



From Projective to Euclidean Reconstruction

Frédéric Devernay, Olivier Faugeras

► To cite this version:

Frédéric Devernay, Olivier Faugeras. From Projective to Euclidean Reconstruction. CVPR - International Conference on Computer Vision and Pattern Recognition - 1996, 1996, San Francisco, United States. pp.264 - 269, 10.1109/CVPR.1996.517084 . hal-00816706

HAL Id: hal-00816706

<https://inria.hal.science/hal-00816706>

Submitted on 22 Apr 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

From Projective to Euclidean Reconstruction

Frédéric Devernay Olivier Faugeras
INRIA Sophia Antipolis, 2004 route des Lucioles
BP 93, 06902 Sophia Antipolis Cedex

Abstract

To make a Euclidean reconstruction of the world seen through a stereo rig, we can either use a calibration grid, and the results will rely on the precision of the grid and the extracted points of interest, or use self-calibration. Past work on self-calibration is focussed on the use of only one camera, and gives sometimes very unstable results.

In this paper, we use a stereo rig which is supposed to be weakly calibrated using a method such as the one described in [1]. Then, by matching two sets of points of the same scene reconstructed from different points of view, we try to find both the homography that maps the projective reconstruction [2] to the Euclidean space and the displacement from the first set of points to the second set of points.

We present results of the Euclidean reconstruction of a whole object from uncalibrated cameras using the method proposed here.

1 Introduction

This article is concerned with the following problem. Given a weakly calibrated stereo rig, i.e. a pair of camera with known epipolar geometry, we know that we can obtain 3-D reconstructions of the environment up to an unknown projective transformation [2, 5]. We call such a reconstruction a *projective* reconstruction. In particular, no affine or euclidean information can a priori be extracted from it unless some further information is available [4]. The problem is then to determine what is the information that is missing and how can it be recovered. We provide a very simple answer to both questions: with one rigid displacement of the stereo rig, the three-dimensional structure of the scene can be in general uniquely recovered up to a similitude transformation using some elementary matrix algebra, assuming that reliable correspondences between the two projective reconstructions obtained from the two viewpoints can be established. We call such a reconstruction a *euclidean* reconstruction. A similar result was obtained [8] but the resulting scheme was a closed form solution computed from two views of

the scene, whereas this method can be used with many more views, giving more stability on the solution.

This result does not contradict previous results, for example [7, 6] which showed that the *intrinsic* parameters of a camera could be in general recovered from two displacements of the camera because we are using simultaneously two cameras. The method developed here avoids any reference to the intrinsic parameters of the cameras and does not require solving the nonlinear Kruppa equations which are defined in the previous references.

2 Goal of the method

Our acquisition system consists of a pair of cameras. This system can be calibrated using a weak calibration method [1], so that we can make a projective reconstruction [2] of the scene in front of the stereoscopic system, by matching features (points, curves, or surfaces) between the two images.

Projective reconstruction roughly consists of choosing five point matches between the two views and choosing these five points as a projective basis to reconstruct the scene. The five point matches can be either real points (i.e. points that are physically present in the scene) or virtual points. The virtual point matches are calculated by choosing a point in the first camera, and then choosing any point on its epipolar line in the second camera as its correspondent, thus these points satisfy the epipolar constraint but are not the images of a physical point. Let us call $\mathcal{P}$ the resulting projective basis which is thus *attached* to the stereo rig.

Let us now consider a real correspondence (m_1, m'_1) between the two images. We can reconstruct the 3-D point M_1 in the projective basis $\mathcal{P}$. Let us now suppose that after moving the rig to another place, the correspondence has become (m_2, m'_2) , yielding a 3-D reconstructed point M_2 in the projective basis $\mathcal{P}$. We know from the results of [2, 5] that the two reconstructions are related by a collineation of $\mathcal{P}^3$ which is represented by a full rank 4×4 matrix $\mathbf{H}_{12}$ defined up to a scale factor. We denote by the symbol $\cong$ the

equality *up to a scale factor*. Thus we have

$$\mathbf{M}_2 \cong \mathbf{H}_{12} \mathbf{M}_1$$

where $\mathbf{M}_1$ and $\mathbf{M}_2$ are homogeneous coordinate vectors of M_1 and M_2 in $\mathcal{P}$.

Let us now imagine for a moment that an orthonormal frame of reference $\mathcal{E}$ is attached to the stereo rig. The change of coordinates from $\mathcal{P}$ to $\mathcal{E}$ is described by a full rank 4×4 matrix $\mathbf{H}_{12}$, also defined up to a scale factor. In the coordinate frame $\mathcal{E}$ the two 3-D reconstructions obtained from the two viewpoints are related by a rigid displacement, not a general collineation. This rigid displacement is represented by the following 4×4 matrix $\mathbf{D}_{12}$:

$$\mathbf{D}_{12} \cong \begin{bmatrix} \mathbf{R}_{12} & \mathbf{t}_{12} \\ \mathbf{0} & 1 \end{bmatrix}$$

where $\mathbf{R}_{12}$ is a rotation matrix. It is well known and fairly obvious that the displacement matrixes form a subgroup of $SL(4)$ which we denote by $E(3)$.

We can now relate the three matrixes $\mathbf{H}_{12}$, $\mathbf{H}$, and $\mathbf{D}_{12}$ (see figure 1):

$$\mathbf{H}_{12} \cong \mathbf{H}^{-1} \mathbf{D}_{12} \mathbf{H} \quad (1)$$

Since the choice of $\mathcal{E}$ is clearly arbitrary, the matrix $\mathbf{H}$ is defined up to an arbitrary displacement. More precisely, we make no difference between matrix $\mathbf{H}$ and matrix $\mathbf{D}\mathbf{H}$ for an arbitrary element $\mathbf{D}$ of $E(3)$. In mathematical terms, this means that we are interested only in the quotient $SL(4)/E(3)$ of the group $SL(4)$ by its subgroup $E(3)$. Therefore, instead of talking about the matrix $\mathbf{H}$ we talk about its equivalence class $\bar{\mathbf{H}}$. The basic idea of our method is to select in the equivalence class a *canonical element* $\hat{\mathbf{D}}\hat{\mathbf{H}}$, which is the same as selecting a special euclidean frame $\hat{\mathcal{E}}$ among all possible ones and show that equation (1) can be solved in general uniquely for $\hat{\mathbf{H}}$ and $\hat{\mathbf{D}}' \cong \hat{\mathbf{D}}^{-1} \mathbf{D}_{12} \hat{\mathbf{D}}$.

3 Colineations modulo a displacement

3.1 First method

Finding a unique representative of the equivalence classes of the group $SL(4)$ modulo a displacement in $E(3)$ is equivalent to finding a unique decomposition of a collineation (which depends upon 15 parameters) into the product of a displacement (which depends upon 6 parameters) and a member of a subgroup of dimension $15 - 6 = 9$. In fact, we are looking for something similar to the well-known QR or QL decompositions of a matrix into an orthogonal matrix and an upper or lower triangular matrix, where “orthogonal” would be replaced by “displacement”.




Figure 1: Given the collineation $\mathbf{H}_{12}$ we want to find the collineation $\mathbf{H}$ that maps the projective reconstruction to the euclidean reconstruction and the displacement $\mathbf{D}_{12}$.

Let us thus consider an element $\mathbf{H}$ of $SL(4)$ and assume that the element h_{44} is non zero. We define the 3×1 vector $\mathbf{t}$ by

$$\mathbf{t} = [h_{14}/h_{44}, h_{24}/h_{44}, h_{34}/h_{44}]^T, \quad (2)$$

and write $\mathbf{H}$ as

$$\mathbf{H} = h_{44} \begin{bmatrix} \mathbf{I}_3 & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{1}^T & 1 \end{bmatrix} \quad (3)$$

Note that since $\det \mathbf{H} = h_{44} \det \mathbf{A} \neq 0$, this implies that $\det \mathbf{A} \neq 0$. Then there is a unique QL decomposition of $\mathbf{A}$, so that

$$\mathbf{H} = h_{44} \begin{bmatrix} \mathbf{Q} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} cc\mathbf{L} & \mathbf{0} \\ \mathbf{1}^T & 1 \end{bmatrix} \quad (4)$$

where $\mathbf{Q}$ is orthogonal and $\mathbf{L}$ is lower triangular with strictly positive diagonal elements. Thus the group $SL(4)$ modulo the displacements $E(3)$ is isomorphic to the group of the lower triangular matrices with strictly positive diagonal elements. $\mathbf{Q}$ is a rotation if $\det \mathbf{H} > 0$, or a plane symmetry if $\det \mathbf{H} < 0$ (remember that the sign of $\det \mathbf{H}$ cannot be changed because $\mathbf{H}$ is of dimension 4).

If we want to decompose $\mathbf{H}$ into a rotation and a translation, we have to remove the constraint on the sign of one the elements of the diagonal of $\mathbf{L}$, e.g. there is no constraint on the sign of the first element of $\mathbf{L}$. In practice, the decomposition will be done using a standard QL decomposition, and then if $\mathbf{Q}$ is a plane symmetry rather than a rotation we just have to change the sign of the first element of $\mathbf{L}$ and of the first column of $\mathbf{Q}$, so that the multiplication of both matrices gives the same result and $\mathbf{Q}$ becomes a rotation.

3.2 Second method

Another way to find a unique representative of the equivalence classes of the group of collineations modulo a displacement is to build these representatives by applying constraints on the group of collineations corresponding to the degrees of freedom of a displacement. A simple representative is the one such that the image of the origin is the origin (i.e. the translational term of the collineation is zero), the z axis is globally invariant (i.e. the axis of the rotational term is the z axis), and the image of the y axis is in the yz plane the sign of the y coordinate being invariant (i.e. the angle of the rotation is zero).

These constraints correspond to constraints on the form of matrix $\mathbf{H}$. The image of the origin by $\mathbf{H}$ is the origin itself iff:

$$\mathbf{H} [0, 0, 0, 1] = [0, 0, 0, a] \quad (5)$$

The z axis is globally invariant iff:

$$\mathbf{H} [0, 0, 1, 0] = [0, 0, b, c] \quad (6)$$

And the last constraint (the angle of the rotation is zero) corresponds to:

$$\mathbf{H} [0, 1, 0, 0] = [0, d, e, f] \quad (7)$$

and a , d , and f have the same sign. Consequently, $\mathbf{H}$ being defined up to a scale factor and non-singular, it can be written as:

$$\mathbf{H} = \begin{bmatrix} g & 0 & 0 & 0 \\ h & d & 0 & 0 \\ j & e & b & 0 \\ k & f & c & 1 \end{bmatrix} \quad (8)$$

with $d > 0$ and $f > 0$. Thus equation 1 becomes:

$$\mathbf{H}_{12} \cong \mathbf{L}^{-1} \mathbf{D}_{12} \mathbf{L} \quad (9)$$

where $\mathbf{L}$ is a lower triangular matrix with the second and third coordinates of the diagonal positive and the last set to 1.

4 Back to the Euclidean world

In this section we show how to recover partly the Euclidean geometry from two projective reconstruction of the same scene. The only thing we have to do is to solve equation 1 for a lower triangular $\mathbf{H}$. Let us first establish some properties of the colineation between the two reconstructions.

Proposition 1 *Let $\mathcal{A}$ and $\mathcal{B}$ be two projective reconstructions in $\mathcal{P}_3$, the projective space of dimension 3, of the same scene using the same projection matrices*

from different points of view. Let $\mathbf{H}_{12}$ be the projective transformation (or colineation) from $\mathcal{B}$ to $\mathcal{A}$. Then The eigenvalues of $\mathbf{H}_{12}$ are α (with order of multiplicity 2), $\alpha e^{i\theta}$, and $\alpha e^{-i\theta}$, with $\alpha = \sqrt[4]{\det \mathbf{H}_{12}}$, and the last coordinate of $\mathbf{H}_{12}$, h_{44} , is not zero.

Equation 1 yields that $\mathbf{H}_{12}$ and $\mathbf{D}_{12}$ are conjugate (up to a scale factor), then $\mathbf{H}_{12}/\sqrt[4]{\det \mathbf{H}_{12}}$ and $\mathbf{D}_{12}$ have the same eigenvalues, which are: 1 with order of multiplicity two, $e^{i\theta}$, and $e^{-i\theta}$.

Before continuing, we have to prove the following lemma:

Lemma 2 *for each 3×3 real matrix $\mathbf{A}$ whose eigenvalues are $(1, e^{i\theta}, e^{-i\theta})$, there exists a 3×3 lower triangular matrix $\mathbf{L}$ ($l_{ik} = 0$ for $k > i$) with $l_{ii} > 0$, $i = 1, 2, 3$ defined up to a scale factor, and a rotation $\mathbf{R}$, satisfying $\mathbf{A} = \mathbf{L}^{-1} \mathbf{R} \mathbf{L}$.*

Since its eigenvalues are either real or conjugate of each other, a real matrix whose eigenvalues are of module one can be decomposed in the form $\mathbf{A} = \mathbf{P} \mathbf{D}_{12} \mathbf{P}^{-1}$, where $\mathbf{D}_{12}$ is a quasi-diagonal matrix of the form:

$$\mathbf{D}_{12} = \begin{bmatrix} B_1 & & 0 \\ & \ddots & \\ 0 & & B_k \end{bmatrix} \quad (10)$$

$$\text{with } B_i = [\pm 1] \text{ or } \begin{bmatrix} \cos \theta_i & -\sin \theta_i \\ \sin \theta_i & \cos \theta_i \end{bmatrix}$$

We can then compute the QL decomposition of $\mathbf{P}^{-1}$, $\mathbf{P}^{-1} = \mathbf{Q} \mathbf{L}$ which gives:

$$\mathbf{A} = \mathbf{L}^{-1} \mathbf{Q}^T \mathbf{D}_{12} \mathbf{Q} \mathbf{L} = \mathbf{L}^{-1} \mathbf{R} \mathbf{L}$$

where $\mathbf{L}$ is a lower triangular matrix with positive diagonal elements, and $\mathbf{R}$ is an orthogonal matrix. Since $\det \mathbf{A} = \det \mathbf{R} = 1$, then $\mathbf{R}$ is a rotation. $\square$

We now have all the tools needed to prove the following theorem.

Theorem 3 *Let $\mathcal{A}$ and $\mathcal{B}$ be two projective reconstructions of the same scene using the same projection matrices from different points of view. Let $\mathbf{H}_{12}$ be the projective transformation (or colineation) from $\mathcal{B}$ to $\mathcal{A}$. $\mathbf{H}_{12}$ can be decomposed in the form $\mathbf{H}_{12} = \lambda \mathbf{L}^{-1} \mathbf{D}_{12} \mathbf{L}$, where $\mathbf{L}$ is lower triangular and $\mathbf{D}$ is a displacement. The set of solutions is a two-dimensional manifold, one dimension being the scale factor on the Euclidean space.*

If we take three reconstructions taken from generic points of view, the full Euclidean geometry can be recovered, up to a scale factor.

Let us suppose that $\det \mathbf{H}_{12} = 1$ to eliminate the scale factor on $\mathbf{H}_{12}$. Let $\begin{pmatrix} l \\ 1 \end{pmatrix}$ be an eigenvector of $\mathbf{H}_{12}^T$ corresponding to the eigenvalue 1. This implies:

$$\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{l}^T & 1 \end{bmatrix} \mathbf{H}_{12} = \begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{l}^T & 1 \end{bmatrix} \quad (11)$$

so that H can be decomposed in the form:

$$\mathbf{H}_{12} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{l}^T & 1 \end{bmatrix} \begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{l}^T & 1 \end{bmatrix} \quad (12)$$

$$\mathbf{H}_{12} = \begin{bmatrix} \mathbf{A} + \mathbf{b}\mathbf{l}^T & \mathbf{b} \\ \mathbf{l}^T((1 - \mathbf{l}^T\mathbf{b})\mathbf{I} - \mathbf{A}) & 1 - \mathbf{l}^T\mathbf{b} \end{bmatrix} \quad (13)$$

Using the lemma 2, $\mathbf{A}$ can be decomposed into:

$$\mathbf{A} = \mathbf{L}^{-1}\mathbf{R}\mathbf{L} \quad (14)$$

and we can write $\mathbf{b}$ as:

$$\mathbf{b} = \mathbf{L}^{-1}\mathbf{t} \quad (15)$$

Thus,

$$\mathbf{H}_{12} = \begin{bmatrix} \mathbf{L}^{-1}\mathbf{R}\mathbf{L} + \mathbf{L}^{-1}\mathbf{t}\mathbf{l}^T & \mathbf{L}^{-1}\mathbf{t} \\ \mathbf{l}^T((1 - \mathbf{l}^T\mathbf{L}^{-1}\mathbf{t})\mathbf{I} - \mathbf{L}^{-1}\mathbf{R}\mathbf{L}) & 1 - \mathbf{l}^T\mathbf{L}^{-1}\mathbf{t} \end{bmatrix} \quad (16)$$

which can be factorized as:

$$\mathbf{H}_{12} = \begin{bmatrix} \mathbf{L}^{-1} & \mathbf{0} \\ -\mathbf{l}^T\mathbf{L}^{-1} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{l}^T & 1 \end{bmatrix} \quad (17)$$

We showed that this decomposition exists, but it is certainly not unique. If we count the parameters on each side, $\mathbf{H}_{12}$ has 16 parameters minus 3 because 2 eigenvalues must be 1 and the two others have one degree of freedom (the angle of the rotation, θ), which makes 13 parameters on the left side of equation 9, and on the right side we have 6 parameters for the displacement and 9 for the lower triangular matrix which makes 15 parameters. Then the solution to this equation is not unique and the set of solutions must be a manifold of dimension 2. One of the two remaining parameters is the scale factor on the Euclidean space, which can not be recovered because we have no length reference. We can eliminate it by setting one of the parameters of the diagonal of $\mathbf{L}$ to 1 (they can never be zero because $\mathbf{L}$ is non singular).

It was shown clearly in [3, 8] that the other parameter represents the incertitude on the choice of the absolute conic from $\mathbf{H}$, because one displacement does not define that conic uniquely, so that we cannot recover

the complete Euclidean structure from one displacement (i.e. two projective reconstructions). One way to deal with it would be to fix one of the intrinsic parameters of the cameras[8], e.g. by saying that the x and y axis of the cameras are orthogonal, but since in our scheme the intrinsic parameters do not appear clearly we could not use this. Another one is to simply use more than one displacement, as we demonstrate it in the next section.

5 Euclidean reconstruction of a whole object using stereo by correlation

To test this method, we took several stereoscopic pairs of images of an object using a stereo rig (Figure 2). In this experiment, we used a mathematical object (called “cyclid”) which equation is known, but the fact that we know its geometry was not used in the recovery of its Euclidean geometry. We performed




Figure 2: One of the ten stereoscopic pairs used for the example

weak calibration[1] on these stereo pairs and computed disparity maps using stereo by correlation. We can show that a disparity map computed from a pair of rectified images can be considered as a 3-D projective reconstruction:

Proposition 4 *Let $d(x, y)$ be a disparity map, where x and y are rectified image coordinates. The projective points formed using by the rectified image coordinates as the two first coordinates, the disparity as the third coordinate, and 1 as the last coordinate, form a projective reconstruction, i.e. the 3-D Euclidean coordinates can be recovered by applying a 3-D collineation $\mathbf{H}$ to the points $(x, y, d(x, y), 1)$.*

Let $\mathbf{P}$ and $\mathbf{P}'$ be the projection matrices corresponding respectively to the rectified reference image and the other rectified image. Since the projection of a 3-D projective point M has the same y coordinate in both images, then only the first line of $\mathbf{P}$ and $\mathbf{P}'$ differ:

$$\mathbf{P} = \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix} \quad \text{and} \quad \mathbf{P}' = \begin{bmatrix} p'_1 \\ p_2 \\ p_3 \end{bmatrix} \quad (18)$$

consequently

$$\mathbf{P}M \cong (x, y, 1) \text{ and } \mathbf{P}'M \cong (x + d(x, y), y, 1) \quad (19)$$

so that finally

$$M \cong \mathbf{H}(x, y, d(x, y), 1)^T \quad (20)$$

with

$$\mathbf{H} \cong \begin{bmatrix} p_1 \\ p_2 \\ p'_1 - p_1 \\ p_3 \end{bmatrix}^{-1} \quad (21)$$

Thus the disparity map $(x, y, d(x, y), 1)$ is a projective reconstruction. $\square$

As we have seen before, we have 8 unknowns for the matrix $\mathbf{L}$ and 6 unknowns for each displacement, which makes $6 + 8(n - 1)$ unknowns, if n is the number of stereo pairs. We compute these parameters using a least-squares minimization technique: We match points between the rectified reference images of overlapping stereo pairs¹, and the error to be minimized is the squared distance between the points of reconstruction i transformed by the matrix $\mathbf{L}^{-1}\mathbf{D}_{ij}\mathbf{L}$ and the matched points of reconstruction j . This error measurement is done in (image+disparity) space, which is not the real 3-D Euclidean space, but since image space is almost Euclidean and disparity is bounded, it should work fine. This minimization is done in two steps: First, only the matches between reconstructions i and $i + 1$ are considered, and the minimization is done over $\mathbf{L}$ and the $\mathbf{D}_{i,i+1}$, $1 \leq i < n$, which are represented as a rotation vector and a translation vector. In practice the error function associated with this minimization is well-conditioned, so that we get rather good estimates, whatever the initial point. Second, all the matches between different reconstructions are considered (especially the one between n and 1, which forces the surface to “fold over itself”), and the minimization is done once again.

In fact we recovered the complete Euclidean geometry of our object. Figure 3 shows the reconstruction from the first stereo pair, as seen when transformed by matrix $\mathbf{L}$, and Figures 4 and 5 show the complete reconstruction of the object from 10 stereo pairs, with lighting or with texture mapping.

6 Conclusion

In this paper we presented a method to recover partly or completely the Euclidean geometry using an

¹This process was done manually in our experiment but could be automated.



Figure 3: The Euclidean reconstruction from the first stereo pair



Figure 4: The complete reconstruction of the object, rendered with lighting




Figure 5: The complete reconstruction of the object, rendered with lighting and texture mapping from the original images

uncalibrated stereo rig. All we need to do this is the fundamental matrix of the stereo rig, which can be calculated by a robust method like [1], and point matches between the different stereo pairs, which could be computed automatically. Using multiple stereo pairs, we increase the stability of the algorithm by adding more equations than unknowns. We presented results on a real object, which was fully reconstructed in Euclidean space using a few stereo pairs.

The possible applications of this method include the possibility to acquire easily 3-D objects using any set of uncalibrated stereo cameras, for example to modelize an object to be used in virtual reality, or autonomous robot navigation.

In the near future we plan to enhance the system in order to make it completely automatic: we must have a way to match points automatically (feature tracking would be a good starting point) a to perform fusion and simplification of the 3-D reconstruction once the registration is done. Furthermore, in order to test the accuracy of the Euclidean reconstruction, we can either compare the intrinsic and extrinsic parameters of the cameras computing a classical camera calibration method with those recovered using this method, or compare the 3-D reconstruction with a mathematical model of the object (which is known in the example presented here).

References

- [1] R. Deriche, Z. Zhang, Q.-T. Luong, and O. Faugeras. Robust recovery of the epipolar geometry for an uncalibrated stereo rig. In J.-O. Eklundh, editor, *Proceedings of the 3rd European Conference on Computer Vision*, volume 800-801 of *Lecture Notes in Computer Science*, pages 567–576, Vol. 1, Stockholm, Sweden, May 1994. Springer Verlag.
- [2] Olivier Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig. In G. Sandini, editor, *Proceedings of the 2nd European Conference on Computer Vision*, volume 588 of *Lecture Notes in Computer Science*, pages 563–578, Santa Margherita Ligure, Italy, May 1992. Springer-Verlag.
- [3] Olivier Faugeras. Non-metric representations in 3-D artificial vision. *Nature*, 1993. Submitted.
- [4] Olivier Faugeras. Cartan’s moving frame method and its application to the geometry and evolution of curves in the euclidean, affine and projective planes. In Joseph L. Mundy, Andrew Zisserman, and David Forsyth, editors, *Applications of Invariance in Computer Vision*, volume 825 of *Lecture Notes in Computer Vision*, pages 11–46. Springer-Verlag, 1994.
- [5] Richard Hartley, Rajiv Gupta, and Tom Chang. Stereo from uncalibrated cameras. In *Proceedings of the International Conference on Computer Vision and Pattern Recognition*, pages 761–764, Urbana Champaign, IL, June 1992. IEEE.
- [6] Tuan Luong and Olivier Faugeras. Active stereo with head movements. In *2nd Singapore International Conference on Image Processing*, pages 507–510, Singapore, September 1992.
- [7] S. J. Maybank and O. D. Faugeras. A theory of self-calibration of a moving camera. *The International Journal of Computer Vision*, 8(2):123–152, August 1992.
- [8] Andrew Zisserman, Paul A. Beardsley, and Ian D. Reid. Metric calibration of a stereo rig. In *Proc. Workshop on Visual Scene Representation*, Boston, MA, June 1995.