

# Omni-rig: Linear Self-recalibration of a Rig with Varying Internal and External Parameters

Assaf Zomet      Lior Wolf      and      Amnon Shashua  
School of Computer Science and Engineering,  
The Hebrew University,  
Jerusalem 91904, Israel  
e-mail: {zomet,lwolf,shashua}@cs.huji.ac.il

## Abstract

*We describe the principles of building a moving vision platform (a Rig) that once calibrated can thereon self-adjust to changes in its internal configuration and maintain an Euclidean representation of the 3D world using only projective measurements. We term this calibration paradigm “Omni-Rig”. We assume that after calibration the cameras may change critical elements of their configuration, including internal parameters and centers of projection. Theoretically we show that knowing only the rotations between a set of cameras is sufficient for Euclidean calibration even with varying internal parameters and unknown translations. No other information of the world is required.*

## 1. Introduction

The projective framework of Structure from Motion (SFM) is supported by a relatively large body of literature on the techniques for taking matching image features (points and lines) across multiple views and producing a projective representation of the three-dimensional (3D) positions of the features in space (the special cases of 2,3 views are theoretically interesting, for a review see [10]). There are tasks, such as visual recognition by alignment and image transfer in Photogrammetry in which a projective representation is sufficient (for example, [18, 13]). However, in mensuration and model-based computer graphics an Euclidean representation is necessary. To that end a growing body of literature is dedicated to the problem of obtaining an Euclidean representation [7, 17, 12, 24, 1]. We divide the various approaches into the following classes:

- **Euclidean Input:** In the most straightforward realization of this approach one requires a minimal set of 5 control points whose positions in the desired Euclidean space are known. This approach is obvious due to the

fact that a projective basis is defined by 5 points and that the projective group includes the Euclidean group as a subgroup. A more advanced variation on this theme is to recover the projective to Euclidean (“Proj-2-Euc”) transformation from Euclidean cues, such as from known angles between lines (such as orthogonal lines) and known distances between points [4].

- **Fixed Internal Parameters Assumption:** Here the goal is to recover the internal parameters of the camera explicitly by making the assumption of a single moving camera whose internal parameters are unknown but remain fixed throughout. These methods are based on recovering invariants to Euclidean transformations, notably the absolute conic and absolute quadric [7, 16, 21] (and [1, 14, 9] in the case of restricted camera motion). More recent approaches [12, 24] assume a stereo rig (of two cameras) moving rigidly.
- **Varying Subset of Internal Parameters:** Using the machinery developed for recovering the absolute quadric [21, 11], one can tolerate variation in a subset of the internal parameters as the number of views increases beyond three. For example, by counting arguments one can show that in the absence of skew 8 views are required for solving for the remaining 4 internal parameters per camera. This idea has been so far described fully and implemented successfully for purely varying focal length [17]. Due to the highly non-linear equations involved, the implementation of these approaches are most suited for relatively large sequences of views, as reported in [17].

In this paper, we pose the question of Projective-to-Euclidean relationship differently. Instead of asking how to calibrate (i.e. recover the internal parameters) a rig of cameras, we ask *under what conditions can we restore the*

*Euclidean representation of a rig under a non-rigid motion of the cameras?*

To appreciate the difference between the two questions, consider the following thought experiment. We are given a rig of  $n \geq 2$  cameras placed in room  $\mathcal{A}$ . Assume that we have obtained somehow the Proj-2-Euc transformation, either through direct input of Euclidean information coming from the features of the room, or through one of the self-calibration methods (by moving the rig rigidly while assuming that the internal and external parameters remain fixed). As long as the rig moves rigidly in space we can obtain an Euclidean representation of space because the projective representation is fixed (or can be made fixed), thus the Proj-2-Euc transformation remains valid. Also, if we apply a non-rigid motion, i.e., change the relative positioning (and internal parameters) among the cameras, but remain in the same room, then by using the overlap between the features of room  $\mathcal{A}$  seen before and after the non-rigid change one can chain the projective-to-projective transformations together (since they form a group) to obtain back the Euclidean representation.

However, consider placing the rig in a different room  $\mathcal{B}$ , such that there is no overlap between the features in room  $\mathcal{A}$  and the features in room  $\mathcal{B}$ . Once the rig is in room  $\mathcal{B}$ , apply a non-rigid change to the rig and take a single snapshot of the room. Can we recover an Euclidean representation of room  $\mathcal{B}$  using only projective calculations, i.e., without using any Euclidean input from room  $\mathcal{B}$ ? Note that existing approaches for self-calibration are not designed to handle such a situation. In room  $\mathcal{B}$  we do not have the freedom to take multiple snapshots while applying a rigid motion to the rig, nor can we assume that the internal parameters remain fixed (because we have a multiple camera situation), and since there is no overlap with the previous room  $\mathcal{A}$  we cannot chain together projective transformations. Yet, the question we have posed is less ambitious than the self-calibration paradigm because we do allow for the possibility of using Euclidean input from room  $\mathcal{A}$  in order to calibrate once (and for all) the Proj-2-Euc transformation. The question of restoring the Proj-2-Euc transformation from dynamically changing projective representations (due to non-rigid changes of the rig) has an interesting practical source. A constructive answer to the question above provides means for designing a rig of cameras that can change its internal configuration, such as focus, zoom, during mensuration. The only requirement is that the rig be calibrated once (i.e., by obtaining the Proj-2-Euc transformation), then during mensuration the field of view and the working distance can change considerably without the need for Euclidean input or overlap with image input taken prior to the change of camera configuration.

Our working assumption is that from the application point of view it is relatively manageable to start with a

known Proj-2-Euc (by simply taking a snapshot of a calibration object prior to the mensuration process). Yet, it is unrealistic to expect for Euclidean input each time the camera configuration has changed, and it may be unrealistic (or very restrictive) to ask for matching features across images before and after the change of rig configuration (for example, change of field of view is large, or the rig is positioned far a way from its previous location, etc.).

In [20], the omni-rig scheme was presented and solved for the case of arbitrary 2D projective transformation per image plane. It was assumed that the internal parameters and camera orientation may change, but the mutual displacement between the centers of projection remains fixed. This assumption is theoretically appealing, but may not hold in most practical situations. Changing the zoom or focus of a camera usually results in a change in the center of projection (COP), and thus changes the mutual displacement between the COPs. Similarly, it is practically very hard to rotate the cameras such that the camera COP is incident with the rotation axis. The system also required 5 cameras with their COPs positioned on a simplex. These requirements pose a practical challenge in terms of occlusions and common field of view.

This work presents a different, more practical approach to the omni-rig scheme. It is assumed that relative orientation between the cameras remains fixed, while the internal parameters and the displacements between the COPs may change freely.

Theoretically it is shown that knowing the rotations between a set of  $n > 2$  cameras is sufficient for upgrading a projective reconstruction to an Euclidean one.

### 1.1 Formal Statement of the Problem

A pinhole camera projects a point  $P$  in 3-D projective space  $\mathcal{P}^3$  to a point  $p$  in the 2-D projective plane  $\mathcal{P}^2$ . The projection can be written as a 3x4 homogeneous matrix  $M$ :

$$p \cong MP$$

where  $\cong$  marks equality up to a scale factor. When the camera is calibrated, it can be factored (by QR decomposition):

$$M_E = K[R; T]$$

where  $R$  and  $T$  are the rotation and translation of the camera respectively, and  $K$  is a 3x3 upper diagonal matrix containing the internal parameters of the camera. The most general form of the internal parameters matrix  $K$  is:

$$K = \begin{pmatrix} f & \gamma & u_0 \\ 0 & \alpha f & v_0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1)$$

where  $f$  is the focal length,  $\alpha$  is the aspect ratio,  $(u_0, v_0)$  is the principle point and  $\gamma$  is the skew. It is practical to model

K by a reduced set of internal parameters, for example assume zero skew.

Generally, given projections of  $m$  3-D points  $\{P_j\}_{j=1}^m$  to  $n$  images, it is possible to estimate the location of the 3-D points and the camera matrices  $\{M\}_{i=1}^n$  up to a projective transformation (collineation) represented by a 4x4 matrix H:

$$p \cong MHH^{-1}P \quad (2)$$

When the internal parameters of the cameras  $\{K_i\}_{i=1}^n$  are known, then H can be recovered up to a 3-D similarity transformation. The goal of (internal) calibration is to recover  $\{K_i\}_{i=1}^n$ , or equivalently recover the 4x4 Collineation  $H$  up to a similarity transformation.

This work follows the omni-rig two rooms calibration scheme: A camera rig is calibrated in room  $\mathcal{A}$ , and then placed in room  $\mathcal{B}$ . During the transfer to room  $\mathcal{B}$ , the cameras may change their zoom/focus. Note that a change in zoom results in a change of the internal parameters (focal length and principle point) as well as a change in the position of the camera center of projection (the translational component of camera position). The problem we address is how to re-calibrate the rig in room  $\mathcal{B}$  using only projective measurements.

It is assumed in this work, that relative rotations between the cameras do not change in the transition from room  $\mathcal{A}$  to room  $\mathcal{B}$ . This assumption is verified by experiments described below.

## 2. Recovery of the Projective-to-Euclidean Collineation in Room $\mathcal{B}$

Given point correspondences between the different views in room  $\mathcal{B}$ , we recover the projective structure of the scene, and the camera matrices up to a projective transformation. This can be done in various ways [10].

Let  $M_i, i = 1..n$  be the projective camera matrices in room  $\mathcal{B}$ , computed from point correspondences. The projective coordinate system can be chosen such that

$$M_1 = [I_{3 \times 3}, 0_3]$$

The Projective-to-Euclidean Collineation  $H$  satisfies, for each  $i = 1..n$ :

$$M_i H \cong K_i [R_i; T_i] \quad (3)$$

where  $R_i, T_i, K_i$  are the rotation, translation and internal parameters of the  $i$ -th camera in some Euclidean coordinate system. The Euclidean coordinate system can be chosen such that  $T_1 = 0$ . Let  $\hat{H}$  be the  $4 \times 3$  matrix composed from the first 3 columns of H. Following equation 3, for each  $i = 1..n$ :

$$M_i \hat{H} \cong K_i R_i \quad (4)$$

so there exist  $\lambda_i, i = 1..n$  such that:

$$M_i \hat{H} R_i^T = \lambda_i K_i \quad (5)$$

$M_i$  are computed from the images, and  $R_i$  is assumed to be known from room  $\mathcal{A}$ . Each known entry of  $K_i, i=1..n$ , contributes one homogeneous equation in the twelve entries of  $\hat{H}$  and the  $n$  scale factors  $\lambda_i$ .

The assumptions about K determine the minimal number of cameras required to solve for the calibration. By counting equations and variables, the number of cameras  $n$  should satisfy:  $n + 12 \leq h * n + 1$  or  $n \geq \frac{11}{h-1}$  Where  $h$  is the number of known entries in  $K_i, i = 1..n$ .

For example:

- **Full calibration:**  $K_i$  contains 4 known entries (3 zeros + one), so the minimal numbers of cameras is 4.
- **Known skew:**  $K_i$  contain 5 known entries (3 zeros + skew + one), so the minimal numbers of cameras is 3.
- **Known principal point, skew:**  $K_i$  contain 7 known entries (3 zeros + 3 known parameters + one), so the minimal numbers of cameras is 2.

Once  $\hat{H}$  is solved, the internal parameters can be recovered by equation 5.

Note that only the first three columns of  $H$  were recovered. By choosing  $T_1 = 0$ , and  $M_1 = [I_{3 \times 3}, 0_3]$ , the last column of  $H$  is  $(0 \ 0 \ 0 \ \delta)^T$ . The unknown parameter  $\delta$  determines the global scale of the reconstructed scene. The reason for the scale ambiguity is that for every  $\delta \neq 0$  define

$$\hat{H}_\delta = H \begin{pmatrix} I_{3 \times 3} & 0 \\ 0 & \delta \end{pmatrix}$$

If equation 3 holds, then for every  $i = 1..n$ ,

$$M_i \hat{H}_\delta \cong K_i [R_i; \delta T_i] \quad (6)$$

Thus the common scale of the camera translations  $T_i, i = 1..n$ , and the scene cannot be recovered by knowing the internal parameters of the cameras.

## 3. Geometrical Analysis: Calibration by invariants.

In this section a geometrical interpretation is given to the method described in the previous section. This interpretation uses points on the plane at infinity  $\pi_\infty$  whose projection to the images are invariant to the changes in translation and internal parameters.

**Theorem 1** *Let*

$$\begin{pmatrix} f & \gamma & u_0 \\ 0 & \alpha f & v_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r_1^T & T_1 \\ r_2^T & T_2 \\ r_3^T & T_3 \end{pmatrix}$$

be a calibrated camera matrix, where  $r_k^T$ ,  $k=1..3$ , are the rows of the rotation matrix  $R$ . Let  $Q$  be the point  $\begin{pmatrix} r_1 \\ 0 \end{pmatrix}$ , and let  $S$  be the point  $\begin{pmatrix} r_2 \\ 0 \end{pmatrix}$ . Then the projection of  $Q$  to the image plane of this camera is  $(1 \ 0 \ 0)^T$ , and the projection of the line passing through  $S$  and  $Q$  is the line of infinity  $(0 \ 0 \ 1)^T$ .

*Proof:*

$$\begin{pmatrix} f & \gamma & u_0 \\ 0 & \alpha f & v_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r_1^T & T_1 \\ r_2^T & T_2 \\ r_3^T & T_3 \end{pmatrix} \begin{pmatrix} r_1 \\ 0 \end{pmatrix} =$$

$$\begin{pmatrix} f & \gamma & u_0 \\ 0 & \alpha f & v_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \cong \begin{pmatrix} f \\ 0 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} f & \gamma & u_0 \\ 0 & \alpha f & v_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r_1^T & T_1 \\ r_2^T & T_2 \\ r_3^T & T_3 \end{pmatrix} \begin{pmatrix} r_2 \\ 0 \end{pmatrix} =$$

$$\begin{pmatrix} f & \gamma & u_0 \\ 0 & \alpha f & v_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \cong \begin{pmatrix} \gamma \\ \alpha f \\ 0 \end{pmatrix}$$

and finally note that  $(f, 0, 0)^T \times (\gamma, \alpha f, 0)^T \cong (0, 0, 1)^T$ . ■

The points  $Q$  and  $S$  are on  $\pi_\infty$  and are determined by the rotation of the matrix, and their projection to the image plane does not depend on the camera translation and internal parameters. They are defined per camera. For a set of cameras  $M_i$ ,  $i = 1..n$ , with known rotations, one can define  $Q_i$ ,  $S_i$ .

Let  $\hat{H}$  be the  $4 \times 3$  matrix which maps  $(X \ Y \ Z)^T$  to  $(A \ B \ C \ D)^T$ , where  $(A \ B \ C \ D)^T$  is the point in the projective coordinate system of room  $\mathcal{B}$  corresponding to the point  $(X \ Y \ Z \ 0)^T$  in some Euclidean coordinate system of room  $\mathcal{B}$ . If  $H$  is a Euclidean-to-projective transformation of the second room, then  $\hat{H}$  is composed from the first three columns of  $H$ . This definition of  $\hat{H}$  is consistent with the definition of  $\hat{H}$  given in the previous section.

Knowing the invariant 3D points/lines constrains  $\hat{H}$ . For example, the invariant point  $\hat{H}Q_i$  lies on the line of sight of point  $(1 \ 0 \ 0)$  in camera  $M_i$ . This contributes two constraints on  $\hat{H}$ . Similar constraints can be used with the invariant lines. For example, The line passing through  $\hat{H}Q_i$  and  $\hat{H}S_i$  intersects the line of sight of point  $(0 \ 1 \ 0)$  in the  $i$ -th camera. This contributes a constraint on  $\hat{H}$ . Therefore, for a full set of 5 internal parameters per camera, we have 3 constraints per camera and thus need 4 cameras to uniquely define  $\hat{H}$ . In the case of vanishing skew, we have an additional invariant as the point  $S$  is projected onto  $(0 \ 1 \ 0)$ , thus each camera would provide 5 constraints on  $\hat{H}$ .

The Euclidean coordinate systems in room  $\mathcal{A}$  and room  $\mathcal{B}$  can be chosen such that the coordinates of the 3D invariant points/lines in room  $\mathcal{B}$  and room  $\mathcal{A}$  coincide. Thus the coordinates of the 3D invariant points/lines can be measured in room  $\mathcal{A}$ , and then used in room  $\mathcal{B}$ . For example, for every camera  $M_i$ ,  $i = 1..n$ , point  $Q_i$  as defined in theorem 1 is the intersection of the plane at infinity with the line of sight of the point  $(1 \ 0 \ 0)$  in the  $i$ -th camera in room  $\mathcal{A}$ .

Having the recovered  $\hat{H}$ , consider the  $3 \times 3$  matrix  $M_i \hat{H}$ . It transforms a point in the plane at infinity, given in the reduced Euclidean coordinate frame, to its projection on the  $i$ -th image in room  $\mathcal{B}$ . i.e This is the homography between the plane at infinity and the  $i$ -th image.

$$K_i R_i \cong M_i \hat{H}$$

$K_i$  can be recovered by QR decomposition, or, as  $R_i$  is known, by  $K_i \cong M_i \hat{H} R_i^T$ .

## 4. Results

To test the applicability of the proposed method, several experiments were conducted. The cameras used were of a  $720 \times 480$  pixels of resolution, where the combination of field of view and distance to object made a pixel occupy a millimeter square in space.

The first experiment was to test the basic assumption that changes of zoom leaves the relative orientation unchanged yet varies all other parameters (internal parameters and position of center of projection). In each zoom setting the calibration pattern (see Fig. 1) was used for a direct estimation of the Euclidean camera matrices using the known 3D positions of the calibration points (the corners of the checkerboard pattern). Table 4 displays the values of the calibration parameters at four different zoom settings (in the range of  $1 - 2$  of scale factor). Note that the internal parameters and the position of the center of projection vary considerably (where  $T$  varies mostly along the  $Z$  axis), yet the camera orientation (measured in Euler angles  $R_x$ ,  $R_y$ ,  $R_z$ ) remains relatively stable.

The accuracy of the omni-rig solution for a 3-camera rig (assuming vanishing skew) was evaluated as follows. The Euclidean camera matrices were recovered from the calibration pattern and the relative camera orientations were estimated. The effective accuracy of the reconstruction (combining together the uncertainty of point matching and the depth uncertainty due to rig geometry) was on average 1 millimeter (distance of reconstructed points from the control points) and the back projection of the reconstructed points onto the image space yielded an average error of 0.4 pixels.

After capturing the first set of images (room  $\mathcal{A}$ ), the zoom of the cameras was changed freely, and another set

Parameter	Zoom 1	Zoom 2	Zoom 3	Zoom 4
$R_X$	169.376	169.267	168.923	168.975
$R_Y$	145.232	144.721	145.106	145.253
$R_Z$	170.295	170.204	169.930	169.954
$T_X$	-102.025	-96.648	-100.095	-103.335
$T_Y$	-149.359	-147.564	-146.491	-146.152
$T_Z$	-693.530	-688.638	-698.129	-718.160
$f_x$	1546.778	1448.106	1335.391	1177.927
$\gamma$	-2.618	-3.118	-2.332	-1.755
$f_y$	1424.203	1333.413	1228.614	1084.241
$u_0$	376.842	387.107	376.397	369.073
$v_0$	232.628	237.015	240.425	241.727

**Table 1.** Camera parameters values in different zoom/focus states. The seems fixed in comparison with the other parameters.

of images of the calibration object was captured (room  $\mathcal{B}$ ). Next, the corresponding points in the images of “Room  $\mathcal{B}$ ” were used in order to compute the projective camera matrices and the projective 3D coordinates of the points. This was done without using any a priori knowledge of the scene. Finally the Projective-to-Euclidean transformation  $H$ , and the internal and external parameters of the cameras were computed, based on the rotations/invariants from “Room  $\mathcal{A}$ ” and the projective camera matrices. The accuracy of calibration was measured both in terms of distance from the control points and, since the calibration pattern consists of two perpendicular planes, we used also the angle between these planes as a measure of calibration quality.

The results of the omni-rig calibration were compared to three other Euclidean reconstructions:

- Proj2Euc: Computing the optimal least squares transformation mapping the projective reconstruction in room  $\mathcal{B}$  to the known 3D points. The quality of this computation depends on the accuracy of the image measurements, and on the quality of the projective reconstruction algorithm. The omni-rig system estimates this transformation without knowing the 3D points. Thus the omni-rig solution is at most as accurate as this solution.
- room-A: Assuming the camera parameters have not changed in the transitions between the rooms. This computation was expected to yield the worst results
- Known 3D: Direct linear computation of the Euclidean cameras by using the knowledge of the 3D coordinates. This computation was expected to yield the best results because the optimization is performed directly in 3D space without going through the projective reconstruction which provides a less meaningful optimization criteria.

Exp.	Proj2Euc.	Known-3D	Omni-rig	Room-A
1	90.0471	89.9585	89.6611	78.2593
2	90.1609	89.8293	89.0388	77.4602
3	89.9949	89.9886	90.5147	72.8573

**Table 2.** The angles between the planes of the calibration object in the reconstructed points (see text for further details). A typical image of the calibration object is shown in Fig. 1

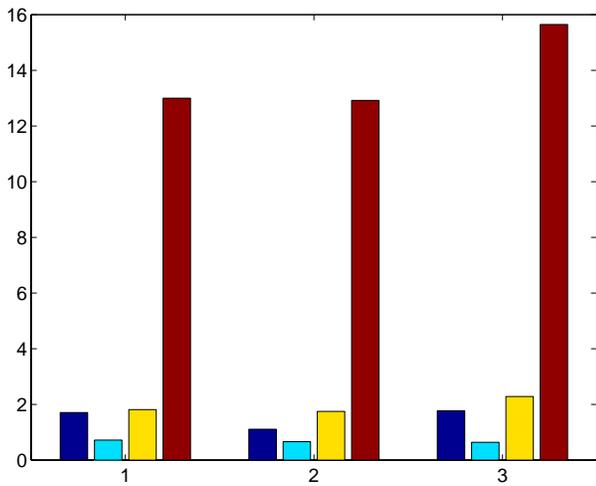


**Figure 1.** The calibration object, composed of two perpendicular planes with a checker-board pattern.

Fig. 2 and table 2 summarize the results for the case of zero-skew (3-camera rig). Fig. 2 describes the mean of distances (in millimeters) between the known 3D points and the reconstructed points. The bar chart presents the results of three experiments, each experiment on different set of images. The high right bar in each experiment is the result of using the original cameras from room  $\mathcal{A}$ . Then, from right to left are the results of the omni-rig, the calibration using known 3D points, and the optimal Projective-to-Euclidean transformation. It is noticeable that the results of the omni-rig are only slightly worse than the optimal Proj-to-Euclidean. This implies that the process (including the assumption that relative orientation remains fixed) did not add a source of error (in addition to the projective reconstruction). Note also that compared to the best one can do in these circumstances (“known-3D”) the omni-rig process is at most twice as worse from optimum, where most of the error is attributed to the projective reconstruction stage and not to the re-calibration stage.

Table 2 describes the angles between the reconstructed planes of the orthogonal calibration object. As expected, the worst results are achieved when using the camera parameters from room  $\mathcal{A}$ , and the best results are achieved when the known 3D points are used. The quality of the omni-rig calibration is only slightly inferior to the Proj2Euc calibration. This proves the accuracy and the applicability of the omni-rig calibration scheme.

Finally, the applicability of omni-rig for visualization applications is demonstrated. Given 3 images, and dense correspondences, the scene was reconstructed up to a



**Figure 2.** The mean of distances (in millimeters) between the reconstructed points and the corresponding known 3D coordinates in 3 experiments. Right to left: Room  $\mathcal{A}$ , omni, known 3D, proj-to-Euc. The omni rig error is almost as small as the optimal least squares Proj-2-Euc error.

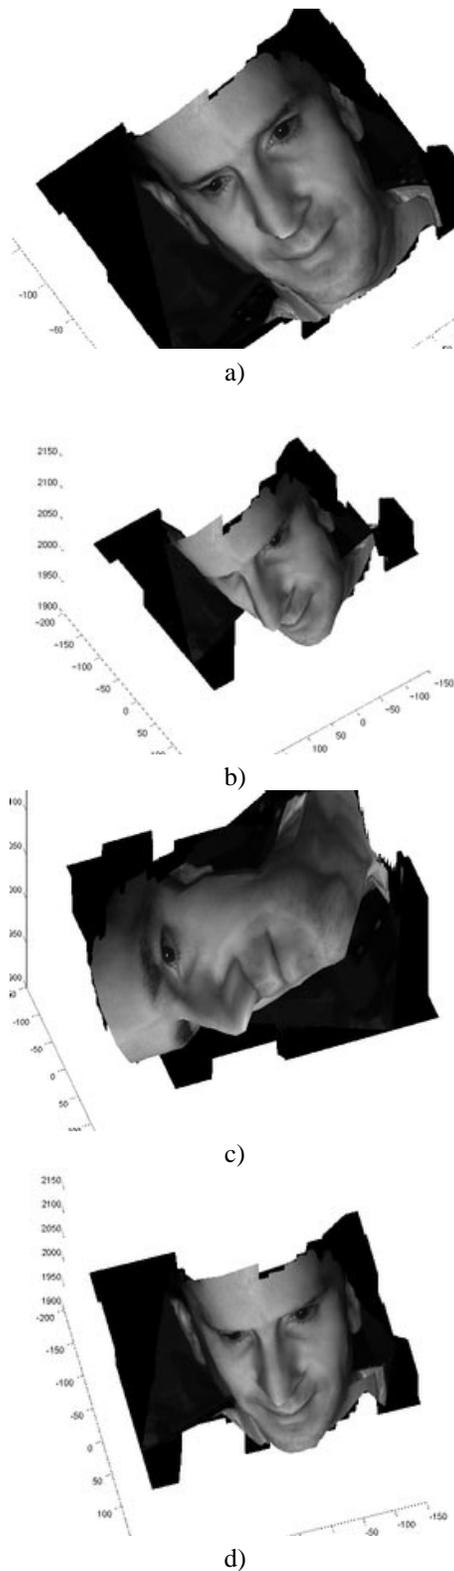
projective reconstruction. Then, using the known rotation, the Projective-to-Euclidean collineation was computed by the omni-rig method. Having the Euclidean coordinate system enables to construct a texture-mapped 3D model, and to rotate it by any angle. Rotating a non-Euclidean reconstruction would result in affine or projective distortion in the images. The results of this experiment are presented in figure 3.

## 5. Summary

We have presented a simple linear method to re-calibrate an  $n$ -camera rig. Our only assumption is that the camera bodies are fixed relative to one another, hence their mutual rotations remain unchanged. We have shown that for a full re-calibration of internal parameters and position of center of projection of each camera, a rig of 4 cameras is required. In the case of vanishing skew a rig of 3 cameras would be sufficient, and in case of vanishing skew and known principle point then a stereo rig would suffice.

The experimental setup verified our basic assumption that changes of zoom does not affect the relative rotational component of camera positions, yet affects all other calibration parameters. We have compared our re-calibration procedure to the optimal solutions (using knowledge of control points) and have found a close agreement in terms of accuracy.

In the future we plan to make the recalibration process more robust, without adding new information in room  $\mathcal{B}$ . One Example of this would be to measure the axis on which



**Figure 3.** A 3D model of a face in various rotations and zooms. The model was reconstructed projectively, and upgraded to Euclidean by the omni-rig method. A correct Euclidean reconstruction enables to rotate the 3D model without introducing projective distortions to the images.

each camera center moves.

Another future direction would be to describe all the parameters of the cameras in the omni rig as a function of a minimal set of parameters. Every camera is allowed to change its focus and zoom, and it seems reasonable to model the entire configuration of the camera as a function of these parameters. This process would require more demanding measurements in room  $\mathcal{A}$ , but might produce a more accurate recalibrated rig.

## Acknowledgments

The authors would like to thank Cognitens LTD. for providing the data for the experiment described in figure 3, and especially to Dr. Tamir Shalom for his help in preliminary experiments.

## References

- [1] M. Armstrong, A. Zisserman, and R.I. Hartley. Self-calibration using image triplets. In *Proceedings of the European Conference on Computer Vision*, LNCS-1064, pages 3–16. Springer, April 1996.
- [2] S. Avidan and A. Shashua. Novel view synthesis by cascading trilinear tensors. *IEEE Transactions on Visualization and Computer Graphics*, 4(3), 1998. Short version can be found in CVPR'97.
- [3] M. Born and E. Wolf. *Principles of Optics*. Permagon Press, 1965.
- [4] B. Boufma, R. Mohr, and F. Veillon. Euclidean constraints for uncalibrated reconstruction. In *Proceedings of the International Conference on Computer Vision*, pages 466–470, Berlin, Germany, May 1993.
- [5] S. Carlsson. Duality of reconstruction and positioning from projective views. In *Proceedings of the workshop on Scene Representations*, Cambridge, MA., June 1995.
- [6] O.D. Faugeras. Stratification of three-dimensional vision: projective, affine and metric representations. *Journal of the Optical Society of America*, 12(3):465–484, 1995.
- [7] O.D. Faugeras, Q.T. Luong, and S.J. Maybank. Camera self calibration: Theory and experiments. In *Proceedings of the European Conference on Computer Vision*, pages 321–334, Santa Margherita Ligure, Italy, June 1992.
- [8] F.R. Gantmacher. *The theory of matrices*, volume I. Chelsea publishing, NY, 1990.
- [9] R. Hartley. Self calibration from multiple views with a rotating camera. In *Proceedings of the European Conference on Computer Vision*, pages 471–478, Stockholm, Sweden, May 1994.
- [10] R.I. Hartley and A. Zisserman. *Multiple View Geometry*. Cambridge University Press, 2000.
- [11] A. Heyden and K. Astrom. Euclidean reconstruction from constant intrinsic parameters. In *Proceedings of the International Conference on Pattern Recognition*, pages 339–343, Vienna, 1996.
- [12] R. Horaud and G. Csurka. Self-calibration and Euclidean reconstruction using motions of a stereo rig. In *Proceedings of the International Conference on Computer Vision*, pages 96–103, Bombay, India, January 1998.
- [13] S. Leveau and O.D. Faugeras. 3-D scene representation as a collection of images and fundamental matrices. Technical report, INRIA, Feb. 1994.
- [14] T. Moons, L. Van Gool, M. Van Diest, and E. Pauwels. Affine reconstruction from perspective image pairs. In *The 2nd European Workshop on Invariants*, Ponta Delagada, Azores, October 1993.
- [15] S.K. Nayar and S. Baker. Catadioptric image formation. In *Proceedings of the DARPA Image Understanding Workshop*, pages 1431–1437, New Orleans, LA, May 1997.
- [16] M. Pollefeys and L. Van Gool. A stratified approach to metric self-calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 407–412, Puerto Rico, June 1997. IEEE Computer Society Press.
- [17] M. Pollefeys, R. Koch, and L. Van Gool. Self-calibration and metric reconstruction in spite of varying and unknown camera parameters. In *Proceedings of the International Conference on Computer Vision*, Bombay, India, January 1998.
- [18] A. Shashua. Algebraic functions for recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):779–789, 1995.
- [19] A. Shashua. Trilinear tensor: The fundamental construct of multiple-view geometry and its applications. In G. Sommer and J.J. Koenderink, editors, *Algebraic Frames For The Perception Action Cycle*, number 1315 in Lecture Notes in Computer Science. Springer, 1997. Proceedings of the workshop held in Kiel, Germany, Sep. 1997.
- [20] A. Shashua. Omni-Rig Sensors: What Can be Done With a Non-Rigid Vision Platform?. In *Proc. of the Workshop on Applications of Computer Vision (WACV)*, Princeton, Oct. 1998.
- [21] B. Triggs. Autocalibration and the absolute quadric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 609–614, Puerto Rico, June 1997. IEEE Computer Society Press.
- [22] M. Watanabe and S.K. Nayar. Telecentric optics for computational vision. In *Proceedings of the DARPA Image Understanding Workshop*, pages 781–785, New Orleans, LA, May 1997.
- [23] D. Weinshall, M. Werman, and A. Shashua. Shape tensors for efficient and learnable indexing. In *Proceedings of the workshop on Scene Representations*, Cambridge, MA., June 1995.
- [24] A. Zisserman, P.A. Beardsley, and I.D. Reid. Metric calibration of a stereo rig. In *Proc. of IEEE Workshop on Representation of Visual Scenes*, pages 93–100, Cambridge, MA, June 1995.