

Critical Motion Sequences for the Self-Calibration of Cameras and Stereo Systems with Variable Focal Length

Peter F Sturm*

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

Abstract.

We consider the self-calibration problem for a moving camera whose intrinsic parameters are known, except the focal length, which may vary freely across different views. The conditions under which the determination of the focal length's values for an image sequence is not possible, are derived. These depend only on the camera's motions. We give a complete catalogue of the so-called *critical motion sequences*. This is then used to derive the critical motion sequences for stereo systems with variable focal lengths.

1 Introduction

One of the major goals of computer vision is the recovery of spatial information about the environment. Classical approaches assume that the cameras are *calibrated* beforehand, but a great interest in *uncalibrated* vision and on-line calibration has arisen during the last decade. A key result is that even with completely uncalibrated cameras, spatial information – *projective structure* – can be obtained: the scene can be reconstructed up to an unknown projective transformation [5, 8]. Furthermore, a moving camera can *self-calibrate*, i.e. the calibration parameters can be estimated solely from feature correspondences between several images [9, 11]. This allows the projective ambiguity in the reconstruction to be reduced to a Euclidean one (up to a similarity transformation), and we speak of *uncalibrated Euclidean reconstruction*.

It is known that several types of camera motion prevent self-calibration, i.e. the calibration parameters can not be determined uniquely. Accordingly, Euclidean structure can not be obtained, although reconstruction at some level between projective and Euclidean is generally possible. These degeneracies are inherent, i.e. they can not be resolved by any algorithm without additional knowledge. Sequences of camera motions that imply such degeneracies will be referred to as *critical motion sequences*. By “sequences” we mean that not only the motion between two successive views, but that over a complete image sequence, is critical.

For the basic self-calibration scenario, a moving camera with *fixed* calibration, the critical motion sequences were derived by Sturm [13]. In this paper, we study the case of a moving camera with *variable* and unknown focal length, but whose other intrinsic parameters are known. A practical self-calibration algorithm was proposed by Azarbajejani and Pentland [1]. Algorithms and closed-form solutions for the two-view case were given by Hartley and Brooks et al. [7, 3, 4, 10]. Newsam et al. derived the critical motions for the two-view case [10]. In this paper, we derive a complete characterization of critical motion *sequences*, for any number of views, and the critical motions for stereo systems.

The paper is organized as follows. In §2 we provide some theoretical background for our approach. The problem of deriving critical motion sequences is formulated in §3. The critical motion sequences are derived in §4. A summary of the derivations is given in §5 and comments are made in §6. The critical motions for stereo systems are derived in §7 and conclusions are drawn in §8. Many proofs are omitted in this paper; please contact the author for more details.

2 Background

The definitions in this section are mainly taken from [2] and [12]. Some of the results for general quadrics are presented only for central conics.

Notation. We refer to the *plane at infinity* as the *ideal plane* and denote it by Π_∞ . \mathcal{P}^n is the n -dimensional projective space and \sim means equality up to a scalar factor, accounting for the use of homogeneous coordinates. We use the abbreviation *PVC* for proper virtual conics (see below).

*This work is supported by the EPSRC funded project GR/K89221 (Vector).

Pinhole Camera Model. We use perspective projection to model cameras. A projection may be represented by a 3×4 projection matrix P that maps points of 3-space to points in 2-space: $\mathbf{q} \sim P\mathbf{Q}$. We consider only the case of perfect perspective projection, i.e. the projection center does not lie on Π_∞ .

With regard to physical cameras, the projection matrix may be decomposed into a *calibration matrix* K and a *pose matrix*. The pose matrix represents the position and orientation of the camera in some world coordinate frame. In general, we distinguish 5 intrinsic parameters for the perspective projection model: the (effective) focal length f , the aspect ratio τ , the principal point (u_0, v_0) and a skew factor accounting for non rectangular pixels. The skew factor is usually very close to 0 and we ignore it in the following. The calibration matrix may be written as:

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

We decompose the projection matrix as follows:

$$P = K \underbrace{\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}}_{P_c} \begin{pmatrix} R & -R\mathbf{t} \\ \mathbf{0}^T & 1 \end{pmatrix}$$

The matrix P_c is the *canonical projection*, and we call its destination the *metric image plane*. The canonical projection depends only on the camera's extrinsic parameters – a rotation matrix R representing its orientation and a 3-vector \mathbf{t} representing its position. The calibration matrix K describes an invertible affine transformation from the metric image plane to pixel coordinates.

Quadrics and Conics. A *quadric* in \mathcal{P}^n is a set of points satisfying a homogeneous quadratic equation. Each quadric can be represented by a symmetric $(n+1) \times (n+1)$ matrix. A *virtual quadric* is a quadric with no real point and a *proper quadric* is a quadric whose matrix has a non-zero determinant. *Conics* are planar quadrics; we will not distinguish between a conic and its matrix. A conic in \mathcal{P}^3 , or *3D conic*, is defined by its *supporting plane* and the conic's equation in that plane. All proper virtual conics (*PVC*) are central [2] and hence can be transformed to *Euclidean normal form* by a Euclidean transformation (principal axis transformation). The Euclidean normal form is a diagonal matrix of the conic's eigenvalues.

Cones. By *cones* we mean rank-3 quadrics in \mathcal{P}^3 with vertex not on Π_∞ . A cone is uniquely defined by its vertex and any (conic) section by a plane not containing the vertex. Cones are used in this paper through the notion of the *projection cone* of a 3D conic, i.e. the cone traced out by the projection rays of the perspective projection of the conic. The Euclidean normal form of a cone is a diagonal matrix $\text{diag}(\lambda_1, \lambda_2, \lambda_3, 0)$, with non-zero λ_i . If the λ_i are all distinct then the cone is an *elliptic cone*. If exactly two of the λ_i are equal, the cone is *circular* (or *right*). For an *absolute cone*, all three λ_i are equal.

Absolute Quadric and Absolute Conic. The *absolute quadric* of \mathcal{P}^n is defined by the equations $x_1^2 + \dots + x_n^2 = x_{n+1} = 0$. The *absolute conic* Ω is the absolute quadric of \mathcal{P}^3 . Ω is a proper virtual conic in the ideal plane whose position uniquely defines the Euclidean structure of 3-space. The calibration of a camera is equivalent to determining the image ω of Ω , respectively its dual ω^* [9, 6]. From the relation $\omega^* \sim KK^T$, the calibration matrix K can be uniquely recovered by Cholesky decomposition.

3 Problem Formulation

We consider a sequence of n views, generally taken from different positions and with different orientations. The focal lengths for the views may all be different, but the other intrinsic parameters (aspect ratio and principal point) are known (they need not be equal for all the views). The problem at hand is to determine the n different values for the focal length, which allows in general to obtain a Euclidean reconstruction of the scene. In the following, we describe this problem in geometrical terms, in analogy to [13].

First, *calibration* of a camera is equivalent to the determination of the image of the absolute conic, as “produced” by that camera. *Self-calibration* means the same, but with the conotation that information used to calibrate does not stem from e.g. known metric 3D structure. *Euclidean reconstruction* is equivalent to the determination of the position of the absolute conic in 3D. The problem of Euclidean reconstruction is slightly more general than that of self-calibration: degeneracy of self-calibration implies degeneracy of Euclidean reconstruction while the reciprocal is not always true (e.g. self-calibration of a camera rotating about its optical center is in general possible, while any 3D reconstruction, Euclidean included, is impossible). The derivations that follow refer to degeneracies of Euclidean reconstruction.

To determine the position of the absolute conic, some constraints are needed. We will describe these constraints in the following paragraph, but first we give a straightforward informal definition for degeneracy of Euclidean reconstruction: the Euclidean reconstruction problem is degenerate exactly if there is a conic in 3D, not identical with the absolute conic, that satisfies the mentioned constraints [13]. All such conics will be called *potential absolute conics*.

We now describe the constraints that may be used to determine the absolute conic. First, the absolute conic must be a proper virtual conic. Second, the image of the absolute conic by any perfect perspective projection, is also a proper virtual conic. Third, the knowledge of some intrinsic parameters constrains the *projections* of the absolute conic in a given set of views, which in turn gives us constraints on the absolute conic itself. For the scenario considered here, we make these constraints explicit in the following.

The image of the absolute conic, as “produced” by a camera with calibration matrix K , is given by:

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

Since f may vary, there is, for each view, exactly one family of possible images of the absolute conic. Consider now a conic Φ in 3D and its projection ϕ in one view. For Φ being a potential absolute conic, its projection ϕ must be of the form (1), for some non zero real value a , possibly different from f :

$$\phi \sim \begin{pmatrix} 1 & 0 & -u_0 \\ 0 & \tau^2 & -\tau^2 v_0 \\ -u_0 & -\tau^2 v_0 & u_0^2 + \tau^2 v_0^2 + \tau^2 a^2 \end{pmatrix}$$

It is easy to show that a conic has this form exactly if, in the metric image plane, the conic is a virtual circle, centered in the origin. To see this, we map ϕ from pixel coordinates to the metric image plane:

$$\phi_m \sim K^T \phi K \sim \begin{pmatrix} f^2 & 0 & 0 \\ 0 & f^2 & 0 \\ 0 & 0 & a^2 \end{pmatrix}$$

For any non zero real value of a , this represents a proper virtual circle whose center is the coordinate origin. It is important to note that this statement is independent of the actual true value f of the focal length. Since all other intrinsic parameters are known, the only important parameters for the consideration of degeneracy of Euclidean reconstruction, are the extrinsic parameters of the views in a given sequence.

We now summarize the discussion.

Proposition 1. *Consider a sequence of n views with known aspect ratio and principal point but unknown and possibly different values for the focal length. Let P_{ci} be the canonical projection for view $i, i = 1, \dots, n$.*

Euclidean reconstruction is degenerate exactly if there is at least one 3D conic Φ , not identical with the absolute conic, such that:

- Φ is a proper virtual conic.
- the $\phi_i, i = 1, \dots, n$, where ϕ_i is the projection of Φ by P_{ci} , are proper virtual circles, centered in the origin. This is equivalent to the ϕ_i being represented by diagonal matrices whose diagonal elements are all non-zero real values of the same sign, the first two elements being equal.

Definition 2. Any Φ as defined in proposition 1 is called a **potential absolute conic**.

Consider a sequence of n views. Let (R_i, t_i) be the extrinsic parameters of view $i, i = 1, \dots, n$. If Euclidean reconstruction is degenerate for the sequence of views, we say that $\{(R_i, t_i) | i = 1, \dots, n\}$ is a **critical motion sequence** for Euclidean reconstruction. ■

The aim of the following section is to derive all generic critical motion sequences, i.e. all configurations where there is no unique solution to Euclidean reconstruction.

4 Derivation of the Critical Motion Sequences

In this section, the critical motion sequences are derived, based on proposition 1 and definition 2. We proceed in a constructive manner: given a generic 3D proper virtual conic Φ , we determine all possible extrinsic parameters that form a critical motion sequence with respect to Φ , i.e. for which Φ is a potential absolute conic. The derivations are divided into two parts, considering potential absolute conics Φ which lie / do not lie on Π_∞ . The results are summarized in §5.

4.1 Potential Absolute Conics on Π_∞

Let Φ be a PVC on the ideal plane. Its canonical projection ϕ by a camera with extrinsic parameters (\mathbf{R}, \mathbf{t}) is given by:

$$\phi \sim \mathbf{R}\Phi\mathbf{R}^\top \quad (2)$$

Like it is the case for all geometric entities on the ideal plane, the projection depends only on the orientation of the camera, not on its position.

We now determine all orientations \mathbf{R} for which ϕ is a proper virtual circle centered in the origin. This implies that ϕ is a diagonal matrix of the form: $\phi \sim \text{diag}(b, b, 1)$. If we choose the free scale factor for Φ such that $\det \Phi = b^2 = \det \phi$, the \sim in equation (2) can be replaced by an equality sign:

$$\begin{pmatrix} b & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & 1 \end{pmatrix} = \mathbf{R}\Phi\mathbf{R}^\top \quad (3)$$

Equation (3) implies that Φ has a double eigenvalue b and the single eigenvalue 1. The case $b = 1$ is of no interest here, because this would mean that Φ is the *true* absolute conic (the absolute conic is the only conic on the ideal plane with a triple eigenvalue).

From equation (3), we derive:

$$\mathbf{R}^\top \begin{pmatrix} b & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = \Phi\mathbf{R}^\top \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

and further:

$$\mathbf{R}^\top \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = \Phi\mathbf{R}^\top \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

Hence, the vector

$$\mathbf{v}_R = \mathbf{R}^\top \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

is an eigenvector of Φ to the eigenvalue 1. Since 1 is a single eigenvalue, all its associated eigenvectors are equal up to scale. This means that for all rotation matrices \mathbf{R} in a critical motion sequence, the vectors \mathbf{v}_R are equal up to scale.

It is easy to show that $(\mathbf{v}_R, 0)$ is nothing else than the ideal point of the optical axis for a camera with orientation \mathbf{R} . All \mathbf{v}_R being equal up to scale is thus equivalent to the optical axes of all views being parallel. We conclude that a set of views with parallel optical axes is critical for Euclidean reconstruction. The proof of this being a necessary and sufficient condition for a set of views to be critical with respect to a potential absolute conic on the ideal plane, is omitted here.

4.2 Potential Absolute Conics not on Π_∞

Contrary to conics on the ideal plane, the projection of conics not on Π_∞ depends on both camera position and orientation. First, we deal with position, then with orientation.

4.2.1 Position

Let Φ be a PVC *not* on the ideal plane. Consider a view with optical center at position \mathbf{t} . Let ϕ be the canonical projection of Φ . Let Λ be the projection cone of ϕ (cf. §2).

One condition for Φ being a potential absolute conic is that ϕ is a circle, centered in the origin of the metric image plane. Note that the origin of the metric image plane is the camera's principal point, i.e. the intersection of the optical axis with the image plane. Since the optical axis is perpendicular to the image plane, the projection cone Λ must be circular. The vertex of the projection cone being the optical center \mathbf{t} , we obtain constraints on the possible camera positions in a critical motion sequence: for a potential absolute conic Φ , all possible camera positions are the vertices of circular cones that contain Φ . These are summarized in the following (proof omitted).

If Φ is a virtual circle, the locus of possible camera positions in a critical motion sequence is the line \mathbf{L} perpendicular to the circle's supporting plane and passing through the circle's center (see figure 1). If Φ is a virtual ellipse, then the locus of camera positions is the union of a real ellipse Ψ_e and a real hyperbola Ψ_h (see figure 1). The supporting plane of the three conics Φ , Ψ_e and Ψ_h are mutually perpendicular (there are further relations between these conics, but they are omitted here).




Figure 1: Locus of camera positions in a motion sequence critical with respect to a conic Φ not on the ideal plane. The conic Φ is shown in dotted style to illustrate that it is virtual and can in fact not be drawn. On the left hand side, the case of Φ being an ellipse is shown and on the right hand side the case of a circle. The figures are further explained in §4.2.

4.2.2 Orientation

The results of the previous paragraph are necessary conditions for camera positions in a critical motion sequence. To obtain sufficient conditions, we now consider the *orientation* of cameras. First note that rotations about the optical axis are not important here: if the projection of Φ is a circle centered in the origin (the principal point), then any rotation about the optical axis will preserve this property. Hence, the only part of camera orientation that matters is the *orientation of the optical axis*.

For the camera positions derived in the previous paragraph, we have to determine the orientations of the optical axis such that Φ is projected onto a circle centered in the origin. The proofs of the following statements are omitted due to lack of space.

If Φ is a virtual circle (cf. right hand part of figure 1), there are two particular camera positions on \mathbf{L} for which the optical axis might take any orientation. These positions are symmetric with respect to the supporting plane of Φ . For all other camera positions on the line \mathbf{L} , the optical axis coincides with \mathbf{L} . To summarize, critical motion sequences with respect to a *circle* Φ consist of collinear optical centers and optical axes passing through all optical centers, except that at two positions, the optical axis may be oriented arbitrarily (even several views might be taken from these two positions, by a camera rotating about its optical center).

If Φ is a virtual ellipse (cf. left hand part of figure 1), and the optical center lies on the ellipse Ψ_e (respectively the hyperbola Ψ_h), then the optical axis is the tangent of Ψ_e (respectively Ψ_h) at the optical center.

5 Summary

The following camera positions/orientations constitute critical motion sequences for Euclidean reconstruction:

Case 1. Arbitrary position of optical centers but parallel optical axes. This means that camera motions are pure translations, possibly combined with an arbitrary rotation about the optical axis and a reversal of the gaze direction. ■

Case 2. Collinear optical centers. The optical axes at two positions may be oriented arbitrarily, all others coincide with the line joining the optical centers. This means that camera motions are pure forward translations, with two exceptions where the translation may be followed by an arbitrary rotation about the optical center. ■

Case 3. The optical centers lie on an ellipse/hyperbola pair as shown in figure 1. At each position, the optical axis is tangent to the ellipse/hyperbola. A necessary condition derived from this is: the views may be partitioned into at most two sets for which the centers and optical axes are all coplanar. In addition, these two sets define perpendicular planes.

If we consider only one of the two conics Ψ_e or Ψ_h , we can describe the critical motion sequences as follows: a camera moving on a trajectory that may be described as (arc of) a conic, and always gazing in the direction of motion. An example would be a camera mounted on a vehicle taking a roundabout or another “smooth” turn. ■

Newsam et al. derived degenerate configurations for two-view self-calibration [10]. Their results are of course contained in the above list.

6 Comments

6.1 Two Cameras in General Position

From case 2 in the previous section, it follows that with two views in general position and with general orientation, Euclidean reconstruction is degenerate. In fact, it is known that in this case there is a two-fold ambiguity for the absolute conic: let ω_1 and ω_2 be the projections of the absolute conic in the two views. The projection cones of ω_1 and ω_2

intersect of course in the absolute conic, but in general also in a second conic Φ , hence the ambiguity. However, self-calibration can be achieved, since Φ has the same projections as the absolute conic (it lies on the same projection cones). This illustrates that Euclidean reconstruction and self-calibration are not exactly equivalent problems. In practice, the ambiguous solution for Euclidean reconstruction can often be ruled out by imposing that the reconstructed scene lies in front of both cameras.

6.2 Camera Rotating about its Optical Center

A camera rotating about its optical center while possibly changing its focal length, can always be calibrated from two views whose optical axes do not coincide. This is briefly explained in the following. The two views are critical for Euclidean reconstruction according to case 2 (cases 1 and 3 allow only one optical axis per camera position). The position \mathbf{t} of the two views is one of the two exceptional points. What we did not mention above is the following fact: the cone spanned by Φ and with vertex \mathbf{t} contains the absolute conic (it is an absolute cone). Hence, the projections of Φ are identical with the image of the absolute conic, which means that (self-) calibration has a unique solution.

6.3 Planar Motion

By planar motion we understand that both optical center and optical axis are bound to move in a single plane. It has been shown by Brooks et al. that for two views, this is always a critical configuration for self-calibration [3]. If the two optical axes are parallel, then the configuration is critical according to both cases 1 and 3. If not, then the configuration is critical according to case 3. In that case, we can describe the configuration as that of a moving camera, that fixates a finite point (the intersection of the two optical axes).

Adding other views that fixate the same point, makes the whole configuration non critical. This can be shown by contradiction using a simple geometrical argument. Suppose the configuration were still critical. According to case 3, and because the camera moves in a single plane, the camera positions have to lie on a conic (the ellipse Ψ_e or the hyperbola Ψ_h). In addition, the optical axes are tangents to the same conic. However, more than two views fixate the same point, i.e. the same point lies on more than two optical axes. This is not possible, since from one point, we can draw at most two tangents to a conic [12].

6.4 On Case 3

The case of a camera moving on a single conic, has been described in the previous paragraph. At first sight, the case of a motion sequence containing both conics Ψ_e and Ψ_h does not seem to be of much practical importance, since the optical axes of views are bound to lie in two perpendicular planes. However, this does not exclude that there is a large overlap between views, as figure 2 illustrates. This case might well occur in practical situations. Of course, the case of coplanar optical axes should be more important, e.g. for the consideration of stereo systems.




Figure 2: On case 3: on the left hand side, possible orientations of the optical axes of two views are shown, one's center lying on Ψ_e and the other's on Ψ_h . On the right hand side, a pair of optical axes is shown for which the fields of view will overlap.

7 Stereo Systems

The results presented in the previous sections allow us to study degeneracies of Euclidean reconstruction for stereo systems. Here, a stereo system consists of two cameras with coplanar optical axes and symmetric but possibly variable vergence angles. The distance between the two cameras is fixed and the system is assumed to move in a single plane. The focal lengths of the two cameras are not constrained to be equal, and they may vary freely between different images. We define the *axis of a stereo system* as the line perpendicular to and passing through the midpoint of the baseline.

A single pair of images taken by such a stereo system is of course critical. In the following, we reveal the conditions for two stereo pairs being critical. This is done for the cases of fixed and variable vergence angles.

7.1 Fixed Vergence Angles

For the trivial case of a stationary stereo system, any number of stereo pairs are critical. Non-trivial cases are discussed in the following.

Parallel optical axes. If the vergence angles are of 90° , i.e. if the two optical axes are parallel, then each stereo pair is critical according to cases 1 and 3. According to case 1, the combination of two stereo pairs is only critical if the stereo system undergoes a pure translation or a rotation by 180° in the plane of motion, i.e. a reversal of gaze direction. Case 3 is dealt with as in the following paragraph.

Convergent optical axes. The stereo pairs can only be critical according to case 3 if there is an ellipse or a hyperbola Ψ containing all 4 optical centers and having all optical axes as tangents. It is easy to show that due to symmetry of the vergence angles, Ψ must be symmetric with respect to the axes of the stereo pairs, but that Ψ cannot be a circle. This implies that the axes of the stereo pairs are either identical or perpendicular. It can be shown that in case they are perpendicular, there is no possibility to place the two stereo pairs in a way that a conic Ψ as described above exists. In case the axes are identical, a conic Ψ exists exactly if the two stereo pairs have opposite gaze direction.

Summary. Two stereo pairs with fixed and symmetric vergence angles are critical in exactly the following situations. If the optical axes of the stereo system are parallel, the stereo pairs are critical if they are related by a pure translation, possibly followed by a reversal of gaze direction. If the optical axes are convergent, the stereo pairs are only critical if they gaze in opposite directions and if their axes are identical. The only case of practical importance is pure translation of a stereo system with parallel optical axes.

7.2 Variable Vergence Angles

Due to varying vergence angles, at least one stereo pair has convergent optical axes. Hence, the combination of the stereo pairs can only be critical according to case 3. As stated above, the conic Ψ must be symmetric with respect to the axes of the stereo pairs, which implies that these are either identical or perpendicular. It can be shown that if they are identical, no conic Ψ as described above can exist. For perpendicular axes, there is a one-dimensional family of possibilities for placing the stereo pairs and setting their vergence angles relative to each other such that the combination of the stereo pairs is critical.

In summary, if the vergence angles for the two stereo pairs are different, the configuration can only be critical if the axes of the stereo pairs are perpendicular. The exact conditions for being critical are omitted due to lack of space.

8 Conclusions

We have derived all motion sequences that are critical for Euclidean reconstruction from image sequences with variable and unknown focal length, whose other intrinsic parameters are known. The critical motion sequences are described geometrically. The most important case is probably that of a camera moving on a straight line or a conic, and gazing in the direction of motion. Our results allowed us to study the critical motions of stereo systems. We showed that there are very few critical motions and that these, besides the case of pure translation, are not likely to occur in practice.

For details and proofs that have been omitted in this paper, please contact the author.

References

- [1] A. Azarbayejani and A.P. Pentland, "Recursive Estimation of Motion, Structure, and Focal Length," *IEEE Transactions on PAMI*, Vol. 17, No. 6, 562-575, 1995.
- [2] W. Boehm and H. Prautzsch, "Geometric Concepts for Geometric Design," A.K. Peters, 1994.
- [3] M.J. Brooks, L. de Agapito, D.Q. Huynh and L. Baumela, "Direct Methods for Self-Calibration of a Moving Stereo Head," *ECCV, UK*, 415-426, 1996.
- [4] L. de Agapito, D.Q. Huynh and M.J. Brooks, "Self-Calibrating a Stereo Head: An Error Analysis in The Neighbourhood of Degenerate Configurations," *ICCV, India*, 747-753, 1998.
- [5] O. Faugeras, "What Can Be Seen in Three Dimensions with An Uncalibrated Stereo Rig ?", *ECCV, Italy*, 563-578, 1992.
- [6] O. Faugeras, "Stratification of Three-Dimensional Vision: Projective, Affine and Metric Representations", *Journal Optical Society of America A*, 465-484, 1995.
- [7] R.I. Hartley, "Estimation of Relative Camera Positions for Uncalibrated Cameras," *ECCV, Italy*, 579-587, 1992.
- [8] R. Hartley, R. Gupta and T. Chang, "Stereo from Uncalibrated Cameras", *ECCV, Italy*, 761-764, 1992.

- [9] S.J. Maybank and O.D. Faugeras, "A Theory of Self Calibration of a Moving Camera," *International Journal on Computer Vision*, Vol. 8, No. 2, 123-151, 1992.
- [10] G.N. Newsam, D.Q. Huynh, M.J. Brooks and H.P. Pan, "Recovering Unknown Focal Lengths in Self-Calibration: An Essentially Linear Algorithm and Degenerate Configurations," *Proc. of the XVIII ISPRS-Congress, Vienna, Austria*, Vol. XXXI, Part B3, 575-580, 1996.
- [11] M. Pollefeys, R. Koch and L. Van Gool, "Self-Calibration and Metric Reconstruction in Spite of Varying and Unknown Internal Camera Parameters," *ICCV, India*, 90-95, 1998.
- [12] J.G. Semple and G.T. Kneebone, "Algebraic Projective Geometry", Oxford Science Publication, 1952.
- [13] P. Sturm, "Critical Motion Sequences for Monocular Self-Calibration and Uncalibrated Euclidean Reconstruction," *CVPR, Puerto Rico*, 1100-1105, 1997.