Sequence-to-Sequence Self Calibration

Lior Wolf and Assaf Zomet

School of Computer Science and Engineering,
The Hebrew University,
Jerusalem 91904, Israel
{lwolf,zomet}@cs.huji.ac.il

Abstract. We present a linear method for self-calibration of a moving rig when no correspondences are available between the cameras. Such a scenario occurs, for example, when the cameras have different viewing angles, different zoom factors or different spectral ranges. It is assumed that during the motion of the rig, the relative viewing angle between the cameras remains fixed and is known. Except for the fixed relative viewing angle, any of the internal parameters and any of the other external parameters of the cameras may vary freely. The calibration is done by linearly computing multilinear invariants, expressing the relations between the optical axes of the cameras during the motion. A solution is then extracted from these invariants. Given the affine calibration, the metric calibration is known to be achieved linearly (e.g. by assuming zero skew). Thus an automatic solution is presented for self calibration of a class of moving rigs with varying internal parameters. This solution is achieved without using any correspondences between the cameras, and requires only solving linear equations.
Keywords: Self-Calibration, Multi-View Invariants.

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.

For many tasks such as computer graphics, projective representation is not sufficient, and an Affine/Metric representation is required. In recent years there has been much progress in the theory and algorithms of Self/Auto-Calibration, the upgrading of a projective structure to an Affine/Metric one, without using knowledge of the viewed scene. These algorithms must make assumptions about the parameters of the camera to achieve a unique solution (For a review see [6]).

There are roughly two approaches to self-calibration: The non-stratified approach solves directly for metric upgrade. This can be done for example by solving Kruppa's equations [4], or by solving for the absolute quadric [11]. In both cases the solution of a set of non-linear equations is required. The stratified approach first upgrades the projective representation to an affine representation by solving for the homographies through the plane at infinity and then upgrades

A. Heyden et al. (Eds.): ECCV 2002, LNCS 2351, pp. 370–382, 2002.

the affine representation to a metric one [3,8]. It was pointed by several authors (e.g. [6,7]) that the first stage, the affine calibration, is the most challenging in stratified methods.

By using a stereo rig, a stable solution for both affine and metric calibration can be computed linearly from two or more images from each camera ([7], following [12]). This method assumes that the rig structure and the Zoom/Focus are fixed. Another approach for rig calibration achieves a linear solution by assuming known relative orientations of the cameras, but allowing the Zoom/Focus and the relative displacements to vary [13]. Both methods require correspondences between the cameras.

In this paper we focus on a self-calibration scenario which is of interest on both practical and theoretical fronts. We consider a rig of two or more video cameras, each capturing an image sequence, where there is no or little spatial overlap between the fields of view of the cameras (such as when the cameras have different zoom settings or pointing to largely different directions). While in motion, each camera may change its internal parameters and zoom factor.

Applying conventional self calibration algorithms for each camera separately is always challenging with varying internal parameters, especially the affine calibration stage. Applying one of the existing rig algorithms [7,13], on the other hand, is not possible since no correspondences are available between the cameras.

Therefore, we seek to solve for the calibration of the cameras, while exploiting the *rigidity constraints between them.* We refer to this problem as "Sequence to Sequence Self-Calibration", since every sequence requires a different calibration, and only the relations between the sequences can be exploited to recover these calibrations. A similar scenario, of exploiting the rigidity between multiple cameras, was presented recently in the context of image alignment [1].

A previous approach for handling self calibration without correspondences between the cameras [2] assumed highly constrained conditions: In addition to knowing one degree of freedom about the rotation between the cameras, that solution also assumed fixed internal parameters, where only two of which were unknown. It also required the solution of highly non-linear equations over a large set of variables. The work presented in this paper solves for affine calibration linearly, without using any assumption but one known d.o.f. of the relative rotation between the cameras.

In order to solve for the calibration, we use assumptions about the relative orientations of the cameras in the rig. We divide the problem to three different cases:

– Two cameras with parallel image planes. The cameras may view the same direction, or the opposite directions (Back to back).
– Two cameras with non-parallel image planes, with the special case of orthogonal image planes.
– Cameras with varying orientations, each rotating about its Y axis. Their Y axes are assumed to be parallel.

In the following sections we present a linear solution for affine self-calibration for each of the cases above, assuming the cameras are synchronized. Later in

Section. 5 we show how such synchronization can be achieved automatically. The calibration is done by computing multilinear tensorial invariants, expressing the relations between the axes of the cameras during the motion. A solution is then extracted from these tensors. Having solved for the challenging stage of affine calibration, the metric calibration can be achieved linearly by imposing further constraints, e.g. zero skew (see [6]). In case the cameras view the same/opposite directions, they share the same affine ambiguity. Thus metric constraints on both cameras can be simultaneously used for calibrating both of the cameras.

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 3×4 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):

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

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

$$K = \begin{bmatrix} f & \gamma & u_0 \\ 0 & \alpha f & v_0 \\ 0 & 0 & 1 \end{bmatrix} \tag{1}$$

where f is the *focal length*, α is the *aspect ratio*, (u_0, v_0) is the *principle point* and γ 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 4×4 matrix H:

$$p \cong MH^{-1}HP \tag{2}$$

The matrices $\{C_i = M_i H^{-1}\}_{i=1}^n$ are called the projective camera matrices. The points HP are the points in the projective coordinate system. We refer to the matrix H^{-1} as the projective to Euclidean matrix.

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 4×4 collineation H up to a similarity transformation.

In all following sections it is assumed that the projective camera matrices $\{C\}_{i=1}^n$ were already computed for each one of the sequences using for example point correspondences between the different frames. This can be done in various ways (for a review see [6]).

2 Parallel Optical Axes

We first analyze the case of two cameras attached in a rig, viewing the same direction, or the opposite directions (back-to-back). In addition, each camera may rotate freely about its Z axis. Note that there are no constraining assumption about neither the internal parameters of the cameras, nor the relative displacements of the centers of projections.

The first Euclidean camera matrix is given by:

$$M_i^1 = K_i^1[R_i; t_i^1]$$

and the second Euclidean camera is given by:

$$M_i^2 = K_i^2[R_0 * R_i; t_i^2]$$

Where R_0 is the relative rotation between the cameras, and the third row of R_0 equals to $[0, 0, 1]$ up to sign.

Given the projective cameras C_i^1, C_i^2, we seek for two projective transformations H_1, H_2 mapping the projective cameras to the Euclidean ones (up to similarity transformation).

Let r_i be the third row of R_i, and let c_i^j be the third row of C_i^j. Let $\hat{H}_j, j = 1..2$ be the 4×3 matrices composed from the first 3 columns of the matrices H_j.

The internal parameters matrices K_i^j are upper triangular matrices and so $\hat{H}_1^T c_i^1 \cong r_i$. Having also the special structure of the third row of $R_0 : \hat{H}_2^T c_i^2 \cong r_i$. Thus for every i:

$$\hat{H}_1^T c_i^1 \cong \hat{H}_2^T c_i^2 \cong r_i \tag{3}$$

The constraint above holds also for $D\hat{H}_1, D\hat{H}_2$ for every 3×3 matrix D. This results in an ambiguity that would later on express itself as an affine ambiguity on the cameras, the same affine ambiguity for both cameras (up to similarity).

The above Eqns. 3 have a form we call "equivalence after projection". In the next section we present the solution for such equations.

Given the solution to these equations, the homographies through the plane at infinity can be recovered for each sequence and between the sequences. A point $[a^T, 0]^T$ on the plane at infinity is projected to $C_i^j \hat{H}_j a$ in the ith frame of the jth camera. Hence the homography at infinity H_{ijkl}^∞ between frame i of camera j and frame k of camera l is given by

$$H_{ijkl}^\infty = C_k^l \hat{H}_l (C_i^j \hat{H}_j)^{-1}$$

The solutions to Eqn. 3 are up to a common 3×3 matrix D which cancels out, and so:

$$H_{ijkl}^\infty = C_k^l \hat{H}_l (C_i^j \hat{H}_j)^{-1} = C_k^l \hat{H}_l D (C_i^j \hat{H}_j D)^{-1}$$

Similarly, the coordinates of the plane at infinity L_j in the projective representations of camera j can be recovered. Since L_j satisfies $L_j^\top \hat{H}_j a = 0$ for every a, then L_j is the null space of $\hat{H}_j{}^\top$, which is also the null space of $D\hat{H}_j{}^\top$.

2.1 The "Equivalence after Projection" Problem

We define the equivalence after projection problem as follows: Given two sets of points in P^3 $\{P_i\}_{i=1}^n$,$\{Q_i\}_{i=1}^n$, determine whether there exist two 3×4 projection matrices A, B such that for every i:

$$AP_i \cong BQ_i$$

If such A and B exist solve for them up to a multiplication on the left by a 3×3 matrix. We next show a solution for this problem which requires only solving linear systems of equations.

Let C_B be the center of the projection matrix B, i.e C_B is the null space of B. Let B^+ be a pseudo-inverse of B, i.e. a mapping from the image plane of B to a plane in $\mathcal{P}^3$ such that $BB^+p \cong p$.

The points C_B, B^+AP_i are incident with the line of sight of the projection matrix B associated with the image point AP_i. Since $AP_i \cong BQ_i$, Q_i is also on this line. This defines a constraint on the projection of these points by an arbitrary projection matrix O: Let $e_O = OC_B$. Then the point OQ_i resides on the line $e_O \times OB^+AP_i$. Therefore:

$$(OQ_i)^T([e_O]_x OB^+A)P_i = 0 \quad (4)$$

By choosing an arbitrary camera matrix O, the 3×4 matrix $F_O = ([e_O]_x OB^+A)$ can be recovered linearly from the pairs of points $\{OQ_i, P_i\}_{i=1}^n$. The introduction of O enables to extract a unique bi-linear invariant with a minimal number of parameters. There exists a set of invariants described by 4×4 matrices G, satisfying: $(Q_i)^T(G)P_i = 0$. These matrices span a linear space of dimension 4. By using the matrix O, we get a unique invariant with 12 elements.

For each O, the image e_O of C_B can be recovered as the null space of F_O^T. Using two such projection matrices O, the 3D position of C_B can be recovered by triangulation.

A, B can be recovered up to some 3×3 matrix D. All cameras of the form DB share the same center of projection. Thus B can chosen as any 3×4 matrix whose left null space is the recovered C_B. In order to solve for A, $\{BQ_i\}_{i=1}^n$ are first computed. Since for all i, $AP_i \cong BQ_i$ and BQ_i, P_i are known, A can be recovered. This can be done by any method for recovering a camera from the projections of known 3D points [6].

3 Non-parallel Optical Axes

In the previous section, we studied the case of two cameras viewing the same or opposite directions. In this section, it is assumed that the two cameras have a known constant angle between their optical axes.

The directions of the optical axes of the cameras are given by $c_i^{jT}\hat{H}_j$, and the angle between the axes of the two cameras α satisfies:

$$\cos^2\alpha = \frac{(c_i^{1T}\hat{H}_1\hat{H}_2^T c_i^2)^2}{(c_i^{1T}\hat{H}_1\hat{H}_1^T c_i^1)(c_i^{2T}\hat{H}_2\hat{H}_2^T c_i^2)} \quad (5)$$

This defines a nonlinear metric constraint on $\hat{H}_1, \hat{H}_2$.

We solve for the case of orthogonal optical axes, where $\cos^2 \alpha = 0$, i.e. by Eqn. 5:

$$c_i^{1T} \hat{H}_1 \hat{H}_2^T c_i^2 = 0$$

Each image pair provides one constraint on the matrix $F = \hat{H}_1 \hat{H}_2^T$. From the matrix F, one can extract using SVD the matrices $\hat{H}_1 D$ and $\hat{H}_2 D^{-\top}$ for some unknown 3×3 matrix D.

The planes at infinity of the two projective reconstructions are the null spaces of $\hat{H}_1^\top, \hat{H}_2^\top$, as in the previous section. The homography through the plane at infinity between frames i, k of the same sequence j can be computed similarly to previous case by $C_k^j \hat{H}_j (C_i^j \hat{H}_j)^{-1}$. Since the two matrices $\hat{H}_1, \hat{H}_2$ do not have the same ambiguity, the homography through the plane at infinity between frames of different sequences cannot be computed.

4 Non-fixed Parallel Rotation Axes

In previous sections we analyzed self calibration for rigs in which the angle between the optical axes of the cameras remains fixed. In some cases it is useful to enable the cameras to rotate. For example, they may need to be reoriented such that the object of interest appears in the image.

In this section we explore another type of constraints for self calibration of rigs without correspondences between the cameras. It is assumed that the cameras in the rig may rotate arbitrarily about their Y axis (or similarly about its X axis). It is further assumed that their Y axes are parallel. The rig as a whole can rotate and translate freely in space, and the internal parameters of the cameras may vary freely. We describe two cases. In the first case three cameras are used and the internal cameras are not constrained. In the second case two cameras are used, but we assume their skew is 0.

4.1 Three Cameras

Let σ_i^j be plane X-Y of camera j in time instance i. Since the three cameras rotate about their Y axes, and the Y axes are parallel, then in every time instance i, all planes $\{\sigma_i^j\}_{j=1,2,3}$ intersect in a point on the plane at infinity. We next show how we express this constraint as a multilinear equation in the projective camera matrices.

Let H_j^{-1} be the transformation mapping the projective coordinate system of camera j to a common affine coordinate system in which the plane at infinity is given by $L = [0 \quad 0 \quad 0 \quad 1]$. Let $\{\pi^j\}_{j=1,2,3}$ be three planes, where π^j is given in the coordinate system of the j-th camera. Then if these planes intersect in a point at infinity, the determinant of the following matrix vanishes:

$$\begin{bmatrix} (\pi^1)^\top H_1 \\ (\pi^2)^\top H_2 \\ (\pi^3)^\top H_3 \\ L \end{bmatrix} \tag{6}$$

This determinant can be written using a $4 \times 4 \times 4$ tensor J^{abc}, as

$$\pi_a^1 \pi_b^2 \pi_c^3 J^{abc} = 0$$

The reader is assumed to be familiar with tensor notations. See for example in [9]

The tensor J^{abc} can now be used to express the calibration constraints. In each time i all planes $\{\sigma_i^j\}_{j=1,2,3}$ intersect in a point on the plane at infinity. As σ_i^j is given by $[0,0,1]C_i^j$, this is expressed by:

$$([0,0,1]C_i^1)_a([0,0,1]C_i^2)_b([0,0,1]C_i^3)_c J^{abc} = 0$$

Given the projective camera matrices C_i^j, every time instance provides a linear constraint on J^{abc}. Thus J^{abc} can be computed linearly from the projection matrices.

In order to extract the homographies through the plane at infinity between different images, We identify points at infinity in each coordinate frame. Let Y_i^3 be a double contraction of the form $Y_i^3 = ([0,0,1]C_i^1)_a([0,0,1]^\top C_i^2)_b J^{abc}$. Let N be any plane in the third coordinate system intersecting σ_i^1, σ_i^2 in a point on the plane at infinity. Then:

$$([0,0,1]C_i^1)_a([0,0,1]^\top C_i^2)_b N_c J^{abc} = N^\top Y_i^3 = 0$$

Hence Y_i^3 is the point of intersection of $\{\sigma_i^j\}_{j=1,2,3}$ in the coordinate frame of the third camera. Similarly this point can be extracted in the two other coordinate systems: Y_i^1, Y_i^2.

The set of points $\{Y_i^j\}_{i=1..n}$ are sufficient to determine the plane at infinity at the jth coordinate system. We can use these matching points on the plane at infinity to compute the homography at infinity between all frames: The homography at infinity H_{stuv}^∞ between frames s and u of cameras t and v can be computed using the pairs of matching points $\{(C_s^t Y_i^t, C_u^v Y_i^v)\}_{i=1..n}$.

4.2 Two Cameras with Zero Skew

In this section we show that when the cameras have zero skew, two cameras in the above settings are sufficient for linear recovery of the affine calibration. Then a Metric upgrade can be achieved linearly using standard methods [6].

For every i, j the projective camera matrix C_i^j satisfies:

$$C_i^j \hat{H}_j \cong K_i^j R_j^i R_i$$

where R_j^i is the rotation of the j-th camera with respect to the first camera in time instance i, and R_i is the rotation of the rig in the same time. Again we base the derivation on the structure of a rotation matrix R_j^i about the Y axis. Since the first and third rows of R_j^i have a vanishing second coordinate, then so is their linear combination. Assuming zero skew, the first and last rows of $C_i^j \hat{H}_j$ are linear combinations of the first and third rows of R_i

Let a_i^j be the first row of the projective camera matrix C_i^j. Then the following determinant vanishes for every i:

$$\begin{bmatrix} (a_i^1)^\top \hat{H}_1 \\ (c_i^1)^\top \hat{H}_1 \\ (c_i^2)^\top \hat{H}_2 \end{bmatrix}$$

This constraint can be expressed by a trilinear tensor:

$$([1,0,0]C_i^1)_a([0,0,1]C_i^1)_b([0,0,1]C_i^2)_c K^{abc} = 0$$

The tensor K^{abc} can be solved linearly from the projection matrices. As in the previous section, this constraint can be interpreted geometrically as the intersection of three planes with the plane at infinity. Hence the plane at infinity of the coordinate systems of the two cameras can be extracted by a method similar to the one in the previous section.

Note that similar tensors can be derived by choosing other three-combinations of $\hat{H}_1^\top a_i^1, H_1^\top c_i^1, hatH_2^\top a_i, hatH_2^\top c_i$

5 Sequences Synchronization

In the previous sections it was assumed that the sequences are synchronized. This enabled to compute the multilinear constraints: F_O of section 2.1, F of section 3 and J^{abc} of Section 4. However the existence of these constraints may be used to establish the synchronization between the sequences.

Consider for example the case of the orthogonal optical axes presented in section 3. The existence of a matrix F such that for every frame i: $c_i^{1\top} F c_i^2 = 0$ is not guaranteed if the two sequences are not temporally aligned. Let A be the estimation matrix of F, i.e A is the matrix whose ith row is composed from the Kronecker product of c_i^1 and c_i^2. If such an F exists then the rank of the matrix A is not more than 8 and the vector composed from the elements of F lies in the null space of A. In practice due to noise A is always of full rank, and we use a least squares solution, choosing F to be composed out of the elements of the eigenvector of $AA^\top$ with the smallest eigenvalue. Let f be this eigenvector. We define the magnitude of Af as the algebraic error of the estimation of F. When the sequences are synchronized this magnitude is expected to be small.

In order to synchronize the sequences this algebraic error is measured for each temporal shift between the sequences. The shift which produces the minimal algebraic error is chosen as the solution. This measure is not optimal for a number of reasons. First it has no real geometrical meaning. Second it depends on the number of the frames which can bias in short sequences toward the ends of the sequence. Third, changing the coordinates of c_i^j changes the measure. However in practice we find minimizing the algebraic error to work well on our sequences.

6 Experiments

We have conducted several experiments for testing the quality of the affine self-calibration and the sequence synchronization. In all our experiments we have used an object modeling framework: The cameras were static, and an object to be modeled was rotating in the scene. Points on this object were tracked, and the virtual camera motion with respect to the object was estimated. Since no calibration was available, the cameras and points had a projective ambiguity to be recovered in the experiments. Then the proposed algorithm was used, estimating an Affine representation of the structure and motion.

In the first experiment we tested the sequence synchronization application. The algorithm was applied for several temporal shifts, and the sequences were temporally aligned by finding the shift with the minimal residual error, as described in Section. 5. Figures 1-a,b present the results of this experiment for opposite and orthogonal directions respectively. The true shift was estimated using hand waving in front of the camera. Note that this is an integer estimation for the non-integer temporal shift of the sequences, and thus it is accurate up to 0.5 a frame. In the experiments we tested 10 frames shifts in each temporal direction on sequences containing 40 frames. Indeed a shift close to the estimated by no more than 1 frame yielded the minimal estimation error.

One way to verify the accuracy of affine calibration is to test the accuracy of the homographies through the plane at infinity. In the second experiment we tested the accuracy of the homography through the plane at infinity by mapping vanishing points between the images. We have marked points lying on parallel lines, and tracked them along the sequences. The vanishing points at each frame were computed as the intersections of the parallel lines defined by the tracked points. We then mapped the vanishing points from the input images to a common coordinate system. The quality of the homographies was measured by the proximity of the mapped vanishing points. We have conducted this test for two scenarios: One for cameras viewing opposite directions (Section. 2), and one for cameras viewing orthogonal directions (Section. 3). Figure 2 shows the results of these experiments. Errors in these results are combined from errors in the Affine calibration, errors in the projective camera matrices (the input to our algorithm), and errors in the estimations of the vanishing points in the images due to drifts in the points tracking. However it is visible that the algorithm does indeed align the vanishing points.

Finally, the quality of the homographies can also be tested visually. Warping a sequence of images of a moving camera to a common reference by the homographies through the plane at infinity cancels the rotations of the cameras, as well as the variations in the internal parameters. The result is a sequence in which the camera virtually moves in pure translation and constant internal parameters. Figure. 3 shows the results of applying the computed homographies through the plane at infinity to several images. Note that while the original motion included a rotation, the homographies through the plane at infinity canceled the rotations, leaving only a translational component.

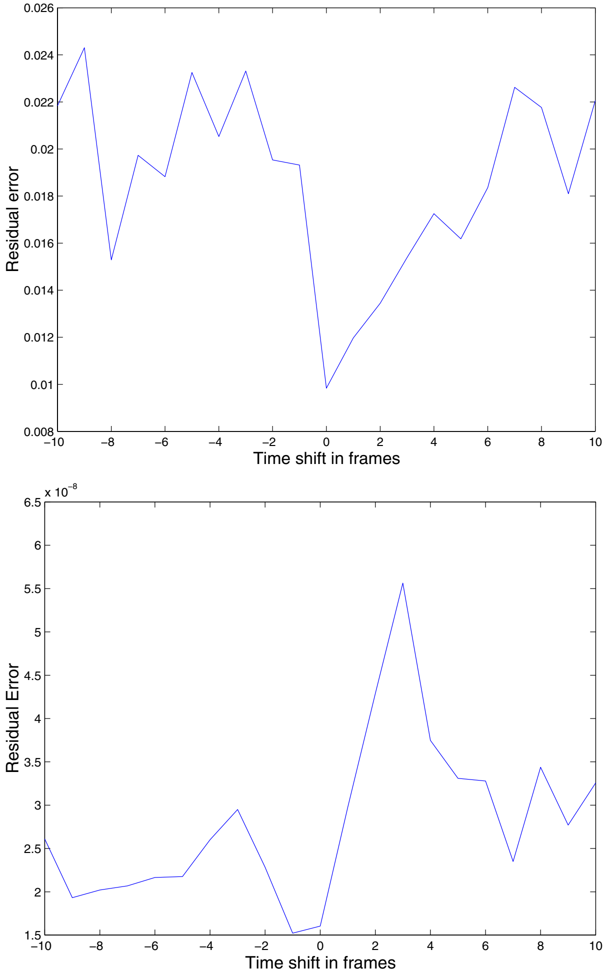

a)

b)

Fig. 1. Sequences synchronization for cameras viewing opposite directions (Fig. 1-a) and for cameras viewing orthogonal directions (Fig. 1-b). The residual error of the estimation matrix is presented as a function of the temporal shift between the sequences, where 0 shift corresponds to the manual integer estimation of the real shift. The residual error for the orthogonal case is averaged over four random choices of the camera matrix O.

7 Summary and Future Work

We have analyzed self calibration of rigs with varying internal parameters, when no correspondences are available between the cameras. The only assumption used for the affine calibration, was that the angle between the viewing directions of the cameras is known.

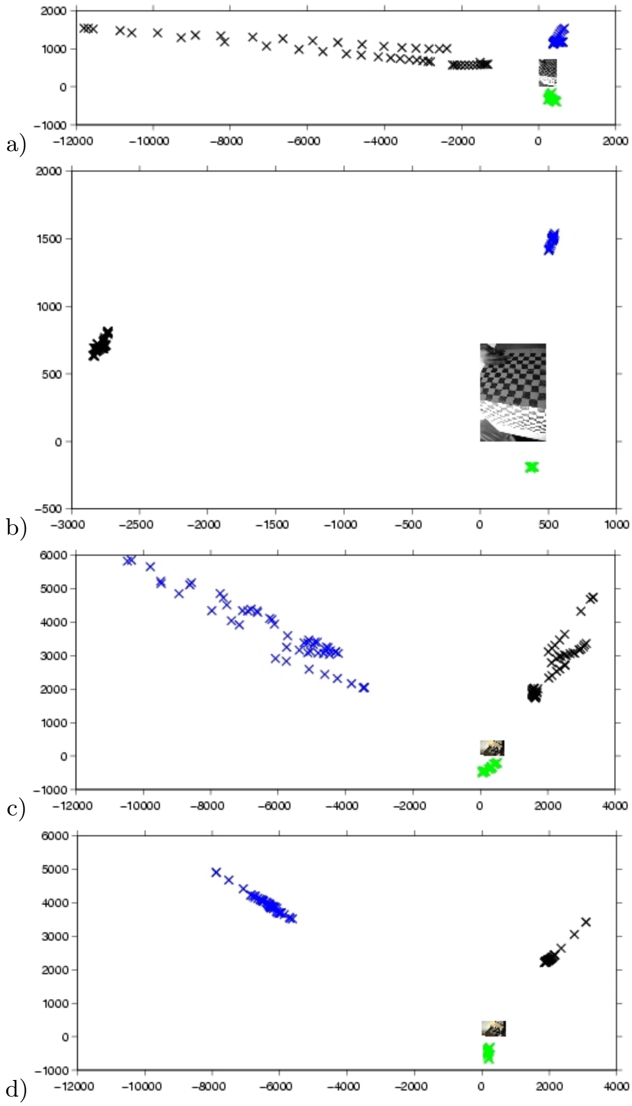

Fig. 2. Mapping vanishing points by the homographies through the plane at infinity extracted by the proposed algorithm. a) The original vanishing points, opposite viewing directions. b) The mapped vanishing points, opposite viewing directions. c) The original vanishing points, orthogonal viewing directions. d) The mapped vanishing points, orthogonal viewing directions. The size of the images in illustrations a,b and c,d visualize the scale differences between the coordinate axes.

We presented constraints for general angles, and solved specific cases of parallel, opposite and orthogonal angles by extracting affine invariants.

Future work can be solving similar cases by deriving metric invariants, e.g.:

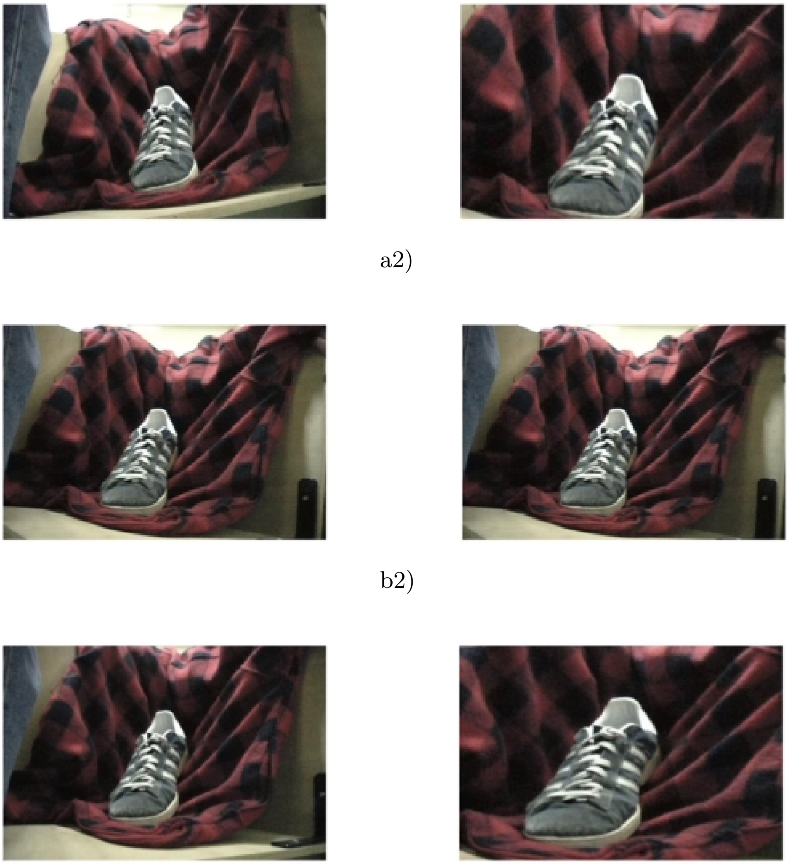

a1) a2)

b1) b2)

c1) c2)

Fig. 3. Using the computed homographies through the planes at infinity to generate a virtual motion of pure translation. Figures b1,b2 contain the same image, to which the input images were warped. Figures a1,c1 to the left are the original input images. Figures a2,c2 to the right are the results of warping a1,c1 by the respective homographies through the plane at infinity.

– Solving the case of arbitrary angles, with or without knowing the angle.
– Assuming that the distances between the centers of projections of the cameras remain fixed during the motion.

We hope that even if such invariants are not compact, they may find use in some application such as cameras synchronization.

We plan to implement such algorithms for self calibration on “domes” containing cameras with fixed orientations and varying Zoom/Focus.

References

1. Y. Caspi, and M. Irani Alignment of Non-Overlapping Sequences In *Proceedings of the International Conference on Computer Vision*, Vol, II, pages 76-83, Vancouver, Canada, July 2001.
2. F. Dornaika, F. Self-Calibration of a Stereo Rig Using Monocular Epipolar Geometry. In *Proceedings of the International Conference on Computer Vision*, Vol, II, pages 467-472, Vancouver, Canada, July 2001.
3. 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.
4. 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.
5. 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.
6. R.I. Hartley and A. Zisserman. *Multiple View Geometry.* Cambridge University Press, 2000.
7. R. Horaud, G. Csurka, and D. Demirdijian Stereo Calibration from Rigid Motions *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(22):1446-1452, December 2000.
8. 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.
9. A. Shashua and L. Wolf. Homography Tensors: On Algebraic Entities That Represent Three Views of Static or Moving Planar Points. in Proc. of the European Conference on Computer Vision (ECCV), June 2000, Dublin, Ireland.
10. 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.
11. 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.
12. A. Zisserman, P.A. Beardsley, and I.D. Reid. Metric calibration of a stereo rig. In *Proc. of IEEE International Conference in Computer Vision*, Vol. I. pages 135-141, Vancouver, Canada, July 2001.
13. A. Zomet, L. Wolf and A. Shashua. Omni-Rig: Linear Self-Recalibration of a Rig with Varying Internal and External Parameters In *Proceedings of the International Conference on Computer Vision*, Vol. I, pages 135-141, Vancouver, Canada, July 2001.