

A Quasi Linear Reconstruction Method from Multiple Perspective Views

Stéphane Christy, Radu Horaud

► **To cite this version:**

Stéphane Christy, Radu Horaud. A Quasi Linear Reconstruction Method from Multiple Perspective Views. IEEE International Conference on Intelligent Robots and Systems (IROS '95), Aug 1995, Pittsburgh, United States. IEEE Computer Society Press, 3, pp.374–380, 1995, Human Robot Interaction and Cooperative Robots. <10.1109/IROS.1995.526244>. <inria-00590051>

HAL Id: inria-00590051

<https://hal.inria.fr/inria-00590051>

Submitted on 5 May 2011

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.

A Quasi Linear Reconstruction Method from Multiple Perspective Views *

Stéphane Christy and Radu Horaud
LIFIA & INRIA Rhône-Alpes
46, avenue Félix Viallet 38031 Grenoble FRANCE

Abstract

In this paper we describe a method for solving the Euclidean reconstruction problem with a perspective camera model by incrementally performing an Euclidean reconstruction with a weak perspective camera model. With respect to other methods that compute shape and motion from a sequence of images with a calibrated perspective camera, this method converges in a few iterations, is computationally efficient, and does not suffer from the non linear nature of the problem. With respect to factorization and/or affine-invariant methods, this method solves for the sign (reversal) ambiguity in a very simple way and provides much more accurate reconstructions results.

1 Introduction and background

The problem of computing 3-D shape and motion from a long sequence of images has received a lot of attention for the last few years. Previous approaches attempting to solve this problem fall into several categories, whether the camera is calibrated or not, and/or whether a projective or an affine model is being used. With a calibrated camera one may compute Euclidean shape up to a scale factor using either a perspective model [8], or a linear model [9], [10], [6], [7]. With an uncalibrated camera the recovered shape is defined up to a projective transformation [2], or up to an affine transformation. One can therefore address the problem of either Euclidean, affine, or projective shape reconstruction. In this paper we are interested in Euclidean shape reconstruction with a calibrated camera.

The perspective model has associated with it, in general, non linear reconstruction techniques. This naturally leads to non-linear minimization methods which require some form of initialization [8], [2]. If

the initial “guess” is too faraway from the true solution then the minimization process is either very slow or it converges to a wrong solution. Affine models lead, in general, to linear resolution methods [9], [10], [6], [7], but the solution is defined only up to a sign (reversal) ambiguity. Moreover, both these two affine solutions are just approximations of the true solution.




Figure 1: This figure shows a theoretical 3-D shape (top-left) and three reconstructions of this shape from 10 views. The first reconstruction (top-right) was obtained using a perspective camera model and the method described in this paper. The second reconstruction (bottom-left) and its reversal (bottom-right) were obtained using a weak perspective camera model and the factorization method of Tomasi & Kanade.

One way to combine the perspective and affine models could be to use the linear (affine) solution in

*This work has been supported by “Société Aérospatiale” and by DRET.

order to initialize the non-linear minimization process associated with perspective. However, there are several drawbacks with such an approach. First, such a resolution technique does not take into account the simple link that exists between the perspective model and its linear approximations. Second, there is no mathematical evidence that a non-linear least-squares minimization method is “well” initialized by a solution that is obtained linearly. Third, there are two solutions associated with the affine model and it is not clear which one to choose.

The perspective projection can be modelled by a projective transformation from the 3-D projective space to the 2-D projective plane. Weak perspective and paraperspective are the most common affine approximations of perspective. Weak perspective may well be viewed as a zero-order approximation: $1/(1 + \varepsilon) \approx 1$. Paraperspective [1] is a first order approximation of full perspective: $1/(1 + \varepsilon) \approx 1 - \varepsilon$. Recently, in [4] a method has been proposed for determining the pose of a 3-D shape with respect to a single view by iteratively improving the pose computed with a weak perspective camera model to converge, at the limit, to a pose estimation using a perspective camera model. At our knowledge, the method cited above, i.e., [4] is among one of the first computational paradigms that link linear techniques (associated with affine camera models) with a perspective model. In [5] an extension of this paradigm to paraperspective is proposed. The authors show that the *iterative paraperspective* pose algorithm has better convergence properties than the *iterative weak perspective* one.

In this paper we describe a new Euclidean reconstruction method that makes use of affine reconstruction in an iterative manner such that this iterative process converges, at the limit, to a set of 3-D Euclidean shape and motion parameters that are consistent with a perspective model. The novelty of the method that we propose is twofold: (i) it extends the iterative pose determination algorithms described in [4] and in [5] to deal with the problem of shape and motion from multiple views and (ii) it is a generalization to perspective of the factorization methods [9], [7] and of the affine-invariant methods [10]. More precisely, the *affine-iterative reconstruction* method that we propose here has a number of interesting features:

- It solves the sign (or reversal) ambiguity that is inherent with affine reconstruction;
- It is fast because it converges in a few iterations (3 to 5 iterations), each iteration involving simple linear algebra computations;

- The quality of the Euclidean reconstruction obtained with our method is only weakly influenced by camera calibration errors. The only intrinsic camera parameter that has a crucial effect on the quality of the reconstruction is the ratio between the horizontal pixel size and vertical pixel size – ratio which is known to be very stable;
- It allows the use of either weak (in this paper) or paraperspective (see [3]) camera model approximations which are used iteratively, and
- It can be combined with almost any affine shape and motion algorithm. In particular we show how our method can be combined either with factorization methods [9], [7] or with affine-invariant methods [10], [6].

2 Camera models

Let us consider a pin hole camera model. We denote by P_i a 3-D point with Euclidean coordinates X_i , Y_i , and Z_i in a frame that is attached to the object – the object frame. The origin of this frame may well be the object point P_0 . The relationship between object points and projected points can be written as:

$$x_i = \frac{\mathbf{i} \cdot \mathbf{P}_i + t_x}{\mathbf{k} \cdot \mathbf{P}_i + t_z} \quad (1)$$

$$y_i = \frac{\mathbf{j} \cdot \mathbf{P}_i + t_y}{\mathbf{k} \cdot \mathbf{P}_i + t_z} \quad (2)$$

Recall the relationship between camera coordinates and image coordinates:

$$u_i = \alpha_u x_i + u_c \quad (3)$$

$$v_i = \alpha_v y_i + v_c \quad (4)$$

In these equations α_u and α_v are the vertical and horizontal scale factors and u_c and v_c are the image coordinates of the intersection of the optical axis with the image plane. In an extended version of this paper [3], we show that the reconstruction method described here depends only on the ratio α_u/α_v and that the reconstruction obtained with our method is not sensitive to errors in u_c and v_c .

We divide both the numerator and the denominator of eqs. (1) and (2) by t_z . We introduce the following notations:

- $\mathbf{I} = \mathbf{i}/t_z$ and $\mathbf{J} = \mathbf{j}/t_z$;
- $x_0 = t_x/t_z$ and $y_0 = t_y/t_z$ are the camera coordinates of p_0 which is the projection of P_0 – the origin of the object frame, and

- We denote by ε_i the following ratio:

$$\varepsilon_i = \frac{\mathbf{k} \cdot \mathbf{P}_i}{t_z} \quad (5)$$

We may now rewrite the perspective equations as:

$$x_i = \frac{\mathbf{I} \cdot \mathbf{P}_i + x_0}{1 + \varepsilon_i} \quad (6)$$

$$y_i = \frac{\mathbf{J} \cdot \mathbf{P}_i + y_0}{1 + \varepsilon_i} \quad (7)$$

Whenever the object is at some distance from the camera, the ε_i are small compared to 1. We may therefore introduce the weak perspective model as an approximation of the perspective equations.

Weak perspective assumes that the object points lie in a plane parallel to the image plane passing through the origin of the object frame, i.e., P_0 . This is equivalent to a zero-order approximation:

$$\frac{1}{1 + \varepsilon_i} \approx 1 \quad \forall i, i \in \{1 \dots n\}$$

With this approximation, eqs. (6) and (7) become:

$$x_i^w - x_0 = \mathbf{I} \cdot \mathbf{P}_i \quad (8)$$

$$y_i^w - y_0 = \mathbf{J} \cdot \mathbf{P}_i \quad (9)$$

In these two equations x_i^w and y_i^w are the camera coordinates of the weak perspective projection of the point P_i . By identification with eqs. (6) and (7) we obtain the relationship between the weak perspective and the perspective projections of P_i :

$$x_i^w = x_i(1 + \varepsilon_i) \quad (10)$$

$$y_i^w = y_i(1 + \varepsilon_i) \quad (11)$$

These equations allow us to determine the quality of the weak perspective approximation with respect to the perspective projection. Indeed the error between the weak perspective projection and the “true” projection is:

$$\Delta x^w = |x_i^w - x_i| = |x_i \varepsilon_i| \quad (12)$$

$$\Delta y^w = |y_i^w - y_i| = |y_i \varepsilon_i| \quad (13)$$

Hence the quality of the approximation depends both on the value of ε_i AND on the position of the point in the image. Let’s consider for example a 512×512 image. Approximate values for the intrinsic parameters are: $\alpha_u = \alpha_v = 1000$ and $u_c = v_c = 256$. Therefore using eq. (3) and (4) we have:

$$0 \leq x_i, y_i \leq 0.25$$

We conclude that for objects that are quite closed to the camera, the weak perspective approximation is still valid provided that the object lies in the neighbourhood of the optical axis.

3 Reconstruction with a perspective camera

Let us consider again the perspective equations (6) and (7). These equations may also be written as:

$$x_i(1 + \varepsilon_i) - x_0 = \mathbf{I} \cdot \mathbf{P}_i \quad (14)$$

$$y_i(1 + \varepsilon_i) - y_0 = \mathbf{J} \cdot \mathbf{P}_i \quad (15)$$

If the values of ε_i are set to zero then we obtain equations that approximate the camera with weak perspective. The crucial point is that if the ε_i ’s are set to some fixed non zero values, then the equations remain linear. In particular, there are values for the ε_i ’s for which these equations are consistent – up to some measurement noise – with the full perspective model. The key idea of our reconstruction method is to iteratively estimate values for the ε_i ’s such that the perspective reconstruction problem becomes a “weak perspective iterative” reconstruction problem.

Let us consider now k views of the same scene points. We assume that image-to-image correspondences have already been established. Equations (14) and (15) can be written as:

$$\underbrace{\mathbf{s}_{ij}}_{2 \times 1} = \underbrace{A_j}_{2 \times 3} \underbrace{\mathbf{P}_i}_{3 \times 1} \quad (16)$$

In this formula the subscript i stands for the i^{th} point and the subscript j for the j^{th} image. The 2-vector \mathbf{s}_{ij} is equal to:

$$\mathbf{s}_{ij} = \begin{pmatrix} x_{ij}(1 + \varepsilon_{ij}) - x_{0j} \\ y_{ij}(1 + \varepsilon_{ij}) - y_{0j} \end{pmatrix} \quad (17)$$

In these equations ε_{ij} , i.e., eq. (5) is defined for each point and for each image:

$$\varepsilon_{ij} = \frac{\mathbf{k}_j \cdot \mathbf{P}_i}{t_{z_j}} \quad (18)$$

The reconstruction problem is now the problem of simultaneously solving $2 \times n \times k$ equations of the form of eq. (16). We introduce now a method that solves these equations by linear iterations. More precisely, this method can be summarized by the following algorithm:

1. $\forall i, i \in \{1 \dots n\}$ and $\forall j, j \in \{1 \dots k\}$ set: $\varepsilon_{ij} = 0$;
2. Update the values of \mathbf{s}_{ij} using ε_{ij} ;
3. Perform an Euclidean reconstruction with a weak perspective camera;

4. $\forall i, i \in \{1 \dots n\}$ and $\forall j, j \in \{1 \dots k\}$ estimate new values for ε_{ij} ;
5. Check the values of ε_{ij} :
 - if $\forall(i, j)$ the values of ε_{ij} just estimated at this iteration are identical with the values estimated at the previous iteration, then stop;
 - else go to step 2.

Step 3 of the algorithm can be farther decomposed into: (i) affine reconstruction and (ii) euclidean reconstruction. The problem of affine reconstruction is the problem of determining both A_j and \mathbf{P}_i , for all j and for all i , in eq. (16), when some estimates of \mathbf{s}_{ij} are provided. Such a reconstruction determines shape and motion up to a 3-D affine transformation. Indeed, for any 3×3 invertible matrix T we have:

$$A_j \mathbf{P}_i = A_j T T^{-1} \mathbf{P}_i$$

In order to convert affine shape and motion into Euclidean shape and motion, one needs to consider some Euclidean constraints associated either with the motion of the camera or with the shape being viewed by the camera. Since we deal here with a calibrated camera, we may well use rigid motion constraints in conjunction with weak perspective [7]. Based on the parameters of the Euclidean shape and motion thus computed one can estimate ε_{ij} for all i and for all j using eq. (18) – step 4.

The first iteration of the algorithm performs a 3-D reconstruction using the initial image measurements and a weak perspective camera model. This first reconstruction allows an estimation of values for the ε_{ij} 's which in turn allow the image vectors \mathbf{s}_{ij} to be *modified* (step 2 of the algorithm). The \mathbf{s}_{ij} 's are modified according to eq. (17) such that they better fit the approximated camera model being used.

The next iterations of the algorithm perform a 3-D reconstruction using (i) image vectors that are incrementally modified and (ii) a weak perspective camera model.

At convergence, the equations (16) are equivalent with the perspective equations (14) and (15). In other terms, this algorithm solves for Euclidean reconstruction with a perspective camera by iterations of an Euclidean reconstruction method with a weak perspective camera – a quasi linear algorithm.

The iterative algorithm outlined in this paper is best illustrated as Figure 2. At the first iteration, the algorithm considers the *true* perspective projections of P_i and attempts to reconstruct the 3-D points as if

they were projected in the image using weak perspective. At the second iteration the algorithm considers modified image point positions. At the last iteration, the image point positions were modified such that they fit the weak perspective projections. The relative positions of the perspective projections and of the weak perspective projections verify the theoretical relationship given by eqs. (10) and (11). Notice that the perspective projection of point P_0 is identical to its weak perspective projection and hence the projection of this point is not modified.




Figure 2: The iterative algorithm described in this section modifies the projection of a 3-D point from true perspective to weak perspective (see text).

4 Reconstruction with a weak perspective camera

In this section we develop step 3 of the algorithm outlined in the previous section. Methods that use a linear camera model provide a 3-D affine reconstruction if at least 2 views of 4 non-coplanar points are available and if the motion is not a pure translation. However, 3 views are necessary in order to convert this affine reconstruction into an Euclidean one. While the affine-invariant method allows a more direct analysis of the problem, the factorization method is more convenient from a practical point of view.

Affine-invariant methods have already been described in [6] and [10]. The factorization method is due to [9]. Tomasi & Kanade [9] noticed that *affine shape* and *motion* may be computed simultaneously by performing a singular value decomposition of the $2k \times n$ matrix σ which is formed by concatenation of \mathbf{s}_{ij} (equation (16)):

$$\sigma = AS = O_1 \Sigma O_2 \quad (19)$$

The rank of σ is theoretically equal to 3. Tomasi & Kanade suggested to solve the rank problem by truncating the matrix Σ such that only the 3 largest diagonal values are considered. They claim that this truncation amounts to removing noise present in the measurement matrix. Therefore, one can write the singular value decomposition of the measurement matrix as:

$$\sigma = O_1' \Sigma' O_2' + O_1'' \Sigma'' O_2''$$

where Σ' is a 3×3 diagonal matrix containing the 3 largest diagonal terms of Σ .

Finally, affine shape and affine motion may be given by:

$$S = (\Sigma')^{1/2} O_2' \quad \text{and} \quad A = O_1' (\Sigma')^{1/2}$$

5 Euclidean constraints

One has to determine now Euclidean shape and motion by combining the affine reconstruction methods just described and the Euclidean constraints available with the camera model being used. As already mentioned, one has to determine a 3×3 invertible matrix T such that the affine shape S becomes Euclidean:

$$(\mathbf{P}_1 \quad \dots \quad \mathbf{P}_n) = T^{-1} (\mathbf{S}_1 \quad \dots \quad \mathbf{S}_n)$$

and the affine motion becomes rigid:

$$\begin{pmatrix} A_1^E \\ \vdots \\ A_k^E \end{pmatrix} = \begin{pmatrix} A_1 \\ \vdots \\ A_k \end{pmatrix} T$$

Recall that in the case of weak perspective the 2×3 matrices A_j^E are of the form:

$$A_j^E = \begin{pmatrix} \mathbf{I}_j \\ \mathbf{J}_j \end{pmatrix}$$

with: $\mathbf{I}_j = \mathbf{i}_j/t_{z_j}$ and $\mathbf{J}_j = \mathbf{j}_j/t_{z_j}$. Therefore for all j , the row vectors of the matrices A_j^E verify:

$$\|\mathbf{I}_j\| = \|\mathbf{J}_j\| \quad \text{and} \quad \mathbf{I}_j \cdot \mathbf{J}_j = 0$$

These constraints are homogeneous and non linear in the coefficients of T . In order to avoid the trivial null solution the scale factor must be fixed in advance. For example, one may choose:

$$t_{z_1} = 1$$

With the substitution $Q = TT^T$, we obtain $2k + 1$ linear independent constraints in the coefficients of Q

(6 unknowns), which can be solved if at least 3 views are available. Once Q has been linearly estimated, T can be easily computed. Next we determine the parameters of the Euclidean motion:

$$\begin{aligned} \mathbf{i}_j &= \mathbf{I}_j / \|\mathbf{I}_j\| & t_{z_j} &= 1/2 (1/\|\mathbf{I}_j\| + 1/\|\mathbf{J}_j\|) \\ \mathbf{j}_j &= \mathbf{J}_j / \|\mathbf{J}_j\| & t_{x_j} &= x_{0_j} t_{z_j} \\ \mathbf{k}_j &= \mathbf{i}_j \times \mathbf{j}_j & t_{y_j} &= y_{0_j} t_{z_j} \end{aligned}$$

6 Solving the reversal ambiguity

The algorithm outlined in Section 3 solves for Euclidean reconstruction with a perspective camera by iterations of an Euclidean reconstruction method with a weak perspective camera. In this section we show how this iterative algorithm has to be modified in order to solve the reversal ambiguity problem which is inherent with any affine camera model. Indeed, let us consider again the affine shape and motion equation, i.e., eq. (19). This equation is defined up to a sign:

$$\sigma = AS = (-A)(-S)$$

Moreover, the use of Euclidean constraints does not solve this ambiguity:

$$\sigma = A^E P = (-A^E)(-P)$$

Because of this sign ambiguity that is associated with the affine model, there are 2 solutions for the ε_{ij} 's at each iteration of the perspective reconstruction algorithm.

The first solution yields:

$$\varepsilon_{ij}^1 = \mathbf{k}_j \cdot \mathbf{P}_i / t_{z_j}$$

The second solution is different because the vector \mathbf{k} remains unchanged when the signs of the row vectors of A^E are changed:

$$(-\mathbf{i}_j) \times (-\mathbf{j}_j) = \mathbf{i}_j \times \mathbf{j}_j = \mathbf{k}_j$$

Therefore we obtain:

$$\varepsilon_{ij}^2 = \mathbf{k}_j \cdot (-\mathbf{P}_i) / t_{z_j} = -\varepsilon_{ij}^1$$

At each iteration of the perspective reconstruction algorithm two values for ε_{ij} are thus estimated. Therefore, after N iterations there will be 2^N possible solutions. All these solutions are not, however, necessarily consistent with the image data and a simple verification technique allows to check this consistency and to avoid the explosion of the number of solutions. Finally, a unique solution is obtained.




Figure 3: A strategy for selecting a unique solution (see text).

Let $S^{(1)}$ be the positive shape computed at the first iteration of the algorithm and $R^{(1)}$ be the negative shape ($R^{(1)} = -S^{(1)}$). At each one of the next iterations one has to deal with four shapes: $S_1^{(i)}$ and $S_2^{(i)}$ that are issued from the positive solution and $R_1^{(i)}$ and $R_2^{(i)}$ that are issued from the negative solution. The S -shape and the R -shape the most consistent with the shapes selected at the previous iterations are selected. Finally, a unique solution is selected on the basis of consistency with the image data. This solution selection process is best illustrated on Figure 3.

7 Experimental results

In this section we describe two types of experiments: (i) experiments with synthetic data which allow us to study both the accuracy of the 3-D reconstruction and the convergence of the iterative algorithm, and (ii) experiments with real data.

Let us consider again the synthetic house of Figure 1. We designate by D the distance between this object center and the camera center of projection divided by the object size – D is therefore a relative distance. For a fixed value of D we consider 10 camera motions, each motion being composed of 15 images. Each such motion is further characterized by a translation vector and a rotation axis and angle. The directions of the translation vector and rotation axis are randomly chosen. The angle of rotation between two images is equal to 2° . Moreover, the image data ob-

tained by projecting this object onto the image plane is perturbed by adding gaussian noise with a standard deviation equal to 1.

The accuracy of the reconstruction is measured by the difference between the theoretical 3-D points of the object and the reconstructed 3-D points. Figures 4 summarize the results where D is allowed to vary from 3 to 19 and for two methods: the factorization method with a weak perspective camera (small squares) and our reconstruction method (small triangles). One may easily notice that the precision associated with these two algorithms converge to the same value as D increases. However, the discrepancy between the precision associated with the two methods increases as D decreases.




Figure 4: These plots show the mean values of the distance between object points and reconstructed points as a function of D . The behaviour of the factorization method (small squares) is compared with the behaviour of the quasi linear method described in this paper (small triangles).

We consider now an experiment performed with 5 images of a “cube”. The first image of this sequence is shown on Figure 5 (top-left). There are 38 feature points that are tracked over the image sequence. The camera parameters have been fixed to the following values: $\alpha_u/\alpha_v = 1.5$ and $u_0 = v_0 = 256$.

The bottom of Figure 5 shows the top views of the reconstructed cube obtained with the factorization method (left) and with the method described in this paper (right).

8 Discussion

In this paper we described a method for solving the Euclidean reconstruction problem with a perspective




Figure 5: This figure shows one image (top-left) out of a sequence of 5 images grabbed with a moving camera, the result of reconstructing the scene with the quasi linear method (top-right and bottom-right), and the result of reconstructing the scene with the factorization method (bottom-left).

camera by incrementally performing an Euclidean reconstruction with a weak perspective camera model. The method converges, on an average, in 5 iterations, is computationally efficient, and it produces accurate results even in the presence of image noise and/or camera calibration errors. The method may well be viewed as a generalization to perspective of shape and motion computation using factorization and/or affine-invariant methods. It is well known that with a linear camera model, shape and motion can be recovered only up to a sign (reversal) ambiguity. The method that we propose in this paper solves for this ambiguity and produces a unique solution even if the camera is at some distance from the scene.

Although the experimental results show that there are little convergence problems, we have been unable to study the convergence of the algorithm from a theoretical point of view. We studied its convergence based on some numerical and practical considerations which allow us to determine in advance the optimal experimental setup under which convergence can be guaranteed [3]. In the future we plan to study more

thoroughly the convergence of this type of algorithms.

We also plan to generalize our method such that it can deal with an uncalibrated camera. In that case one has to investigate ways to incrementally compute shape and motion without any explicit Euclidean constraints. The result of such a method could be, for instance, an affine reconstruction with a projective camera model.

References

- [1] Y. Aloimonos. Perspective approximations. *Image and Vision Computing*, 8(3):177–192, August 1990.
- [2] B. Boufama, R. Mohr, and F. Veillon. Euclidean constraints for uncalibrated reconstruction. In *Proceedings Fourth International Conference on Computer Vision*, pages 466–470, Berlin, Germany, May 1993. IEEE Computer Society Press, Los Alamitos, Ca.
- [3] S. Christy and R. Horaud. Euclidean shape and motion from multiple perspective views by affine iterations. Technical Report RR-2421, INRIA, December 1994. Available by anonymous ftp on imag.fr, directory pub/MOVI, file RR-2421.ps.gz.
- [4] D. F. Dementhon and L. S. Davis. Model-based object pose in 25 lines of code. *International Journal of Computer Vision*, 1995. In press.
- [5] R. Horaud, S. Christy, F. Dornaika, and B. Lamiroy. Object pose: Links between paraperspective and perspective. In *Proceedings Fifth International Conference on Computer Vision*, Cambridge, Mass., June 1995. IEEE Computer Society Press, Los Alamitos, Ca.
- [6] J. Koenderink and A. van Doorn. Affine structure from motion. *J. Opt. Soc. Amer. A*, 8(2):377–385, 1991.
- [7] C. J. Poelman and T. Kanade. A paraperspective factorization method for shape and motion recovery. In Jan-Olof Eklundh, editor, *Computer Vision – ECCV 94, Proceedings Third European Conference on Computer Vision, Stockholm, Sweden, May 1994*, volume 2, pages 97–108. Springer Verlag, May 1994.
- [8] C. J. Taylor, D. J. Kriegman, and P. Anandan. Structure and motion in two dimensions from multiple images: A least squares approach. In *IEEE Workshop on Visual Motion*, pages 242–248, Princeton, New Jersey, USA, October 1991.
- [9] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: a factorization method. *International Journal of Computer Vision*, 9(2):137–154, November 1992.
- [10] D. Weinshall. Model-based invariants for 3-d vision. *International Journal of Computer Vision*, 10(1):27–42, February 1993.