

On Plane-Based Camera Calibration: A General Algorithm, Singularities, Applications

Peter F. Sturm and Stephen J. Maybank

Computational Vision Group, Department of Computer Science, The University of Reading
Whiteknights, PO Box 225, Reading, RG6 6AY, United Kingdom
{P.F.Sturm, S.J.Maybank}@reading.ac.uk

Abstract

We present a general algorithm for plane-based calibration that can deal with arbitrary numbers of views and calibration planes. The algorithm can simultaneously calibrate different views from a camera with variable intrinsic parameters and it is easy to incorporate known values of intrinsic parameters. For some minimal cases, we describe all singularities, naming the parameters that can not be estimated. Experimental results of our method are shown that exhibit the singularities while revealing good performance in non-singular conditions. Several applications of plane-based 3D geometry inference are discussed as well.

1 Introduction

The motivations for considering planes for calibrating cameras are mainly twofold. First, concerning calibration in its own right, planar calibration patterns are cheap and easy to produce, a laser printer output for example is absolutely sufficient for applications where highest accuracy is not demanded. Second, planar surface patches are probably the most important twodimensional “features”: they abound, at least in man-made environments, and if their metric structure is known, they carry already enough information to determine a camera’s pose up to only two solutions in general [4]. Planes are increasingly used for interactive modeling or measuring purposes [1, 10, 11].

The possibility of calibrating cameras from views of planar objects is well known [7, 12, 14]. Existing work however, restricts in most cases to the consideration of a single or only two planes (an exception is [8], but no details on the algorithm are provided) and cameras with constant calibration. In addition, the study of singular cases is usually neglected (besides in [12] for the simplest case, calibration of the aspect ratio from one view of a plane), despite their presence in common configurations.

It is even possible for cameras to self-calibrate from views of planar scenes with unknown metric structure [13], however several views are needed (Triggs recommends up to 9 or 10 views of the same plane for reliable results) and the “risk” of singularities should be greater compared to calibration from planes with known metric structure.

In this paper, we propose a general algorithm for calibrating a camera with possibly variable intrinsic parameters and position, that copes well with an arbitrary number of calibration planes and camera views. Calibration is essentially done in two steps. First, the 2D-to-2D projections of planar calibration objects onto the image plane(s) are computed. Each of these projections contributes to a system of homogeneous linear equations in the intrinsic parameters, which are hence easily determined. Calibration can thus be achieved by solving linear equations, but can of course be enhanced by subsequent non linear optimization.

In §2, we describe our camera model and projections of planar objects. In §3, we introduce the principle of plane-based calibration. A general algorithm is proposed in §4. Singularities are revealed in §5. Experimental results are presented in §6, and some applications described in §7.

2 Background

Camera Model. We use perspective projection to model cameras. A projection may be represented by a 3×4 projection matrix P that incorporates the so-called extrinsic and intrinsic camera parameters:

$$P \sim KR(I_3 \mid -t) . \quad (1)$$

Here, \sim means equality up to a non zero scale factor, I_3 is the 3×3 identity matrix, R a 3×3 orthogonal matrix representing the camera’s orientation, t a 3-vector representing its position, and K the 3×3 calibration matrix:

$$K = \begin{pmatrix} \tau f & s & u_0 \\ 0 & f & v_0 \\ 0 & 0 & 1 \end{pmatrix} .$$

In general, we distinguish 5 intrinsic parameters for perspective projection: the (effective) focal length f , the aspect ratio τ , the principal point (u_0, v_0) and the skew factor s accounting for non rectangular pixels. The skew factor is usually very close to 0 and we ignore it in the following.

Calibration and Absolute Conic. Our aim is to calibrate a camera, i.e. to determine its intrinsic parameters or its calibration matrix K (subsequent pose estimation is relatively straightforward). Instead of directly determining K ,

we will try to compute the symmetric matrix $\mathbf{K}\mathbf{K}^\top$ or its inverse, from which the calibration matrix can be computed uniquely using Cholesky decomposition [5]. This leads to simple and, in particular, linear calibration equations. Furthermore, the analysis of singularities of the calibration problem is greatly simplified: the matrix $\omega \sim (\mathbf{K}\mathbf{K}^\top)^{-1}$ represents the image of the Absolute Conic whose link to calibration and metric scene reconstruction is exposed for example in [2]. This geometrical view helps us with the derivation of singular configurations (cf. §5).

Planes, Homographies and Calibration. We consider the use of one or several planar objects for calibration. When we talk about *calibration planes*, we mean the supports of planar calibration objects. The restriction of perspective projection to points (or lines) on a specific plane takes on the simple form of a 3×3 homography that depends on the relative position of camera and plane and the camera's intrinsic parameters. Without loss of generality, we may suppose that the calibration plane is the plane $Z = 0$. This way, the homography can be derived from the projection matrix \mathbf{P} by dropping the third column in equation (1):

$$\mathbf{H} \sim \mathbf{K}\mathbf{R} \begin{pmatrix} 1 & 0 & \\ 0 & 1 & -\mathbf{t} \\ 0 & 0 & \end{pmatrix}. \quad (2)$$

The homography can be estimated from four or more point or line correspondences. It can only be sensibly decomposed as shown in equation (2), if the metric structure of the plane is known (up to scale is sufficient), i.e. if the coordinates of points and lines used for computing \mathbf{H} are given in a metric frame.

Equation (2) suggests that the 8 coefficients of \mathbf{H} (9 minus 1 for the arbitrary scale) might be used to estimate the 6 pose parameters \mathbf{R} and \mathbf{t} , while still delivering 2 constraints on the calibration \mathbf{K} . These constraints allow us to calibrate the camera, either partially or fully, depending on the number of calibration planes, the number of images, the number of intrinsic parameters to be computed and on singularities.

3 Principle of Plane-Based Calibration

Calibration will be performed via the determination of the image of the Absolute Conic (IAC), $\omega \sim \mathbf{K}^{-\top}\mathbf{K}^{-1}$, using plane homographies. As mentioned previously, we consider pixels to be rectangular, and thus the IAC has the following form (after appropriate scaling):

$$\omega \sim \begin{pmatrix} 1 & 0 & -u_0 \\ 0 & \tau^2 & -\tau^2 v_0 \\ -u_0 & -\tau^2 v_0 & \tau^2 f^2 + u_0^2 + \tau^2 v_0^2 \end{pmatrix}. \quad (3)$$

The calibration constraints arising from homographies can be expressed and implemented in several ways. For example, it follows from equation (2) that:

$$\mathbf{H}^\top \omega \mathbf{H} \sim \mathbf{H}^\top \mathbf{K}^{-\top} \mathbf{K}^{-1} \mathbf{H} \sim \begin{pmatrix} 1 & 0 & -t_1 \\ 0 & 1 & -t_2 \\ -t_1 & -t_2 & \mathbf{t}^\top \mathbf{t} \end{pmatrix}.$$

The camera position \mathbf{t} being unknown and the equation holding up to scale only, we can extract exactly two different equations in ω that prove to be homogeneous linear:

$$\mathbf{h}_1^\top \omega \mathbf{h}_1 - \mathbf{h}_2^\top \omega \mathbf{h}_2 = 0 \quad \mathbf{h}_1^\top \omega \mathbf{h}_2 = 0, \quad (4)$$

where \mathbf{h}_i is the i th column of \mathbf{H} . These are our basic calibration equations. If several calibration planes are available, we just include the new equations into a linear equation system. It does not matter if the planes are seen in the same view or in several views or if the same plane is seen in several views, provided the calibration is constant (this restriction is relaxed in the next section). The equation system is of the form $\mathbf{A}\mathbf{x} = \mathbf{0}$, with the vector of unknowns $\mathbf{x} = (\omega_{11}, \omega_{22}, \omega_{13}, \omega_{23}, \omega_{33})^\top$. After having determined \mathbf{x} , the intrinsic parameters are extracted via:

$$\begin{aligned} \tau^2 &= \frac{\omega_{22}}{\omega_{11}} & u_0 &= -\frac{\omega_{13}}{\omega_{11}} & v_0 &= -\frac{\omega_{23}}{\omega_{22}} \\ f^2 &= \frac{\omega_{11}\omega_{22}\omega_{33} - \omega_{22}\omega_{13}^2 - \omega_{11}\omega_{23}^2}{\omega_{11}\omega_{22}^2} \end{aligned} \quad (5)$$

4 A General Calibration Algorithm

We describe now how the basic principle can be extended in two important ways. First, we show that prior knowledge of intrinsic parameters can be easily included. Second, and more importantly, we show how the scheme can be applied for calibrating cameras with variable intrinsic parameters.

4.1 Prior Knowledge of Intrinsic Parameters

Let \mathbf{a}_i be the i th column of the design matrix \mathbf{A} of the linear equation system described in the previous section. We may rewrite the equation system as:

$$\omega_{11}\mathbf{a}_1 + \omega_{22}\mathbf{a}_2 + \omega_{13}\mathbf{a}_3 + \omega_{23}\mathbf{a}_4 + \omega_{33}\mathbf{a}_5 = \mathbf{0}.$$

Prior knowledge of, e.g. the aspect ratio τ , allows us via equation (5) to eliminate one of the unknowns, say ω_{22} , leading to the reduced linear equation system:

$$\omega_{11}(\mathbf{a}_1 + \tau^2\mathbf{a}_2) + \omega_{13}\mathbf{a}_3 + \omega_{23}\mathbf{a}_4 + \omega_{33}\mathbf{a}_5 = \mathbf{0}.$$

Prior knowledge of u_0 or v_0 can be dealt with similarly. The situation is different for the focal length f , due to the complexity of equation (5): prior knowledge of f allows to eliminate unknowns only if the other parameters are known, too. However, this is not much of an issue – it is rarely the case that the focal length is known beforehand while the other intrinsic parameters are unknown.

4.2 Variable Intrinsic Parameters

We make two assumptions that are not very restrictive but eliminate useless special cases to deal with. First, we consider the aspect ratio to be constant for a given camera. Second, the principal point may vary, but only in conjunction with the focal length. Hence, we consider two modes of variation: only f varies or f, u_0 and v_0 vary together.

If we take into account the calibration equations arising from a view for which it is assumed that the intrinsic parameters have changed with respect to the preceding view (e.g. due to zooming), we just have to introduce additional unknowns in x and columns in A . If only the focal length is assumed to have changed, a new unknown ω_{33} is needed. If in addition the principal point is supposed to have changed, we add also unknowns for ω_{13} and ω_{23} (cf. equation (3)). The corresponding coefficients of the calibration equations have to be placed in additional columns of A .

Note that the consideration of variable intrinsic parameters does not mean that we have to assume different values for *all* views, i.e. there may be views sharing the same intrinsics, sharing only the aspect ratio and principal point, or sharing the aspect ratio alone.

4.3 Complete Algorithm

The complete algorithm consists of the following steps:

1. Compute plane homographies from feature correspondences.
2. Construct the equation matrix A according to the directions outlined in §§3.4.1 and 4.2.
3. Ensure good numerical conditioning of A (see below).
4. Solve the equation system to least squares by any standard method and extract the intrinsic parameters from the solution as shown in equation (5).

Conditioning. We may improve the conditioning of A by the standard technique of rescaling rows and columns [5]. In practice, we omit row-wise rescaling for reasons explained below. Columns are rescaled such as to have equal norms. The coefficients of the solution vector of the modified equation system have to be scaled accordingly to obtain the solution of the original problem. In our experiments, this rescaling proved to be crucial to obtain reliable results.

As for rescaling rows, this proves to be delicate in our case, since occasionally there are rows with all coefficients very close to zero. Rescaling these rows will hugely magnify noise and lead to unreliable results.

Comments. The described calibration algorithm requires mainly the least squares solution of a single linear equation system. Naturally, the solution may be optimized subsequently using non linear least squares techniques. This optimization should be done simultaneously for the calibration and the pose parameters, that may be initialized in a straightforward manner from the linear calibration results. For higher accuracy, estimation of optical distortion parameters should be included.

Minimal Cases. Each view of a calibration object provides two calibration equations. Hence, in the absence of singularities, the following minimal calibration schemes may be realized: with a single view of a single plane, we might calibrate the aspect ratio and focal length, provided

the principal point is given. With two views of a single plane, or one view of two planes we can fully calibrate the camera. Three views of a single plane, taken by a zooming camera, enable calibration of the 3 different focal lengths, as well as the constant aspect ratio and principal point.

5 Singularities

The successful application of any algorithm requires awareness of singularities. This helps avoiding situations where the result is expected to be unreliable or restricting the problem at hand to a solvable one. We describe here the singularities of calibration from one or two planes.

Due to lack of space, we are only able to give a sketch of the derivations. A first remark is that only the relative orientation of planes and camera is of importance for singularities, i.e. the position and the actual intrinsic parameters do not influence the existence of singularities. A second observation is that planes that are parallel to each other provide exactly the same information as a single plane with the same orientation (except that more feature correspondences may provide a higher robustness in practice). So, as for the case of two calibration planes, we omit dealing with parallel planes and instead refer to the one-plane scenario.

Since the calibration equations are linear, singularities imply the existence of a linear family of solutions for the IAC ω . Hence, there is also a degenerate conic ω' , i.e. a conic consisting of the points on two lines only. Let us note that any conic that satisfies the calibration equations (4), contains the projections of the circular points of the calibration planes. Naturally, this is also valid for ω' . If we exclude the planes of being parallel to each other (cf. the above discussion), the two lines making up ω' are nothing else than the vanishing lines of the calibration planes. There is one point left to consider: since we are considering rectangular pixels, the IAC is required to be of the form (3), i.e. its coefficient ω_{12} is zero. Geometrically, this is equivalent to the conic being symmetric with respect to a vertical and a horizontal line (this is referred to as “reflection constraint” in table 2). Based on these considerations, it is a rather mechanical task to derive all possible singularities.

All singularities for one- and two-plane calibration and for different levels of prior knowledge are described in tables 1 and 2. We reveal which of the intrinsic parameters can/can't be estimated uniquely. The tables contain columns for τf and f which stand for the calibrated focal length, measured in horizontal and vertical pixel dimensions respectively. In some cases it is possible to compute, e.g. τf , but not to compute τ or f individually.

A general observation is that a plane parallel to the image plane, allows to estimate the aspect ratio, but no other parameters. Generally speaking, the more regular the geometric configuration is, the more singularities may occur.

Prior	Position of calibration planes	τ	τf	f	u_0	v_0
None	One plane is parallel to the image plane	cf. case of known τ in table 1				
	General case of planes satisfying reflection constraint (see caption)	-	-	-	-	-
	Both planes are parallel to the u axis Same absolute incidence angle with respect to image plane	-	-	-	+	-
	Both planes are parallel to the v axis Same absolute incidence angle with respect to image plane	-	-	-	+	+
	Vanishing lines intersect "above" the principal pt. i.e. at a point $(u_0, v, 1)$	-	-	-	+	-
	Vanishing lines intersect at a point $(u, v_0, 1)$	-	-	-	-	+
	Both planes are perpendicular to image (and satisfy reflection constraint)	-	-	-	+	+
u_0, v_0	At least one plane is parallel to the image plane	cf. case of known τ, u_0, v_0 in table 1				
	Both planes are perpendicular to the image (and satisfy reflection constr.)	-	-	-	+	+
τ	One plane is parallel to the image plane	cf. case of known τ in table 1				
τ, u_0, v_0	One plane is parallel to the image plane	cf. case of known τ, u_0, v_0 in table 1				

Table 2. Singularities of calibration from two planes. The cases of parallel planes are not displayed, but may be consulted in the appropriate parts of table 1 on one-plane calibration. In all configurations not represented here, all intrinsic parameters can be estimated. By "reflection constraint" we mean that the vanishing lines of the two planes are reflections of each other by both a vertical and a horizontal line in the image.

Prior	Pos. of cal. plane	τ	τf	f	u_0	v_0	
u_0, v_0	Parallel to image pl.	+	-	-	+	+	
	Perpend. to image pl. parallel to u axis parallel to v axis else	-	+	-	+	+	
		-	-	+	+	+	
		-	-	-	+	+	
		+	+	+	+	+	
	Else	parallel to u axis	-	-	-	+	+
		parallel to v axis	-	-	-	+	+
else		+	+	+	+	+	
τ	Parallel to u axis	+	-	-	+	-	
	Parallel to v axis	+	-	-	-	+	
	Else	+	-	-	-	-	
τ, u_0, v_0	Parallel to image pl.	+	-	-	+	+	

Table 1. Singularities of calibration from one plane. Here, parallelism to the image plane's u or v axis means parallelism in 3-space.

6 Experimental Results

We performed a number of experiments with simulated and real data, in order to quantify the performance of our method, to motivate its use in applications described in the following section and to exhibit singularities.

6.1 Simulated Experiments

For our simulated experiments, we used a diagonal calibration matrix with $f = 1000$ and $\tau = 1$. Calibration is performed using the projections of the 4 corner points of squares of size 40cm. The distance of the calibration squares to the camera is chosen such that the projections roughly fill the 512×512 image plane. The projections of the corner points are perturbed by centered Gaussian noise of 0 to 2 pixels variance.

We only display graphs showing the behavior of our algorithm with respect to other parameters than noise; note however that in all cases, the behavior with respect to noise is nearly perfectly linear. The data in the graphs shown stem from experiments with a noise level of 1 pixel. The errors shown are absolute ones (scaled by 1000 for the aspect ratio). Each point in a graph represents the *median* error of 1000 random experiments. The graphs of the mean errors are similar but less smooth.

One plane seen in one view. The scenario and results are shown in the upper part of figure 1. Calibration is performed for different orientations of the square, ranging from 0° (parallel to the image plane) to 90° (perpendicular to the image plane). Given the principal point, we calibrated the aspect ratio and the focal length. An obvious observation is the presence of singularities: the error of the aspect ratio increases considerably as the calibration square tends to be perpendicular to the image plane (90°). The determination of the focal length is impossible for the extreme cases of parallelism and perpendicularity. Note that these observations are all predicted by table 1. In the range of $[30^\circ, 70^\circ]$, the relative error for the focal length is below 1%, while the aspect ratio is estimated correctly within 0.01%.

Two planes seen in one view. Calibration is performed with a camera rotating about its optical axis by 0° to 90° . Two planes with an opening angle of 90° are observed (cf. lower part of figure 1). Plane-based calibration is now done without any prior knowledge of intrinsic parameters. For comparison, we also calibrate with a standard method [3], using full 3D coordinates of the corner points as input.

The standard calibration approach is insensitive to rotation about the optical axis. As for the plane-based method, the singularities for the estimation of the aspect ratio and the

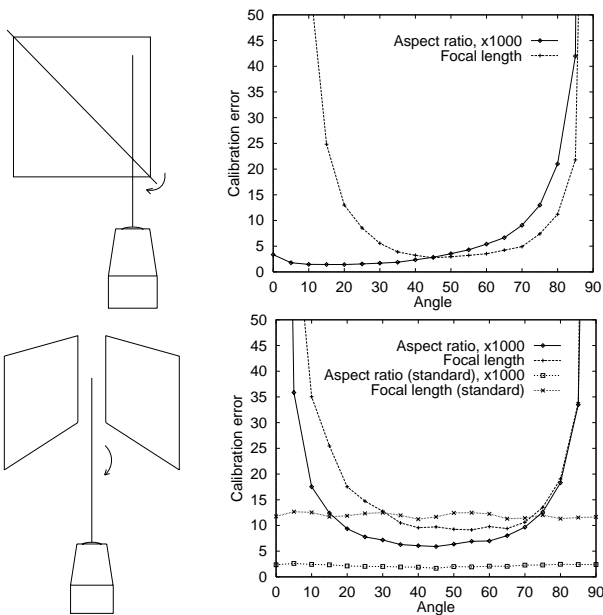


Figure 1. Simulation scenarios and results.

focal length for angles of 0° and 90° are predicted by table 2. As for the intermediate range of orientations, the estimation of the aspect ratio by the plane-based method is 3 to 4 times worse than with the standard approach, although it is still quite accurate. As for the focal length, the plane-based estimate is even slightly better between 30° and 70° . The error graphs for u_0 and v_0 are not shown; for both methods they are nearly horizontal (i.e. there is no singularity), the errors of the plane-based estimation being about 30% lower than with the standard approach.

6.2 Calibration Grid

We calibrated a camera from images of a 3D calibration grid with targets arranged in three planes (cf. figure 2). For comparison, calibration was also carried out using a standard method [3]. We report the results of two experiments. First, 4 images were taken from different positions, but with fixed calibration. The camera was calibrated from single views in different modes: standard calibration using all points or points from two planes only, plane-based calibration from one, two or three planes with different levels of prior knowledge (cf. table 3). Prior values were taken from the results of standard calibration.

Table 3 shows the mean and standard deviation of the results for the focal length, computed over the 4 views and over all combinations of planes. We note that even the one-plane method gives results very close to those of the standard method that uses all points and their full 3D coordinates. The precision of the plane-based results is lower than for full standard calibration, though comparable to standard calibration using two planes. The results are very accurate despite the proximity to singular configurations. This may




Figure 2. Calibration grid and lab scene.

Method	f
Standard calibration from three planes	1041.4 ± 0.6
Standard calibration from two planes	1042.1 ± 3.3
One plane, u_0, v_0 known	1044.5 ± 9.0
One plane, τ, u_0, v_0 known	1041.2 ± 3.7
Two planes, nothing known	1043.6 ± 4.7
Two planes, τ known	1040.7 ± 2.7
Two planes, u_0, v_0 known	1040.2 ± 2.5
Two planes, τ, u_0, v_0 known	1040.3 ± 2.1
Three planes, nothing known	1039.9 ± 0.7

Table 3. Results for calibration grid.

be attributed to the high accuracy of target extraction.

For the second experiment, we took images at 5 different zoom positions. The camera was calibrated using the 5×3 planes simultaneously, where for each zoom position an individual focal length and principal point were estimated. Table 4 shows the results for the focal lengths (a value of 1000 corresponds to about 7.5mm), compared to those of standard calibration, averaged over single views. The deviation increases with the focal length but stays below 1%.

6.3 Lab Scene

A pair of images of an ordinary lab scene were taken. A rectangular part of a computer tower (cf. figure 2) was used for calibration. Subsequently, the pose of the views with respect to the calibration plane was determined. The three points shown in figure 2 were triangulated and their 3D distances measured and compared to hand-measured ones. The differences for the pairs (1,2), (1,3) and (2,3) were 4mm, 3mm and 0mm respectively, for absolute distances of 275mm, 347mm and 214mm. These results are about as good as we might expect: the edges of the rectangular patch are rounded, thus not reliably extracted in the images. The measured point distances are “extrapolated” from this rectangle, thus amplifying the errors of edge extraction. From the views’ calibration and pose, we computed the epipolar geometry and found that the distance of points to corresponding epipolar lines was about 1 pixel, even at the borders of the images.

This simple experiment highlights two issues. First, be-

Method	Focal lengths				
Standard	714.7	1041.4	1386.8	1767.4	2717.2
Planes	709.9	1042.7	1380.2	1782.8	2702.0

Table 4. Results for variable focal length.

sides calibrating the views, we readily obtain their pose in a *metric* 3D frame. Second, we obtain reasonable estimates of matching constraints, potentially for distant views.

7 Applications

Cheap Calibration Tool. Planar patterns are easy to produce, while enabling a reasonably reliable calibration.

Ground Plane Calibration. We have successfully performed experiments with images of traffic scenes. Ground plane calibration from road markers is used to restrict the pose of vehicles to be detected and tracked.

Reconstruction of Piecewise Planar Objects from Single Views. Using geometrical constraints such as coplanarity, parallelism, right angles etc., 3D objects may be reconstructed from a single view (see e.g. [10]). Our calibration method requires knowledge of the metric structure of planes. This requirement may be relaxed by simultaneously determining calibration and plane structure, e.g. one view of a rectangle allows to determine the focal length and the ratio of the rectangle's edge lengths. We are using this in combination with the mentioned geometrical constraints to reconstruct objects from a single image.

Reconstruction of Indoor Scenes. Our calibration method is the central part of ongoing work on a system for interactive multi-view 3D reconstruction of indoor scenes, similar in spirit to the approaches presented in [10, 11]. The main motivation for using plane-based calibration is to make a compromise between requirements on flexibility, user interaction and implementation cost. We achieve flexibility by not requiring off-line calibration: our calibration patterns, planar objects, are omnipresent in indoor scenes. The amount of user interaction is rather little: we usually use rectangles as calibration objects; they have to be delineated in images and their edge lengths measured. By identifying planar patterns across distant views, we not only can simultaneously calibrate many views but also compute a global initial pose of many views to bootstrap, e.g. wide baseline matching. This scheme relies on methods that are relatively simple to implement and might provide a useful alternative to completely automatic techniques such as [9] that are more flexible but more difficult to realise.

Augmented Reality. A nice and useful application of plane-based calibration and pose estimation is presented in [6]. Rectangular plates are used to mark the position of non planar objects to be added to a video sequence, which is in some way a generalisation of "overpainting" planar surfaces in videos by homography projection of a desired pattern.

Plane-based methods may also be used for blue screening; attaching calibration patterns on the blue screen allows to track camera pose and calibration and thus to provide input for positioning objects in augmented reality.

8 Conclusion

We presented a general and easy to implement plane-based calibration method that is suitable for calibrating variable intrinsic parameters and that copes with any number of calibration planes and views. Experimental results are very satisfactory. For the basic cases of one or two planes, we gave an exhaustive list of singularities. Several applications of plane-based calibration were described. An analytical error analysis might be fruitful, i.e. examining the influence of feature extraction errors on calibration accuracy.

An extended version of this paper can be retrieved at <http://www.cvg.cs.reading.ac.uk/~pfs/plane.ps.gz>.

References

- [1] A. Criminisi, I. Reid, A. Zisserman, "Duality, Rigidity and Planar Parallax," *ECCV*, pp. 846-861, 1998.
- [2] O. Faugeras, "Stratification of Three-Dimensional Vision: Projective, Affine and Metric Representations," *Journal of the Optical Society of America A*, 12, pp. 465-484, 1995.
- [3] O. Faugeras, G. Toscani, "Camera Calibration for 3D Computer Vision," *Int. Workshop on Machine Vision and Machine Intelligence*, pp. 240-247, 1987.
- [4] R.J. Holt, A.N. Netravali, "Camera Calibration Problem: Some New Results," *CVIU*, 54 (3), pp. 368-383, 1991.
- [5] A. Jennings, J.J. McKeown, *Matrix Computation*, 2nd edition, Wiley, 1992.
- [6] M. Jethwa, A. Zisserman, A. Fitzgibbon, "Real-time Panoramic Mosaics and Augmented Reality," *BMVC*, pp. 852-862, 1998.
- [7] R.K. Lenz, R.Y. Tsai, "Techniques for Calibration of the Scale Factor and Image Center for High Accuracy 3-D Machine Vision Metrology," *PAMI*, 10 (5), pp. 713-720, 1988.
- [8] F. Pedersini, A. Sarti, S. Tubaro, "Multi-Camera Acquisitions for High-Accuracy 3D Reconstruction," *SMILE Workshop*, pp. 124-138, 1998.
- [9] M. Pollefeys, R. Koch, M. Vergauwen, L. Van Gool, "Metric 3D Surface Reconstruction from Uncalibrated Image Sequences," *SMILE Workshop*, pp. 139-154, 1998.
- [10] H.-Y. Shum, R. Szeliski, S. Baker, M. Han, P. Anandan, "Interactive 3D Modeling from Multiple Images Using Scene Regularities," *SMILE Workshop*, pp. 236-252, 1998.
- [11] R. Szeliski, P.H.S. Torr, "Geometrically Constrained Structure from Motion: Points on Planes," *SMILE Workshop*, pp. 171-186, 1998.
- [12] R.Y. Tsai, "A Versatile Camera Calibration Technique for High-Accuracy 3D Machine Vision Metrology Using Off-the-Shelf TV Cameras and Lenses," *IEEE Journal of Robotics and Automation*, 3 (4), pp. 323-344, 1987.
- [13] B. Triggs, "Autocalibration from planar scenes," *ECCV*, pp. 89-105, 1998.
- [14] G.-Q. Wei, S.D. Ma, "A Complete Two-Plane Camera Calibration Method and Experimental Comparisons," *ICCV*, pp. 439-446, 1993.