Simple Calibration of Non-overlapping Cameras with a Mirror

Ram Krishan Kumar[1], Adrian Ilie[1], Jan-Michael Frahm[1]
[1]Department of Computer Science
University of North Carolina, Chapel Hill, USA
{ramkris, adyilie, jmf, marc}@cs.unc.edu

Marc Pollefeys[1,2]
[2]Department of Computer Science
ETH Zurich, Switzerland
marc.pollefeys@inf.ethz.ch

Abstract

Calibrating a network of cameras with non-overlapping views is an important and challenging problem in computer vision. In this paper, we present a novel technique for camera calibration using a planar mirror. We overcome the need for all cameras to see a common calibration object directly by allowing them to see it through a mirror. We use the fact that the mirrored views generate a family of mirrored camera poses that uniquely describe the real camera pose. Our method consists of the following two steps: (1) using standard calibration methods to find the internal and external parameters of a set of mirrored camera poses, (2) estimating the external parameters of the real cameras from their mirrored poses by formulating constraints between them. We demonstrate our method on real and synthetic data for camera clusters with small overlap between the views and non-overlapping views.

1. Introduction

Camera calibration is a well-studied problem in computer vision. It is a fundamental task that enables many computer vision applications such as 3D reconstruction, motion estimation, augmented reality, visual surveillance and telepresence. All these applications use multiple cameras and benefit greatly of having all of them calibrated in a single coordinate frame. However, an inherent problem when using camera networks is that the cameras may have non-overlapping fields of view (FOVs). Many camera networks calibration techniques assume that cameras have overlapping FOVs and some even require that the cameras be synchronized.

To overcome these limitations, we propose a novel calibration method that relies on a planar mirror to make a single calibration object visible to all cameras. The mirror removes the need for cameras to observe the calibration object directly; instead we move the mirror to create different views of the calibration object. The calibration object remains fixed with respect to the cameras during the cali-


Figure 1. Calibrating with a planar mirror. The camera sees the pattern in the mirror during the calibration step. Left: Image during calibration. Right: Camera poses recovered using our calibration.

bration procedure, providing a unique reference frame for computing a family of mirrored camera poses. This method introduces the four degrees of freedom of each mirror pose (3D normal vector and distance from the camera) as free parameters into the calibration process. A naïve method could place markers on the mirror to estimate its relative pose with respect to the camera. The mirror pose could then be used to correct for the reflection introduced by the mirror. Our proposed method introduces a novel set of constraints that demonstrate that the mirror pose is not needed for calibration. Instead we use the family of mirrored camera poses that are associated with the real camera pose as the camera observes the calibration object through a mirror. We show that two mirrored images of the calibration pattern are sufficient to recover the external calibration of the camera up to a finite number of solutions. However, because the resulting equations are non-linear, we prefer to linearize them by adding more unknowns, which then requires at least five mirrored images for the entire calibration process.

We show that the internal camera parameters and the radial distortion coefficients remain the same when using a mirror during the calibration process. Our calibration method can then use any available camera calibration technique as a first step to determine the internal camera parameters, the radial distortion coefficients and the mirrored camera poses. In the experiments for this paper we use

1

Bouguet's calibration toolbox [4], which is based on the approach proposed by Zhang [23]. The key difference from most calibration techniques is that instead of moving the calibration object or the camera, we move the mirror and leave the camera and calibration object fixed. The internal camera parameters and the mirrored camera poses are used in the second step of our calibration method to recover the real camera pose. The rest of the paper is organized as follows. Section 2 describes the related prior work on calibration of cameras and camera networks. Section 3 briefly outlines the notations we use in the rest of the paper. Section 4 introduces our novel calibration constraints for cameras using a planar mirror and also discusses the details of how our calibration procedure is applied to calibrate a network of non-overlapping cameras. Section 5 shows the result of our calibration method on both synthetic and real data, followed by a detailed analysis of its performance, as well as some practical considerations and rules of thumb to be followed during the calibration procedure. Finally we conclude the paper in Section 6.

2. Related Work

Various methods have been proposed to calibrate cameras using calibration objects. Generally there are two classes of methods: the first class uses a calibration object and provides the external camera parameters within the coordinate system of the calibration object as a byproduct of the internal calibration; the second class does not use any calibration object and delivers an external camera calibration up to scale.

We first review the methods that can be used in conjunction with our method. A variety of methods use fixed 3D geometry [9, 19]. Orthonormal properties of the rotation matrix have been used in plane-based approaches [11, 10, 16, 23, 17]. A planar pattern viewed from at least three different orientations is used in [23]. Other calibration objects include: spheres [2, 22], circles [12], surfaces of revolution [7], shadows [5] and even fixed stars in the night sky [14]. Most of these methods are based on the constraints provided by vanishing points along perpendicular directions.

The method in [1] uses a moving plate to calibrate multi-camera environments and does not require a 3D calibration object with known 3D coordinates. Kitahara et al. [13] calibrated their large scale multi-camera environment by using a classical direct method [19]. The necessary 3D points are collected by a combined use of a calibration board and a 3D laser-surveying instrument. Svoboda et al. [21] have calibrated a system of at least three cameras by obtaining a set of virtual 3D points made by waving a bright spot throughout the working volume. Baker and Aloimonos [3] proposed a calibration method for a multi-camera network which requires a planar pattern with a precise grid. Lee et al. [15] established a common coordinate frame for a sparse set of cameras so that all cameras observe a common dominant plane. They tracked objects moving in this plane and from their trajectories they estimated the external parameters of the cameras in a single coordinate system. Sinha et al. [20] have calibrated a camera network by using epipolar geometry constraints derived from dynamic silhouettes. All the above systems are networks with cameras that have overlapping views or at least pairwise overlapping views. Our technique can be applied without this limitation.

In addition to the above methods for camera calibration through the use of calibration objects there is work on the external camera calibration through camera motion in an unknown scene. Caspi and Irani [6] proposed an approach to calibrate a camera system consisting of two cameras with approximately the same center of projection and non-overlapping views. They capture two image sequences and recover the spatial and temporal transformations between them using the correlated temporal behavior induced by moving the cameras jointly in space. Requiring a common optical center and rigid motion are strong limitations that our approach overcomes. Esquivel [8] proposed an approach to estimate the relative external parameters up to a common scale factor for a multi-camera rig with non-overlapping views of internally calibrated cameras. The approach computes the camera motion up to scale independently for each camera. Afterwards all reconstructions are aligned with a similarity transformation.

3. Notations

In this paper we will denote a 2D image point by $\boldsymbol{m} = [x \quad y]^T$, and a 3D point by $\boldsymbol{M} = [X \quad Y \quad Z]^T$. In the homogeneous coordinates, we write them as $\overline{\boldsymbol{m}} = [x \quad y \quad 1]^T$ and $\overline{\boldsymbol{M}} = [X \quad Y \quad Z \quad 1]^T$ respectively. We model cameras using the pinhole model. The 3D point M and its image projection m into a camera with center of projection $\boldsymbol{C}$ and with orientation $\boldsymbol{R}$ are related by:

$$\overline{\boldsymbol{m}} = \boldsymbol{P}\overline{\boldsymbol{M}}, \tag{1}$$

where $\boldsymbol{P} = \boldsymbol{K}[\boldsymbol{R}^T \quad -\boldsymbol{R}^T\boldsymbol{C}]$ is the camera projection matrix and $\boldsymbol{K}$ is a matrix of the form:

$$\boldsymbol{K} = \begin{bmatrix} \alpha & \gamma & u_0 \\ 0 & \beta & v_0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{2}$$

(u_0, v_0) are the 2D image coordinates of the principal point and α, β are the focal length in pixels in the x and y directions respectively. γ is the skew factor between the two axes, which in practice is almost always set to zero. $\boldsymbol{K}$ represents the intrinsic camera parameters, and $\boldsymbol{R}$ and $\boldsymbol{C}$ represent the extrinsic camera parameters.

4. Calibrating a Camera Cluster

We first introduce our proposed constraints for the calibration using the mirrored images of a calibration pattern. Afterwards we explain the details of our novel calibration technique based on those constraints.

4.1. Using a Planar Mirror to Calibrate

Consider a camera pose characterized by its center of projection C and its orientation R. We can place a mirror inside the camera's field of view to enable it to see a calibration pattern that is outside its field of view. This effectively means the camera observes a mirrored pattern and is equivalent to the mirrored camera observing the original pattern (see Figure 2).

Figure 2. A real camera observing the image of a 3D point in a mirror is equivalent to a mirrored camera observing the real point itself.

We describe the mirror pose by its plane normal n and the distance to the camera center of projection, d. By changing the position and orientation of the mirror, we generate a set of mirror poses Π_i with $i = 1, \ldots, l$ as:

$$\Pi_i = [n_i{}^{T}, -d_i] \quad \text{for } i = 1, \ldots, l \tag{3}$$

For each mirror pose we obtain one image of the calibration pattern while keeping the camera and the pattern fixed. For a particular mirror pose the mirrored camera center C'_i is given by

$$C'_i = C - 2d_i n_i. \tag{4}$$

We also make the following two crucial observations:

1. The mirrored camera is now left-handed if the original camera was right-handed and vice-versa.

2. Due to the duality of points and cameras, a camera placed at point C and looking at point X' is equiv-

alent[1] to a mirrored camera placed at image point C' and looking at the object point X.

It follows that the image coordinates of x (in the mirrored camera image) and x' (in the real camera image) are the same, which means that the intrinsics and the radial distortion coefficients of both cameras are identical. Consequently, applying any standard calibration method to the mirrored images computes the correct intrinsic camera parameters and radial distortion coefficients for the mirrored camera positions, as well as for the real camera, except that the handedness of the coordinate system is changed. The only remaining unknown is the real camera pose.

4.2. Computing the Real Camera Pose from the Mirrored Camera Poses

In this subsection, we develop the constraints that the mirrored camera poses enforce on the real camera pose. Given a camera at position C with orientation R observing a calibration pattern through a moving planar mirror, it follows that the resulting mirrored camera poses belong to a 3-DoF family of camera poses generated by the three degrees of freedom of the mirror plane $\Pi_m = [n_m{}^\top, -d_m]$. The mirrored camera center C' can be computed using (4) as $C' = C - 2d_m n_m$. Additionally, the right-handed orientation of the real camera R is uniquely determined by the left-handed orientations R' of the mirrored cameras.

We now derive the constraints that the mirrored camera poses enforce on the real camera pose. Let r_k and r'_k be k^{th} column vectors of the rotation matrix of the original camera and a mirrored camera, respectively. Then the following proposition holds.


Figure 3. The relationship between a mirrored pose and the real pose of a camera.

---

[1] If we don't consider the change in the handedness of the cameras then images observed by them are flipped L-R with respect to each other.

Proposition 4.1 *The vector connecting point C to C' is orthogonal to $(r_k + r'_k)$, iff r_k is not orthogonal to the mirror plane for $k = \{1, 2, 3\}$.*

Proof $C'C$ is parallel to the normal of the mirror. Let r_k be at an angle $\theta_k (\neq 0)$ with respect to the normal of the mirror. Using the laws of reflection, it follows that r'_k is at an angle $180 - \theta_k$ to the mirror normal. Since $|r_k| = |r'_k|$, $(r_k + r'_k)$ will bisect the angle between them, which is $(180 - 2\theta_k)$ and will enclose an angle of $90 - \theta_k$ with vectors r'_k and r_k. $(r_k + r'_k)$ will then be at an angle of $(180 - \theta_k) - (90 - \theta_k)$ or $\theta_k + (90 - \theta_k)$ to the mirror normal.

Using proposition 4.1, three equations of the following form can be formulated, for $k = \{1, 2, 3\}$:

$$(C'^\top - C^\top)(r'_k + r_k) = 0$$
$$or \; C'^\top r'_k + C'^\top r_k - C^\top r'_k - C^\top r_k = 0 \qquad (5)$$

where all the vectors are of size 3×1. Since C and r_k with $k = 1, \ldots, 3$ are unknowns, the resulting three equations (5) are non-linear. However, if we define

$$C^\top r_k = s_k \qquad (6)$$

for $k = \{1, 2, 3\}$ and substitute into equations (5), we obtain them as linear constraints with three additional unknowns. These are now the basic linear equations for the position of the real cameras center and its orientation, which we can rewrite as:

$$AX = B, \qquad (7)$$

where

$$A = \begin{bmatrix} -r'_1{}^\top & C'^\top & 0 & 0 \\ -r'_2{}^\top & 0 & C'^\top & 0 & -I \\ -r'_3{}^\top & 0 & 0 & C'^\top \end{bmatrix} \qquad (8)$$

$$X = [C^\top, r_1{}^\top, r_2{}^\top, r_3{}^\top, s_1, s_2, s_3]^\top \qquad (9)$$
$$B = -[C'^\top r'_1, C'^\top r'_2, C'^\top r'_3]^\top \qquad (10)$$

For m mirror poses, $3m$ such linear equations can be stacked and we have A as a $3m \times 15$ matrix and B as a $3m \times 1$ matrix. Since there are 15 unknowns and each image provides us 3 constraints, at least 5 images are needed to solve for the position C and orientation R of the real camera. Furthermore, the orientation R of the real camera should satisfy:

$$[r_1 \;\; r_2 \;\; r_3]^\top [r_1 \;\; r_2 \;\; r_3] = I \qquad (11)$$

11 generates an additional six quadratic constraints for the orientation of the real camera:

$$|r_1| = |r_2| = |r_3| = 1 \qquad (12)$$
$$r_1{}^\top r_2 = r_2{}^\top r_3 = r_3{}^\top r_1 = 0 \qquad (13)$$

(11) and (6) lead to nine quadratic constraints. Together with the $3m$ linear constraints, they are used to determine the 15 unknowns. It follows that $m = 2$, or two distinct images are sufficient to obtain the six degrees of freedom of the real cameras pose. Note that the non-linear constraints are of degree two and could provide up to two solutions. During the calibration we typically obtain the solution through an over-constrained linear system of equations. Afterwards we enforce the nine non-linear constraints to optimize the solution.

The next section introduces the details of our calibration method.

4.3. Non-overlapping camera calibration

Our novel camera calibration technique for non-overlapping cameras is a two step process consisting of:

1. **Internal calibration and mirrored external calibration** from the mirrored images captured by leaving the cameras and the calibration pattern fixed, and moving the mirror to enable the cameras to observe the mirrored calibration pattern in various poses. This step can use any standard calibration technique that provides internal and external camera calibration.

2. **Computing the real camera pose** from the family of mirrored camera poses. In this step the proposed constraints are used to obtain the pose of the real camera.

We used the calibration toolbox from [4], which implements the method from [23]. This provides us with the internal parameters and the mirrored camera poses. Next we solve the linear equations (7) for all mirrored cameras C'_k and r_k, s_k, $k = \{1, 2, 3\}$ in a least-squares sense. Once we solved for C and r_k, s_k, $k = \{1, 2, 3\}$ using the linear set of equations, the rotation matrix consisting of $r_k{}'s$ is parameterized into quaternion by approximating the estimated linear rotation matrix $R = [r_1, r_2, r_3]$ with the closest rotation matrix. r_1, r_2, r_3 serve as basis solution vectors for a non-linear optimization to meet (6) and (11).

Bundle adjustment is later used to determine the extrinsic parameters corresponding to the real pose of the camera by minimizing the reprojection error of the mirrored pattern corners. Our bundle adjustment is parameterized in the real camera's parameters and the mirror positions. In order to recover the mirror parameters $(n_k{}^\top, -d_k)$ for each mirror k, we use (4) to compute the initial mirror parameters from the estimated position C and the mirrored camera centers C'_k.

5. Experimental Results

We conducted experiments on both synthetic and real data to test our proposed method. We first review the perfor-
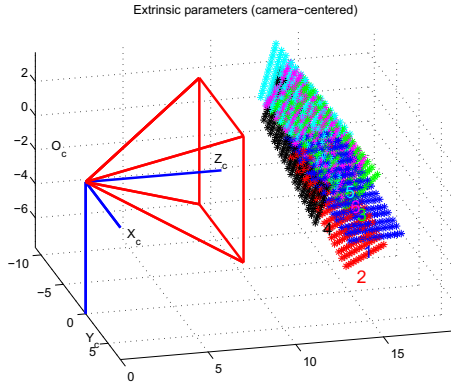
Figure 4. Position of the six mirrored images of the calibration pattern as observed by the camera in the mirror. The mirror normals are randomly generated.

mance evaluation on synthetic data. Afterwards, we discuss our calibration of a real camera system.

5.1. Evaluation on Synthetic Data

We fixed the values of the internal parameters of a synthetic camera as: $\alpha = \beta = 1300, \gamma = 0, u_0 = 320, v_0 = 240$. We set the image resolution to 640×480. We generated images as reflections of a fixed calibration pattern containing the dataset of 256 points from [23]. The planar mirror normal was randomly generated and only the images in which the pattern can be seen by the camera in the mirror were considered for calibration. The position and orientation of the camera were chosen randomly and used as the ground truth in the error analysis.

We added Gaussian noise with zero-mean and standard deviation σ to the generated image points in order to simulate the noise that would be present after reflection in the mirror and subsequent capture by the camera. We then compared the estimated extrinsic camera parameters with the ground truth. Note that the correctness of the intrinsic parameters depends on the calibration method used in the first step of our approach. We varied the noise level from 0.1 to 1.6 pixels, performed 100 independent trials for each noise level and averaged the results. Figures 5 and 6 show the percentage error in the recovery of the extrinsic parameters of the real camera. The errors are around 0.7% for $\sigma = 0.5$, values that occur in practice for real data.

5.2. Evaluation on Real Data

We used our method to calibrate a PointGrey LadyBug2 camera system [18]. It has six Sony CCD cameras with 2.5 mm lenses. Five CCD cameras are positioned in a horizontal ring and one camera points straight up as shown in Figure 7. The image resolution is 1024×768. We used an


Figure 5. The percentage error in recovering the original pose of the real camera from its six mirrored poses.


Figure 6. The percentage error in recovering the rotation vectors of the real camera from its six mirrored poses.

Figure 7. Ladybug: a system of spherical digital video cameras. Five cameras are positioned on a horizontal ring and the sixth camera points upwards

8×7 checkerboard pattern with 42 corner points. The size of the pattern was 649.6 mm $\times$ 648.9 mm. The pattern and the camera were kept fixed. For each camera, the mirror was moved in front of it such that the camera saw the complete pattern in the mirror. Figure 8 shows example images as seen by different cameras. Figure 9 shows a top view of the position of the cameras recovered using our calibration approach. Note the five cameras that lie on the horizontal ring and the sixth camera that points away from the checkerboard pattern. The pattern is located on the $Z = 0$ plane.

Since the reliability of the recovered internal parameters depends on the calibration method used (in our case [23]), we only investigate the stability of our method in terms of

Figure 8. Checkerboard pattern as seen in the mirror by each of the six cameras.

external parameters. In Tables 1 and 2, we show the calibration results (position and orientation, respectively) using 6 through 12 images. The results are very consistent with each other and converge with increasing number of images used. In Figure 10, we compare the mean reprojection er-

Images	3D position (mm)		
6	279.2	-865.1	-1090.5
8	276.8	-866.7	-1099.6
10	277.2	-874.5	-1105.4
11	276.3	-874.6	-1103.6
12	275.8	-874. 1	-1103.4

Table 1. Convergence of position of the real camera after bundle adjustment with regard to the number of images.

Images	Rotation vector (Rodrigues form)		
6	1.7022	-0.8178	1.7300
8	1.7044	-0.8132	1.7366
10	1.6885	-0.8043	1.7449
11	1.7014	-0.8042	1.7452
12	1.7016	-0.8055	1.7430

Table 2. Convergence of orientation of the real camera after bundle adjustment with regard to the number of images.

ror before and after applying bundle adjustment to our initial solution obtained by solving linear system of equations.

5.3. Practical Considerations

Degenerate configurations can arise from both the calibration process using method in [23] and our computation


Figure 9. Relative position of the Ladybug cameras recovered using our calibration method, in the coordinate system defined by the pattern. Note the five cameras that lie on a horizontal ring and the sixth camera that points away from the checkerboard pattern. Top: cameras before bundle adjustment, Bottom: cameras after bundle adjustment


Figure 10. The mean reprojection error before and after applying bundle adjustment.

of the parameters of the real camera from the external parameters of the mirrored cameras presented in section 4.2. In this subsection, we discuss how these degenerate configurations can be avoided in practice.

The calibration process requires that the rotations of each pattern position be independent of each other. Zhang [23] suggests simply changing the orientation of the pattern from one snapshot to another. While this is easy when calibrating a camera directly, it is surprisingly difficult when calibrating using a mirror: the mirror has to be in the field of view of the camera, and oriented such that the pattern's image is reflected into the camera. These requirements may result in little variation in the pattern orientations as seen by the camera in the mirror. In order to avoid this degeneracy, we have found that it is useful to make the pattern small enough (or place it far enough) such that its image in the camera only occupies a fraction of the camera's field of view. Although this may lead to less precision in finding the checkerboard corners, it gives more room to maneuver

the mirror inside the camera's field of view, and allows for larger mirror rotations, which result in more variation in the pattern orientations.

The computation of the real camera parameters from the mirrored camera poses requires that none of the axes of the camera coordinate system be perpendicular to the mirror plane. This is not a common problem for the camera X and Y axes, because the mirror would be parallel to the Z axis, and its image would only occupy a small number of columns or rows of pixels in the camera image. In order to ensure that the camera's Z axis is not perpendicular to the mirror plane, one should try to avoid aiming the mirror in the general direction of the camera. In practice, this means that the pattern should not be placed in close proximity to the cameras.

6. Conclusion

We presented a flexible new method to easily calibrate a cluster of cameras with non-overlapping or barely overlapping views. This technique does not require a camera to see the calibration object directly. The proposed method involves solving a set of linear constraints between a family of mirrored camera poses and the real camera pose. We showed that at least five images are required linearly (otherwise two are enough) in order to recover the real camera position and orientation. We used both synthetic and real data to test the proposed method and showed impressing results. Compared to other methods of calibrating camera clusters, our proposed method is more flexible and does not require any overlapping views between cameras.

7. Acknowledgements

References

[1] Open source computer vision library. *http://www.intel.com/research/mrl/research/opencv/*.

[2] M. Agrawal and L. Davis. Complete camera calibration using spheres: Dual space approach. *IEEE ICCV*, pages 782–789, 2003.

[3] P. T. Baker and Y. Aloimonos. Calibration of a multicamera network. in r. pless, j. santos-victor, and y. yagi, editors. *Omnivis 2003: Omnidirectional Vision and Camera Networks*, 2003.

[4] J.-Y. Bouguet.

[5] X. Cao and H. Foroosha. Camera calibration and light source orientation from solar shadows. *Computer Vision and Image Understanding*, 105:60–72, 2006.

[6] Y. Caspi and M. Irani. Alignment of non-overlapping sequences. *ICCV, Vancouver, BC*, pages II: 76–83, 2001.

[7] C. Colombo, A. Bimbo, and F. Pernici. Metric 3d reconstruction and texture acquisition of surfaces of revolution from a single uncalibrated view. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1:99–114, 2005.

[8] S. Esquivel, F. Woelk, and R. Koch. Calibration of a multicamera rig from non-overlapping views. In *Lecture Notes in Computer Science 4713 (DAGM)*, 2007.

[9] O. Faugeras and Q. Luong. The geometry of multiple images. *MIT Press,Cambridge*, 2001.

[10] P. Gurdjos, A. Crouzil, and R. Payrissat. Another way of looking at plane-based calibration: the centre circle constraint. *Proc. ECCV*, pages 252–266, 2002.

[11] P. Gurdjos and P. Sturm. Methods and geometry for plane-based self-calibration. *IEEE CVPR*, pages 491–496, 2003.

[12] G. Jiang, H. Tsui, L. Quan, and A. Zisserman. Single axis geometry by fitting conies. *IEEE Trans. Pattern Anal. Mach. Intell.*, 10, 2003.

[13] I. Kitahara, H. Saito, S. Akimichia, T. Onno, Y. Ohta, and T. Kanade. *In Computer Vision and Pattern Recognition, Technical Sketches*, 2001.

[14] A. Klaus, J. Bauer, K. Karner, P. Elbischger, R. Perko, and H. Bischof. Camera calibration from a single night sky image. *IEEE CVPR*, pages 151–157, 2004.

[15] L. Lee, R. Romano, and G. Stein. Monitoring activities from multiple video streams: establishing a common coordinate frame. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2000.

[16] D. Liebowitz and A. Zisserman. Combining scene and auto-calibration constraints. *IEEE ICCV*, pages 293–300, 1999.

[17] E. Malis and R. Cipolla. Camera self-calibration from unknown planar structures enforcing the multi-view constraints between collineations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 9:1268–1272, 2002.

[18] PointGrey Research, Inc.

[19] R.Y.Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE J. Robotics Autom*, 3:323–344, 1987.

[20] S. Sinha, M. Pollefeys, and L. McMillan. Camera network calibration from dynamic silhouettes. *IEEE CVPR*, 2004.

[21] T. Svoboda, D. Martinec, and T. Pajdla. A convenient multi-camera self-calibration for virtual environments. *PRESENCE: Teleoperators and Virtual Environments*, 14(4):407–422, August 2005.

[22] H. Zhang, G. Zhang, and K.-Y. Wong. Camera calibration with spheres: linear approaches. *IEEE ICIP*, 2:1150–1153, 2005.

[23] Z. Zhang. A flexible new technique for camera calibration. *IEEE Transcations on Pattern Analysis and Machine Intelligence*, pages 1330–1334, 2000.