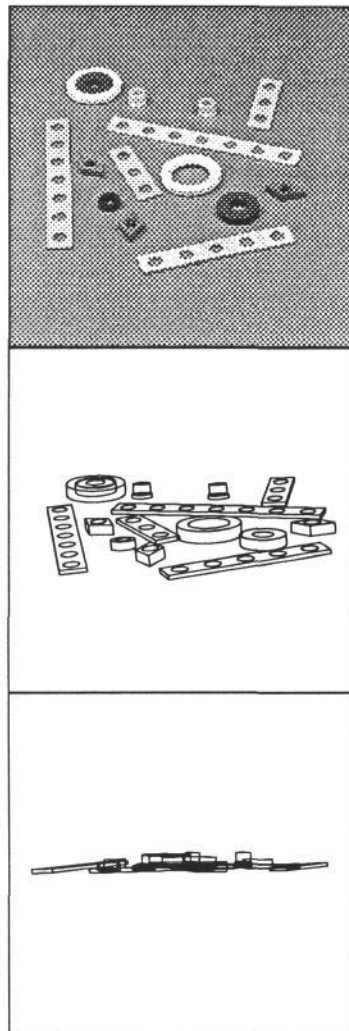

Fig. 2: a) A scene and the result of its 3-D reconstruction b) in a front and c) in a side view to show the accuracy of the reconstruction w.r.t. the planar surface of the table.

Stage II: Applying this method to each detected object, we obtain several estimates of the focal length for each camera and an estimate for the pose of each object relative to each camera. For a given camera the median of all estimates of

---

[3]MAPLE V Release 3 © by Waterloo Maple Software and the University of Waterloo

the focal length from stage I is fixed at this stage and it is used to reconstruct the pose of each object in the scene. So in this stage, better initial estimates for objects' poses are derived for each view of the scene.

Stage III: The median focal length and the resulting objects' poses of stage II are used as initial values for global model fitting. It is possible to estimate the relative pose between different cameras from the object correspondences. But, we found out that the minimisation process is not sensitive to the initial values of these parameters. Therefore, we take a rough estimation of commonly-used camera positions. Monitoring of this global minimisation is not necessary because of the good initial parameters values now available.

4 Experimental Results

Various experiments with the approach outlined in the preceding sections have been performed using real images as well as synthetic data with simulated noise. It is not possible for us to measure the exact distance between the cameras and the scene and the exact focal lengths. We get results for the focal lengths which are very similar to those of our implementation of the algorithm of [Tsai 85]. The results of camera calibration and 3-D reconstruction are evaluated comparing measured and reconstructed distances within the scene. Fig. 2 shows a scene and two views of its reconstruction from a stereo image. The side view indicates the accuracy of the reconstruction w.r.t. the planar surface of the table.

4.1 Accuracy of 3-D reconstruction

In order to measure accuracy in 3-D reconstruction we use a scene of objects with known relative poses shown in Fig. 3. The pose of an object relative to another is uniquely determined by the distances between two points of the first object and one point of the other object and the angle between two surface normals. Table 1 shows the accuracy of 3-D reconstruction comparing reconstructed and measured distances and angle differences between the top surface normals of all objects. The objects are taken from a children's toolkit and are imprecisely manufactured, therefore inaccuracies of measurements up to ±1mm can occur.

Fig. 3: A scene with objects with known relative poses.

The results reflect that the more features available for one object, the better the accuracy of the estimated pose. For the two holed bars (object 1 and 2 in Fig. 3) the four vertices and three or seven model circles are visible on the top surface. The estimated relative pose is very close to the measured one. The ring (object 4) is reconstructed using only the detected ellipse of the hole. The results show that for this object the largest errors occur.
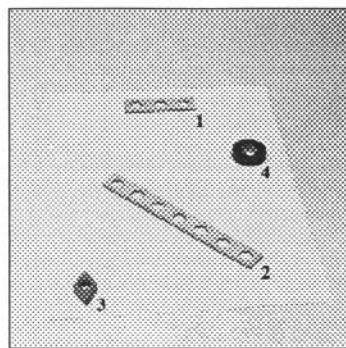
Object	$\Delta\varphi[°]$	d_0[mm]	Δd_0[mm]	$\frac{\Delta d_0}{d_0}$ [%]	d_1[mm]	Δd_1[mm]	$\frac{\Delta d_1}{d_1}$ [%]
1, 2	2.7	293.5	1.9	0.7	191.9	2.1	1.1
1, 3	3.6	351	4.8	1.3	325.9	5	1.5
1, 4	1.4	111.7	3.7	3.4	-	-	-
2, 1	2.7	293.5	1.9	0.7	313.9	2	0.6
2, 3	3.5	191.2	6.8	3.4	177.7	6.3	3.4
2, 4	2.4	198	2.6	1.3	-	-	-
3, 1	3.6	351	4.8	1.3	317.1	3.2	1
3, 2	3.5	191.2	6.8	3.4	171.8	1.8	1
3, 4	4.6	306.3	9.6	3	-	-	-
4, 1	1.4	111.7	3.7	3.4	175.7	2.2	1.3
4, 2	2.4	198	2.6	1.3	190.1	3.5	1.8
4, 3	4.6	306.3	9.6	3	282.2	9.6	3.3

Table 1: Accuracy in 3-D reconstruction of the scene shown in Fig. 3. The table shows the difference of the measured distance between the first point of the first object and one point of the second object Δd_0, the difference in distance between the second point of the first object and the point of the second object Δd_1, and $\Delta\varphi$ which measures the difference between the surface normals of the two objects.

4.2 Sensitivity to errors in image coordinates

Accuracy of feature detection and number of used features

The accuracy of 3-D reconstruction and camera calibration is mainly influenced by the accuracy of image feature detection and the the number of features used for reconstruction and calibration. This is shown using simulated data. Equally distributed noise in the range of ± 0.5 pixel is added to all projected model-features of two three-holed bars (ex. Fig. 2) independently.

The reconstructed distance between the two bars is used to measure accuracy. Fig. 4 shows the results of three different simulations with 1000 runs in each simulation. The errors in the reconstruction from the simulated data seem quite large. This is due to the small number of available features for calibration and reconstruction, compared to usual calibration patterns, and because the added noise is quite large.

Notice the different scaling of the histograms. Fig. 4a) shows the distribution of the reconstructed distances using one image and four points for reconstruction. The results improve if more features are used. Fig. 4b) shows the effect when four points and three circles per object are used for reconstruction and calibration in one image. The mean of the reconstructed distances is 172.5mm and the standard deviation is 16.6. A simulation taking four objects with four points per object in one image leads to similar results, $\mu = 173.1$ and $\sigma = 23.3$. This time a total of 16 features for four objects is used. This is worse than 14 features for two objects as in Fig. 4b), and this is reflected by the results. The best results are achieved using two views with the maximum number of features per object. The mean in Fig. 4c) is very close to the true distance, and the standard deviation is $\sigma = 3.5$.

Radial distortion

Another experiment with synthetic data shows the influence of radial distortion to calibration and reconstruction. Radial distortion resulting in a maximum displacement of $\Delta r_{max} = 4.5$ pixel at the corners of an image is added to a synthetic image. This is a commonly occurring distortion. Table 2 shows a maximal dif-

N° of views	N° of objects	N° of features	$\frac{\Delta d_0}{d_0}$ [%]
1	2	2×4	1.7
1	2	2×7	0.9
2	2	2×4	0.7
2	2	2×7	0.5

Table 2: Influence of radial distortion to the accuracy of reconstruction and calibration, $\Delta r_{max} = 4.5$ pixel.

ference between true and reconstructed object distance of 1.7%. The influence of radial distortion becomes smaller if more features and more views of a scene are used for reconstruction and calibration. [Weng *et al.* 92] report similar results.
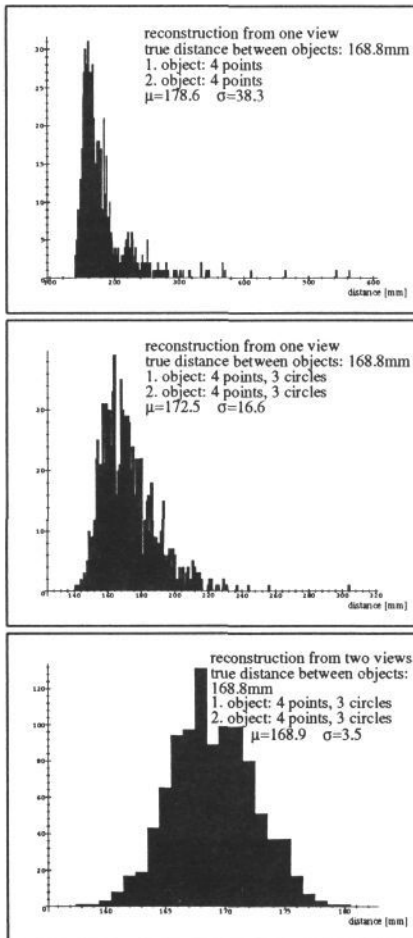
5 Conclusion

Fig. 4: Sensitivity to errors in image point coordinates using different numbers of model features.

A method for camera calibration and metric 3-D reconstruction from one or more uncalibrated images is presented. To this end, several objects of the scene in contrast to a special calibration pattern are used. These objects are modelled as sets of vertices, edges, and circles and the correspondence between model and image features is exploited as well as the correspondence of objects in different views of the scene. The projection of features is modelled explicitly capturing the geometric constraints given by a pin-hole camera model resulting in a minimum number of parameters to be estimated. For the perspective projection of model circles we derive a new formulation using projective invariants. This results in a simple method for the comparison of a projected model circle with a detected image ellipse on the basis of five parameters. To minimise a suitable cost function we apply the Levenberg-Marquardt method in a three-stage process, monitoring the iteration in order to stepwise gain good initial parameter estimates for subsequent minimisations. The accuracy of 3-D reconstruction is shown using real images. The relative error of distance and surface normal between objects in the scene is in the range of few percents. Using simulated data we show that the accuracy of the results is mainly influenced by the accuracy of image feature detection and the the number of features used, while radial distortion shows little impact on the accuracy.

Future work will concern two extensions of our approach. First, the method will be extended to re-calibrate the camera(s) using image sequences. This will allow us to iteratively enhance the estimates of the parameters and to support a vision system with active camera(s). Furthermore, other types of features, like edges, will be incorporated into model-fitting and minimisation in order to exploit more information of the objects in the scene.

References

[Boufama *et al.* 93] B. Boufama, R. Mohr, F. Veillon, Euclidean Constraints for Uncalibrated Reconstruction. *Proc. Int. Conf. on Computer Vision*, Berlin, Germany, 1993, pp. 466–470.

[Crowley *et al.* 93] J. Crowley, P. Bobet, C. Schmid, Auto-Calibration by Direct Observation of Objects. *Image and Vision Computing* **11**:2, 1993, pp. 67–81.

[Dhome *et al.* 90] M. Dhome, J. T. Lapreste, G. Rives, M. Richetin, Spatial Localization of Modelled Objects of Revolution in Monocular Perspective Vision. *Proc. First European Conference on Computer Vision*, 1990, pp. 475–485.

[Faugeras 92] O. Faugeras, What can be seen in three dimensions with an uncalibrated stereo rig ?. *Proc. Second European Conference on Computer Vision*, G. Sandini (Ed.), Lecture Notes in Computer Science 588, Springer-Verlag, Berlin etc., 1992, pp. 563–578.

[Goldberg 93] R. R. Goldberg, Pose determination of parameterized object models from a monocular image. *Image and Vision Computing* **11**:1, January/February 1993, pp. 49–62.

[Li 94] Mengxiang Li, Camera Calibration of a Head-Eye System for Active Vision. *Proc. 3^{rd} European Conference on Computer Vision* **I**, ECCV'94, Stockholm, Sweden, May 2-6, 1994, pp. 543–554.

[Lowe 91] D. G. Lowe, Fitting Parameterized Three-Dimensional Models to Images. *IEEE Trans. Pattern Analysis and Machine Intelligence* **PAMI-13**:5, 1991, pp. 441–450.

[Ma 93] Song De Ma, Conics-Based Stereo, Motion Estimation, and Pose Determination. *International Journal of Computer Vision* **10**:1, 1993, pp. 7–25.

[Mohr *et al.* 93] R. Mohr, F. Veillon, L. Quan, Relative 3D Reconstruction Using Multiple Uncalibrated Images. *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, New York, NY, 1993, pp. 543–548.

[Mohr 93] R. Mohr, Projective Geometry and Computer Vision. C. H. Chen, L. F. Pau, P. S. P. Wang (eds.): *Handbook of Pattern Recognition and Computer Vision*, Ch. 2.4, World Scientific Publishing Company, 1993, pp. 369–393.

[Press *et al.* 88] W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, *Numerical Recipes in C.* Cambridge University Press, 1988.

[Safaee-Rad *et al.* 92] R. Safaee-Rad, I. Tchoukanov, K. C. Smith, B. Benhabib, Constraints on quadratic-curved features under perspective projection. *Image and Vision Computing* **10**:8, October 1992, pp. 532–548.

[Semple & Kneebone 52] J. G. Semple, G. T. Kneebone, *Algebraic Projective Geometry.* Oxford Science Publication, 1952.

[Tsai 85] R. Tsai, A Versatile Camera Calibration Technique for High Accuracy 3D Machine Vision Metrology using Off-the-Shelf TV Cameras and Lenses, Research Report, 1985.

[Weng *et al.* 92] Juyang Weng, Paul Cohen, Marc Herniou, Camera Calibration with Distortion - Models and Accuracy Evaluation. *IEEE Trans. Pattern Analysis and Machine Intelligence* **PAMI-14**:10, 1992, pp. 965–980.

[Yuan 89] J. Yuan, A General Photogrammetric Method for Determining Object Position and Orientation. *IEEE Trans. Pattern Analysis and Machine Intelligence* **PAMI-5**:2, 1989, pp. 129–142.