Katholieke
Universiteit
Leuven

Departement Elektrotechniek
Afdeling ESAT/MI2
Kardinaal Mercierlaan 94
B-3001 Heverlee - Belgium

TECHNISCH RAPPORT - TECHNICAL REPORT

A stratified approach to metric self-calibration

Marc Pollefeys and Luc Van Gool

January 1997

Nr. KUL/ESAT/MI2/9702

A stratified approach to metric self-calibration

M. Pollefeys and L. Van Gool
K.U.Leuven-MI2
Belgium

Abstract

Camera calibration is essential to many computer vision applications. In practice this often requires cumbersome calibration procedures to be carried out regularly. In the last few years a lot of work has been done on self-calibration of cameras, ranging from weak calibration to metric calibration. It has been shown that a metric calibration of the camera setup (up to scale) was possible based on the rigidity of the scene only. In this report a stratified approach is proposed which gradually retrieves the metric calibration of the camera setup. Starting from an uncalibrated image sequence the projective calibration is retrieved first. In projective space the plane at infinity is then identified yielding the affine calibration. This is achieved using a constraint which can be formulated between any two arbitrary images of the sequence. Once the affine calibration is known the upgrade to metric is easily obtained through linear equations.

1 Introduction

In recent years several methods were proposed to obtain the calibration of a camera from correspondences between several views of the same scene. These methods are based on the rigidity of the scene and on the constancy of the internal camera parameters. Most existing methods start from the projective calibration and then immediately try to solve for the intrinsic parameters. However, they all have to cope with the affine parameters (i.e. the position of the plane at infinity).

Faugeras *et al* [4] eliminated these affine parameters yielding two Kruppa equations for each pair of views. A more robust approach was proposed by Zeller *et al* [16]. Heyden [7] proposed a variant of this method. He does not eliminate the affine parameters and has to solve for an additional scale factor for each pair of views he takes into account (which means at least ten unknowns). Hartley [6] does a minimization on all eight parameters to obtain metric camera projection matrices. Most of these methods encounter problems having to solve for many parameters at once from nonlinear equations.

This problem prompted a stratified approach, where starting from a projective reconstruction an affine reconstruction is obtained first and used as the initialization towards metric reconstruction. A similar method has been proposed by Armstrong *et al* [1] based on the work of Moons *et al* [9]. But this method needs a pure translation which is for example not easy to ensure with a hand-held camera. A method which avoids this restriction was proposed by Pollefeys *et al* [11]. This method has the disadvantage of needing at least 4 views to obtain the self-calibration.

In this report a new method is derived based on the generalization of the method proposed in [11]. This method is more robust and can obtain the metric calibration of a camera setup from only 3 images. This report also investigates the possibility of using only 2 views of the scene. It will be shown that combining geometric constraints with internal or external constraints can solve the calibration where none of them separately could.

2 Camera models and geometry

In this section the basic principles and notations used in the rest of the report are introduced. Projective geometry is used throughout the report to describe the perspective projection of the scene onto the images. This projection is described as

$$m \simeq \mathbf{P}M \tag{1}$$

where $\mathbf{P}$ is a 3×4 projection matrix describing the perspective projection process, $M = [X Y Z\, 1]^\top$ and $m = [x\, y\, 1]$ are vectors containing the homogeneous coordinates of the world points respectively image points. Notice that $\simeq$ will be used throughout this report to indicate equality up to a non-zero scale factor.

If $C = [t^\top 1]^\top$ is the optical center of the camera, the projection matrices can in general be written as follow:

$$\mathbf{P} = [\mathbf{H}| \text{-}\mathbf{H}t] \tag{2}$$

with $\mathbf{H}$ the homography describing the projection of points from the reference plane (the plane at infinity in the affine case) to the image plane. When the geometry is only determined up to a projective or affine transformation, the first projection matrix can be chosen as follows $\mathbf{P}_1 = [\mathbf{I}|\mathbf{0}]$. This implies that the camera and scene coordinate frames are aligned. In that case the homographies of equation 2 also describe the transfer from points lying in the reference plane from the first image to the image under consideration.

When the ambiguity on the geometry is metric, the camera projection matrices have the following form:

$$\mathbf{P}_{E1} = \mathbf{K}[\,\mathbf{I}\,|0] \text{ and } \mathbf{P}_{Ek} = \mathbf{K}[\mathbf{R}_k \,|- \mathbf{R}_k t_k] \tag{3}$$

with t_k and $\mathbf{R}_k$ indicating the position and orientation of the camera for view k and $\mathbf{K}$ an upper diagonal 3×3 matrix containing the internal camera parameters:

$$\mathbf{K} = \begin{bmatrix} f_x & s & u_x \\ & f_y & u_y \\ & & 1 \end{bmatrix} \tag{4}$$

where f_x and f_y represent the focal length divided by the pixel dimensions, s is a measure of the skew and (u_x, u_y) is the principal point.

3 Self-calibration from an uncalibrated image sequence

Starting from an uncalibrated image sequence it is possible to obtain a reconstruction up to a projective transformation, based on the assumption that the pinhole camera model is valid[3, 5, 13]. To obtain an affine or metric reconstruction additional knowledge has to be brought in. Geometric as well as internal or external constraints can be used to identify the plane at infinity or to find the camera parameters. Often the assumption is made that the intrinsic camera parameters are constant throughout the sequence (5 constraints per additional view). Restricting the projective ambiguity (15 parameters) to metric (Euclidean+scale=7 parameters) requires 8 or more constraints. These are needed to solve for the plane at infinity and the intrinsic parameters. A quick calculation shows that this must in general be possible for 3 or more images.

3.1 Projective calibration

In a first stage the projective calibration of the whole sequence is calculated together with a reconstruction of the scene. This reconstruction is initialized from the two first views. The fundamental matrix $\mathbf{F}$ can be calculated in a robust way from point matches [15]. The epipoles correspond to the left and right nullspaces of $\mathbf{F}$. Then the homography for some arbitrary plane $P = [a1]^\top$ is obtained as follow: $\mathbf{H} = [e]_\times \mathbf{F} + ea$ (with e the epipole in the left image and $[e]_\times$ the antisymmetric matrix representing the vector product with this epipole e). Hence the following camera projection matrices are obtained for the first two views (transforming the plane P to $[0\,0\,0\,1]^\top$):

$$\mathbf{P}_1 = [\mathbf{I}|\mathbf{0}] \text{ and } \mathbf{P}_2 = [\mathbf{H}|e] \tag{5}$$

The reconstruction of the points can be obtained by minimizing the reprojection error in the images. The camera projection matrices for the following views are then obtained by imposing that some previously reconstructed points which are still visible in the new view are reprojected as closely as possible to the matches:

$$\min_{\mathbf{P}_k} \sum \left[(x_k - \frac{\mathbf{p}_{k1} M}{\mathbf{p}_{k3} M})^2 + (y_k - \frac{\mathbf{p}_{k2} M}{\mathbf{p}_{k3} M})^2 \right] \tag{6}$$

where $\mathbf{p}_{k1}, \mathbf{p}_{k2}$ and $\mathbf{p}_{k3}$ are the rows of $\mathbf{P}_k$. Once $\mathbf{P}_k$ is known, points which are visible in two or more views can be added to the reconstruction or updated. A more complete description of this approach is given by Beardsley *et al* [2].

3.2 Affine calibration

The concern in this step of the calibration is to identify the plane at infinity and to use it as the reference plane of the previous paragraph. This is of course easy when 3 or more points of this plane are known. For the moment, it is assumed that nothing is known about the scene. Therefore vanishing points or parallelism can not be used to identify the plane at infinity.

In this section it will be explained how a property of the infinity homography yields one constraint on the plane at infinity for *any* pair of views. This is an important generalization of the constraint introduced by Pollefeys *et al* [11] where the first image of the pair always had to be a fixed reference image. This generalization not only reduces the minimum number of images to 3 in stead of 4, but –more important– also increases the robustness and accuracy of the method.

What is needed at this point is something which makes the plane at infinity different from all the others so that it can be identified. The infinity homographies can be written as a function of the Euclidean entities of equation 3 or explicitly as functions of projective entities –known at this point– and the parameters of the plane at infinity. Both representations are given in the following equation:

$$\mathbf{H}_{\infty kl} \simeq \underbrace{\mathbf{K}\mathbf{R}_{kl}\mathbf{K}^{-1}}_{\sim Euclidean} \simeq \underbrace{(\mathbf{H}_{1l} + e_{1l}a)(\mathbf{H}_{1k} + e_{1k}a)^{-1}}_{\sim projective} \quad . \tag{7}$$

From the Euclidean representation it follows that $\mathbf{H}_{\infty kl}$ is conjugated with a rotation matrix (up to a scale factor) which implies that the 3 eigenvalues of $\mathbf{H}_{\infty 1k}$ must have the same moduli, hence the modulus constraint [10, 11]. Note from equation 7 that this property requires the intrinsic camera parameters to be constant.

This can be made explicit by writing down the characteristic equation for the infinity homography:

$$\det(\mathbf{H}_{\infty kl} - \lambda \mathbf{I}) = l_3\lambda^3 + l_2\lambda^2 + l_1\lambda + l_0 = 0 \tag{8}$$

It can be shown that the following condition is a necessary condition for the roots of equation 8 to have equal moduli:

$$l_3 l_1^3 = l_2^3 l_0 \tag{9}$$

This yields a constraint on the 3 affine parameters a_1, a_2, a_3 by expressing l_3, l_2, l_1, l_0 as a function of these parameters (using the projective representation in equation 7). The inverse in the projective representation of $\mathbf{H}_{\infty kl}$ can be avoided by using the constraint $\det(\mathbf{H}_{\infty 1l} - \lambda \mathbf{H}_{\infty 1k}) = 0$ which is equivalent to equation 8. Factorizing this expression using the multi-linearity of determinants, l_3, l_2, l_1, l_0 turn out to be *linear* in a_1, a_2, a_3. The explicit expressions can be found in equation 10. The following notations are used: $\mathbf{H}_{1l} = [\mathbf{h}_1\mathbf{h}_2\mathbf{h}_3]$, $\mathbf{H}_{1k} = [\mathbf{h}'_1\mathbf{h}'_2\mathbf{h}'_3]$, $e_{1l} = \mathbf{e}$ and $e_{1k} = \mathbf{e}'$, $|\mathbf{H}|$ means the determinant of $\mathbf{H}$.

$$\begin{aligned}
l_3 =& - |\mathbf{h}'_1\, \mathbf{h}'_2\, \mathbf{h}'_3| - a_1|\mathbf{e}'\, \mathbf{h}'_2\, \mathbf{h}'_3| - a_2|\mathbf{h}'_1\, \mathbf{e}'\, \mathbf{h}'_3| - a_3|\mathbf{h}'_1\, \mathbf{h}'_2\, \mathbf{e}'| \\
l_2 =& + (|\mathbf{h}_1\mathbf{h}'_2\mathbf{h}'_3| + |\mathbf{h}'_1\mathbf{h}_2\mathbf{h}'_3| + |\mathbf{h}'_1\mathbf{h}'_2\mathbf{h}_3|) \\
& + a_1 (|\mathbf{e}\,\mathbf{h}'_2\mathbf{h}'_3| + |\mathbf{e}'\mathbf{h}_2\mathbf{h}'_3| + |\mathbf{e}'\mathbf{h}'_2\mathbf{h}_3|) \\
& + a_2 (|\mathbf{h}_1\mathbf{e}'\mathbf{h}'_3| + |\mathbf{h}'_1\mathbf{e}\,\mathbf{h}'_3| + |\mathbf{h}'_1\mathbf{e}'\mathbf{h}_3|) \\
& + a_3 (|\mathbf{h}_1\mathbf{h}'_2\mathbf{e}'| + |\mathbf{h}'_1\mathbf{h}_2\mathbf{e}'| + |\mathbf{h}'_1\mathbf{h}'_2\mathbf{e}|) \\
l_1 =& - (|\mathbf{h}'_1\mathbf{h}_2\mathbf{h}_3| + |\mathbf{h}_1\mathbf{h}'_2\mathbf{h}_3| + |\mathbf{h}_1\mathbf{h}_2\mathbf{h}'_3|) \\
& - a_1 (|\mathbf{e}'\mathbf{h}_2\mathbf{h}_3| + |\mathbf{e}\,\mathbf{h}'_2\mathbf{h}_3| + |\mathbf{e}\,\mathbf{h}_2\mathbf{h}'_3|) \\
& - a_2 (|\mathbf{h}'_1\mathbf{e}\,\mathbf{h}_3| + |\mathbf{h}_1\mathbf{e}'\mathbf{h}_3| + |\mathbf{h}_1\mathbf{e}\,\mathbf{h}'_3|) \\
& - a_3 (|\mathbf{h}'_1\mathbf{h}_2\mathbf{e}| + |\mathbf{h}_1\mathbf{h}'_2\mathbf{e}| + |\mathbf{h}_1\mathbf{h}_2\mathbf{e}'|) \\
l_0 =& + |\mathbf{h}_1\, \mathbf{h}_2\, \mathbf{h}_3| + a_1\, |\mathbf{e}\, \mathbf{h}_2\, \mathbf{h}_3| + a_2\, |\mathbf{h}_1\, \mathbf{e}\, \mathbf{h}_3| + a_3\, |\mathbf{h}_1\, \mathbf{h}_2\, \mathbf{e}|
\end{aligned} \tag{10}$$

Equation 9 thus yields an order 4 polynomial constraint for each camera pair. This means that for 3 images only a finite number of solutions exist. Checking the feasibility of the solutions is often enough to yield only one solution. In general with more images only one solution exists. In practice these constraints can be solved by general minimization algorithms. Because only 3 parameters have to be estimated convergence towards a global minimum is relatively easily obtained in general.

3.3 Metric calibration

Once the affine calibration is known it is easy to upgrade it to metric [6]. The conic $\mathbf{KK}^\top$ is invariant when multiplied to the left by $\mathbf{H}_\infty$ and to the right by $\mathbf{H}_\infty^\top$ as is seen when using $\mathbf{H}_\infty = \mathbf{KRK}^{-1}$. This property is related to the fact that the absolute conic is invariant under the group of Euclidean transformations. This implies that, when the camera parameters are fixed, the image of the absolute conic and thus also its dual (which is exactly $\mathbf{KK}^\top$) will stay fixed under the corresponding transformations. The following equation results in linear equations in the coefficients of $\mathbf{KK}^\top$ by choosing $\det \mathbf{H}_\infty = 1$:

$$\mathbf{KK}^\top = \mathbf{H}_\infty \mathbf{KK}^\top \mathbf{H}_\infty^\top \tag{11}$$

For a pair of images this set of equations only determines $\mathbf{KK}^\top$ up to a one parameter family of solutions. From 3 images or more $\mathbf{KK}^\top$ is fully determined except when the rotation axes are parallel. Once $\mathbf{KK}^\top$ is known, $\mathbf{K}$ can be extracted from it by Cholesky factorization.

4 Self-calibration from only 2 views?

Is it still possible to use self-calibration in the case where only two views are present? It is clear that in the strict sense self-calibration will not yield a unique solution. Imposing equality of the 5 internal camera parameters for both views doesn't give enough constraints to solve the self-calibration problem (3 affine and 5 intrinsic parameters).

On the other hand several constraints are still available and these can help solving the calibration. In this paragraph it will be shown that with a little extra information it is possible to retrieve the full calibration of the camera from only two views. This extra information can consist of constraints on the internal camera parameters or from external constraints like parallelism or right angles in the scene. Often some of these constraints can be obtained without user interaction.

4.1 Projective calibration

The projective part causes no problem. The same approach can be followed as in paragraph 3.1 yielding two camera projection matrices determined up to a projective transformation. The affine and the metric part are more involved and will be discussed in the following paragraphs.

4.2 Affine calibration

For a pair of images only one modulus constraint exists (see paragraph 3.2) which is not enough to locate uniquely the plane at infinity.

Some scenes contain parallel lines. These result in vanishing points in the images. Techniques have been proposed to automatically detect such points [8]. To identify the plane at infinity 3 of these are necessary. By using the modulus constraint this number can be reduced by one. This reduction can be crucial in practice. For example in the castle sequence (see figure 3) two vanishing points could be extracted automatically in all frames, not three. This is typical for a lot of scenes where one vanishing point is extracted for horizontal lines and one for vertical lines. Even when 3 candidate vanishing points are identified, the modulus constraint can still be very useful by providing a mean to check the hypothesis.

When a vanishing point is identified in the two images it can be used as follows to constraint the homography of the plane at infinity:

$$m_2 \simeq [\mathbf{H} + \mathbf{ea}]\, m_1 \tag{12}$$

This result in one linear equation for the coefficients of $\mathbf{a}$ (from the 3 equations only 2 are independent due to the epipolar correspondence of m_1 and m_2 and one is needed to eliminate the unknown scale factor).

With two known vanishing points we are thus left with a one parameter family of solutions for $\mathbf{H}_\infty$:

$$\mathbf{H}_\infty = \mathbf{H} + \mathbf{e}(\lambda \mathbf{a}_1 + \mathbf{a}_0) \ . \tag{13}$$

Applying the modulus constraint is much easier than in the general case. The coefficients l_3, l_2, l_1, l_0 (see equation 10) can be evaluated for both $\mathbf{a}_1$ and $\mathbf{a}_0$. The modulus constraint in the two view case then takes on the following form:

$$\begin{aligned}&(\lambda l_3(\mathbf{a}_1) + l_3(\mathbf{a}_0))(\lambda l_1(\mathbf{a}_1) + l_1(\mathbf{a}_0))^3 \\ &= (\lambda l_2(\mathbf{a}_1) + l_2(\mathbf{a}_0))^3(\lambda l_0(\mathbf{a}_1) + l_0(\mathbf{a}_0)) \ .\end{aligned} \tag{14}$$

This results in a polynomial of degree 4 in only one variable λ (not degree 6 as Sturm anticipated [14]). Therefore at most 4 solutions are possible. Because equation 14 is only a necessary condition for $\mathbf{H}_\infty$ to be conjugated with a scaled rotation matrix, this property should be checked out. This can eliminate several solutions. If different solutions subsist at this stage some can still be eliminated in the metric calibration part.

4.3 Metric calibration

Once the affine calibration is known equation 11 can be used. This results in a one parameter family of solutions for $\mathbf{KK}^\top$. Additional constraints like some known aspect ratio, perpendicularity of the image axes or scene orientations can be used to restrict $\mathbf{KK}^\top$ to one unique solution. If more than one affine calibration was still under consideration, these constraints can also help out. Also the fact that $\mathbf{KK}^\top$ should be positive definite and that the principal point should be more or less in the center of the image can be used to find the true affine, and thus also metric, calibration.

The application of these constraints is not so hard. Here the case of the orthogonal orientations in the scene will be discussed. This can for example be applied when it is assumed that the extracted vanishing points correspond to such orientations.

The points v and v' are two vanishing points in the first image. The corresponding scene points can be obtained from the following equation:

$$v \simeq \mathbf{K}[\mathbf{I}|\mathbf{0}] \begin{bmatrix} \tilde{v} \\ 0 \end{bmatrix} \tag{15}$$

Thus $\tilde{v} = \mathbf{K}^{-1}v$ represents the associated direction. Orthogonality now means that $\tilde{v}\tilde{v}' = 0$ or

$$v\mathbf{K}^{-\top}\mathbf{K}^{-1}v = 0 \ . \tag{16}$$

Therefore it is more appropriate to use the dual equation of equation 11:

$$\mathbf{K}^{-\top}\mathbf{K}^{-1} \simeq \mathbf{H}_{\infty}^{\top}\mathbf{K}^{-\top}\mathbf{K}^{-1}\mathbf{H}_{\infty} \tag{17}$$

which of course also yields a one parameter family of solutions for $\mathbf{K}^{-\top}\mathbf{K}^{-1}$. Adding equation 16 resolves this ambiguity.

From $\mathbf{K}^{-\top}\mathbf{K}^{-1}$ first $\mathbf{K}^{-1}$ is extracted by Cholesky factorization and subsequently inverted to obtain $\mathbf{K}$ which contains the camera internal parameters.

5 Experiments

Experiments have been done on both real and synthetic data. First the synthetic data give some insights in the behavior of the method depending on the number of views and the presence of noise. Then the feasibility of the method will be illustrated with some calibration/reconstruction work done on a real video sequence,

5.1 Simulations

The simulations were carried out on sequences of 3, 6 and 10 views. The scene consisted of 50 points uniformly distributed in a unit sphere with its center at the origin. For the calibration matrix the canonical form $\mathbf{K} = \mathbf{I}$ was chosen. The views were taken from all around the sphere and were all more or less pointing towards the origin. An example of such a sequence can be seen in figure 1. The scene points were projected into the images. Gaussian noise with an equivalent standard deviations of 0, 0.1, 0.2, 0.5, 1 and 2 pixels for 500×500 images was added to these projections. For every sequence length and noise level ten sequences were generated. The self-calibration method proposed in this report was carried out on all these sequences. The results for the camera intrinsic parameters were compared with the real values and the RMS error is shown in figure 2.

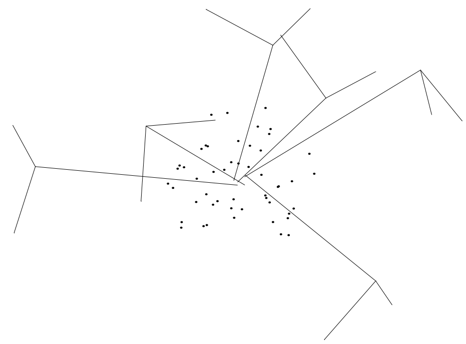

Figure 1: *Example of sequence used for simulations (the views are represented by the image axis and optical axis of the camera in the different positions.)*

When 6 or 10 views were used the accuracy was very good, even for high amounts of noise (around 1% error for 2 pixels noise). For only 3 views the method gives good results for small amounts of noise, but the error grows when more noise is added. This is due to the fact that in the 3 view case no redundancy is present in the equations.

The method almost always converges without problems. Only in the 3 view case the method regularly ends up in an erroneous solution (3 polynomials of degree 4 can have up to 64 possible solutions). By checking the fact that the modulus constraint is indeed satisfied and that the solution for $\mathbf{K}\mathbf{K}^\top$ is positive definite these problems could be reduced. But still around 40% of the 3 view sequences ended up in a wrong minimum. For the 3 view case only the experiments reaching the correct solution were taken into account to calculate the RMS error.

5.2 A real video sequence

In this paragraph results obtained from a real sequence are presented. The quality of the calibration can be evaluated by looking at the reconstruction. The sequence consists of a university building. These were recorded with a video camera. The images used for self-calibration are shown in figure 3. The approach for the calibration and reconstruction was the following:

- First a projective calibration of the sequence was obtained together with a reconstruction of the corners matched throughout (a part of) the sequence. A method similar to that presented by Beardsley [2] was used.

- Next, the projective ambiguity on the geometry was reduced to affine. This was achieved by identifying the plane at infinity with the generalized modulus constraint.

- Using equation 11 the camera intrinsic parameters were obtained and the reconstruction was upgraded to metric.

- The method of Proesmans *et al* [12] was used to obtain a dense correspondence map in spite of discontinuities and illumination changes.

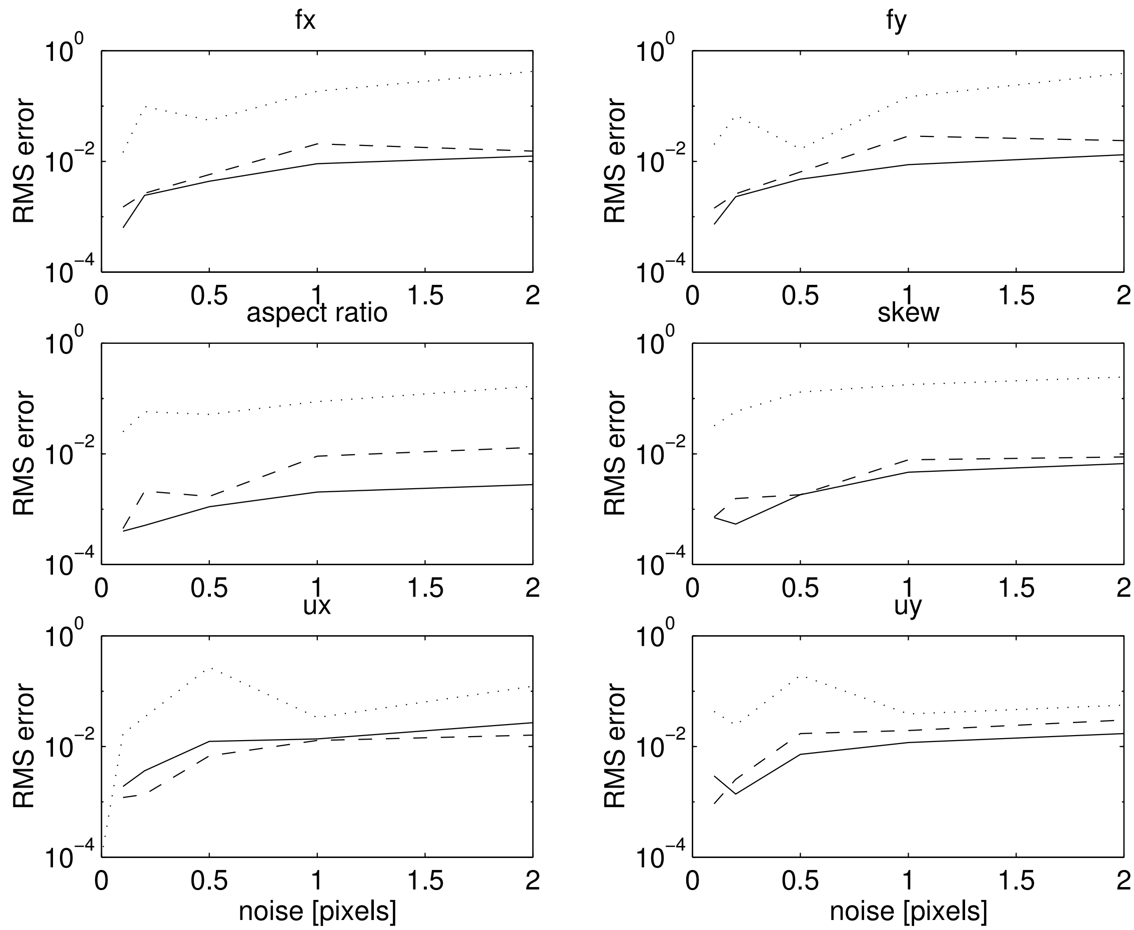

Figure 2: *RMS error on the camera intrinsic parameters for synthetic image sequences (3 views:· · ·, 6 views:- -, 10 views:—)*

- This map was used together with the metric camera projection matrices to generate the final reconstruction. One of the images is used as texture map.

In figure 4 one can see 3 orthographic views of the scene. Parallelism and orthogonality relation clearly have been retrieved. Look for example at the right angles in the top view or at the rectangular windows. Figure 5 and figure 6 contain some perspective views of the reconstruction. Because the dense correspondence map was only obtained from two images there are some inaccuracies left in the reconstruction. This however has nothing to do with the accuracy of the calibration.

6 conclusion

In this report a stratified approach to self-calibration has been proposed. Starting from point correspondences first the projective calibration is obtained, then affine and finally

Figure 3: *Images of the Arenberg castle which were used for self-calibration*

metric. The metric camera projection matrices can then be used with a correspondence map to obtain a dense reconstruction. This was illustrated with a real video sequence.

The results are quite satisfactory but can still be improved. A possible way of doing this is by appending a non-linear refinement method at the end of the self-calibration method. Also imposing some constraints like the absence of skew or a known aspect ratio can be expected to improve the accuracy.

Some ideas were proposed to combine geometric constraints with internal and external constraints. These were developed for the 2 view case where not enough geometric constraints are available for pure self-calibration. A lot of work remains to be done in this area.

Acknowledgments

Marc Pollefeys acknowledges a specialization grant from the Flemish Institute for Scientific Research in Industry (IWT). Financial support from the EU ACTS project AC074 'VANGUARD' and from IUAP-50 project of the Belgian OSTC is also gratefully acknowledged.

Figure 4: *Orthographic views of the reconstruction (notice parallelism and orthogonality)*

References

[1] M. Armstrong, A. Zisserman and P. Beardsley, Euclidean structure from uncalibrated images, *Proc. BMVC'94*.

[2] P. Beardsley, P. Torr and A. Zisserman 3D Model Acquisition from Extended Image Sequences *Proc. ECCV'96*, vol.2, pp.683-695

[3] O. Faugeras, What can be seen in three dimensions with an uncalibrated stereo rig, *Proc. ECCV'92*, pp.563-578.

[4] O. Faugeras, Q.-T. Luong and S. Maybank. Camera self-calibration: Theory and experiments, *Proc. ECCV'92*, pp.321-334.

Figure 5: *Perspective views of the reconstruction*

[5] R. Hartley, Estimation of relative camera positions for uncalibrated cameras, *Proc. ECCV'92*, pp.579-587.

[6] R. Hartley, Euclidean reconstruction from uncalibrated views, Applications of invariance in Computer Vision, LNCS 825, Springer-Verlag, 1994.

[7] A. Heyden, K. Åström, Euclidean Reconstruction from Constant Intrinsic Parameters *Proc. ICPR'96.*

[8] McLean, G.F. and Kotturi, D. Vanishing point Detection by Line Clustering, *PAMI*, pp.1090-1095, november 1995.

Figure 6: *More perspective views of the reconstruction*

[9] T. Moons, L. Van Gool, M. Van Diest and E. Pauwels, Affine reconstruction from perspective image pairs, Applications of Invariance in Computer Vision, LNCS 825, Springer-Verlag, 1994.

[10] M. Pollefeys, L. Van Gool and M. Proesmans, Euclidean 3D Reconstruction from Image Sequences with Variable Focal Lengths, *Proc. ECCV'96*, vol.1, pp. 31-42.

[11] M. Pollefeys, L. Van Gool and A. Oosterlinck, The Modulus Constraint: A New Constraint for Self-Calibration, *Proc. ICPR'96*, pp. 349-353.

[12] M. Proesmans, L. Van Gool and A. Oosterlinck, Determination of optical flow and its discontinuities using non-linear diffusion, *Proc. ECCV'94*, pp. 295-304.

[13] M. Spetsakis and Y.Aloimonos, A Multi-frame Approach to Visual Motion Perception *International Journal of Computer Vision*, 6:3, 245-255, 1991.

[14] P. Sturm and L. Quang, Affine stereo calibration, *Proc. CAIP'95*, pp. 838-843, Prague 1995.

[15] P. Torr, Motion Segmentation and Outlier Detection, *Ph.D.Thesis*, Oxford 1995.

[16] C. Zeller and O. Faugeras, Camera self-calibration from video sequences: the Kruppa equations revisited. Research Report 2793, INRIA, 1996.