**Shape and Motion from Image Streams:**
**a Factorization Method**
**Full Report on the Orthographic Case***

Carlo Tomasi
Takeo Kanade

TR 92-1270
March 1992

Department of Computer Science
Cornell University
Ithaca, NY 14853-7501

*This research was sponsored by the Avionics Laboratory, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U.S. Air Force, Wright-Patterson AFB, Ohio 45433-6543 under Contract F33615-90-C-1465, ARPA Order No. 7597.

# Shape and Motion from Image Streams: a Factorization Method

## Full Report on the Orthographic Case

Carlo Tomasi Takeo Kanade

March 1992

ALSO APPEARS AS CARNEGIE MELLON REPORT CMU-CS-92-104

This research was sponsored by the Avionics Laboratory, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U.S. Air Force, Wright-Patterson AFB, Ohio 45433-6543 under Contract F33615-90-C-1465, ARPA Order No. 7597.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. government.

**Keywords:** computer vision, motion, shape, time-varying imagery

## Abstract

Inferring scene geometry and camera motion from a stream of images is possible in principle, but is an ill-conditioned problem when the objects are distant with respect to their size. We have developed a *factorization method* that can overcome this difficulty by recovering shape and motion without computing depth as an intermediate step.

An image stream can be represented by the $2F \times P$ measurement matrix of the image coordinates of $P$ points tracked through $F$ frames. We show that under orthographic projection this matrix is of rank 3.

Using this observation, the factorization method uses the singular value decomposition technique to factor the measurement matrix into two matrices which represent object shape and camera motion respectively. The method can also handle and obtain a full solution from a partially filled-in measurement matrix, which occurs when features appear and disappear in the image sequence due to occlusions or tracking failures.

The method gives accurate results, and does not introduce smoothing in either shape or motion. We demonstrate this with a series of experiments on laboratory and outdoor image streams, with and without occlusions.

# Preface

In principle, the stream of images produced by a moving camera allows the recovery of both the shape of the objects in the field of view, and the motion of the camera. Traditional algorithms recover depth by triangulation, and compute shape by taking differences between depth values. This process, however, is sensitive to noise for distant scenes.

To overcome this problem, we have developed a factorization method to decompose an image stream directly into object shape and camera motion, without computing depth as an intermediate step. To explore this new method, we designed a series of eleven technical reports, as shown in figure 1, going from basic theory to implementation.

This report summarizes our work for the orthographic case, including the basic mathematical formulation, the treatment of occlusions, and the discussion of implementation issues (parts 2, 8, and 10 in the diagram below).

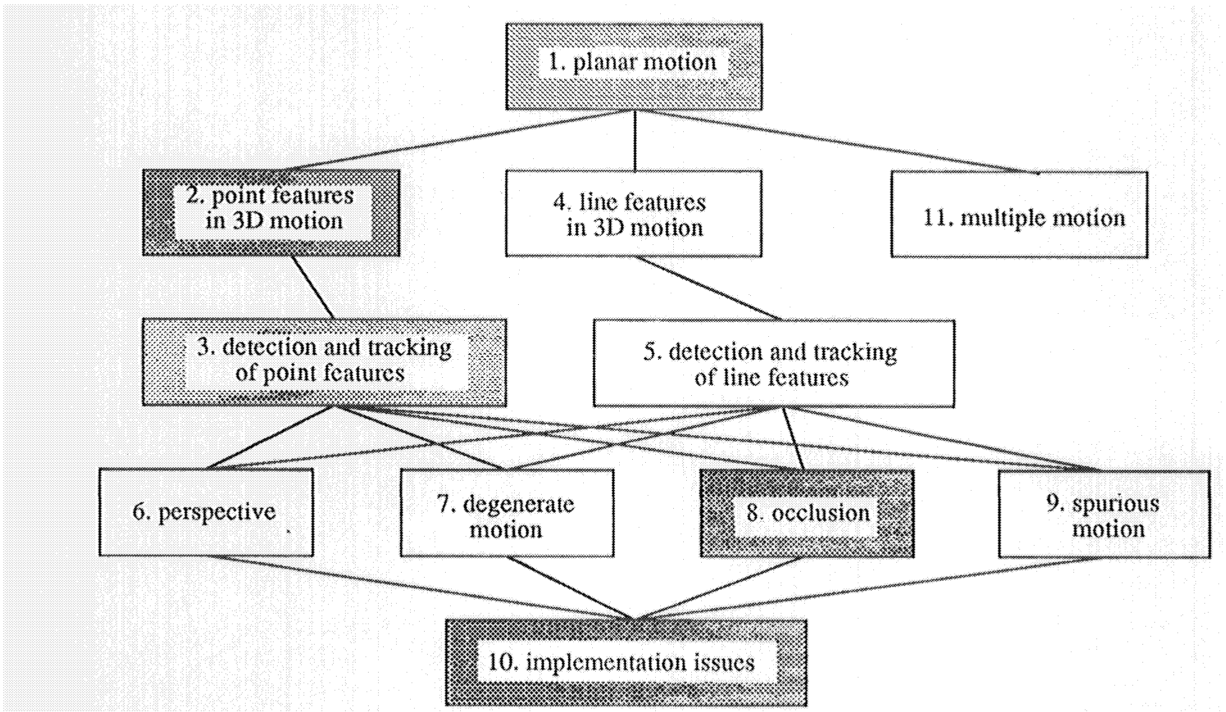


The material covered in the shaded blocks in the diagram has appeared as follows:

1 CMU-CS-90-166, DARPA IUW '90, and IJCV '90

2 CMU-CS-91-105, DARPA IUW '91, and IEEE Workshop on Visual Motion '91

3 CMU-CS-91-132

8 This report and DARPA IUW '91

10 This report and DARPA IUW '91

# Chapter 1

# Introduction

The structure from motion problem – recovering scene geometry and camera motion from a sequence of images – has attracted much of the attention of the vision community over the last decade. Yet it is common knowledge that existing solutions work well for perfect images, but are very sensitive to noise. We present a new method called the factorization method which can robustly recover shape and motion from a sequence of images without assuming a model of motion, such as constant translation or rotation.

More specifically, an image sequence can be represented as a $2F \times P$ measurement matrix $W$, which is made up of the horizontal and vertical coordinates of $P$ points tracked through $F$ frames. If image coordinates are measured with respect to their centroid, we prove the *rank theorem*: under orthography, the measurement matrix is of rank 3. As a consequence of this theorem, we show that the measurement matrix can be factored into the product of two matrices $R$ and $S$. Here, $R$ is a $2F \times 3$ matrix that represents camera rotation, and $S$ is a $3 \times P$ matrix which represents shape in a coordinate system attached to the object centroid. The two components of the camera translation along the image plane are computed as averages of the rows of $W$. When features appear and disappear in the image sequence due to occlusions or tracking failures, the resultant measurement matrix $W$ is only partially filled-in. The factorization method can handle this situation by growing a partial solution obtained from an initial full submatrix into a full solution with an iterative procedure.

The rank theorem precisely captures the nature of the redundancy that exists in an image sequence, and permits a large number of points and frames to be processed in a conceptually simple and computationally efficient way to reduce the effects of noise. The resulting algorithm is based on the singular value decomposition, which is numerically well-behaved and stable. The robustness of the recovery algorithm in turn enables us to use an image sequence with a very short interval between frames (an *image stream*), which makes feature tracking relatively easy.

We have demonstrated the accuracy and robustness of the factorization method in a series of experiments on laboratory and outdoor sequences, with and without occlusions.

# Chapter 2

# Relation to Previous Work

In Ullman's original proof of existence of a solution [Ullman, 1979] for the structure from motion problem under orthography, as well as in the perspective formulation in [Roach and Aggarwal, 1979], the coordinates of feature points in the world are expressed in a world-centered system of reference. Since then, however, this choice has been replaced by most computer vision researchers with that of a camera-centered representation of shape [Prazdny, 1980], [Bruss and Horn, 1983], [Tsai and Huang, 1984], [Adiv, 1985], [Waxman and Wohn, 1985], [Bolles *et al.*, 1987], [Horn *et al.*, 1988], [Heeger and Jepson, 1989], [Heel, 1989], [Matthies *et al.*, 1989], [Spetsakis and Aloimonos, 1989], [Broida *et al.*, 1990]. With this representation, the position of feature points is specified by their image coordinates and by their depths, defined as the distances between the camera center and the feature points, measured along the optical axis. Unfortunately, although a camera-centered representation simplifies the equations for perspective projection, it makes shape estimation difficult, unstable, and noise sensitive.

There are two fundamental reasons for this. First, when camera motion is small, effects of camera rotation and translation can be confused with each other: for example, small rotation about the vertical axis and small translation along the horizontal axis both generate a very similar change in an image. Any attempt to recover or differentiate between these two motions, though doable mathematically, is naturally noise sensitive. Second, the computation of shape as relative depth, for example, the height of a building as the difference of depths between the top and the bottom, is very sensitive to noise, since it is a small difference between large values.

These difficulties are especially magnified when the objects are distant from the camera relative to their sizes, which is usually the case for interesting applications such as site modeling.

The factorization method we present in this paper takes advantage of the fact that both difficulties disappear when the problem is reformulated in world-centered coordinates, unlike the conventional camera-centered formulation. This new (old – in a sense) formulation links object-centered shape to image motion directly, without using retinotopic depth as an intermediate quantity, and leads to a simple and well-behaved solution. Furthermore, the mutual independence of shape and motion in world-centered coordinates makes it possible to cast the structure-from-motion problem as a factorization problem, in which a matrix representing image measurements is decomposed directly into camera motion and object shape.

We first introduced this factorization method in [Tomasi and Kanade, 1990a, Tomasi and Kanade, 1990b], where we treated the case of single-scanline images in a flat, two-dimensional world. In [Tomasi and Kanade, 1991] we presented the theory for the case of arbitrary camera motion in three dimensions and full two-dimensional images. This paper extends the factorization method for dealing with feature occlusions as well as presenting more experimental results with real-world images. Debrunner and Ahuja have pursued an approach related to ours, but using a different formalism [Debrunner and Ahuja, 1990, Debrunner and Ahuja, 1991]. Assuming that motion is constant over a period, they provide both closed-form expressions for shape and motion and an incremental solution (one image at a time) for multiple motions by taking advantage of the redundancy of measurements. Boult and Brown have investigated the factorization method for multiple motions [Boult and Brown, 1991], in which they count and segment separate motions in the field of view of the camera.

# Chapter 3

# The Factorization Method

Given an image stream, suppose that we have tracked $P$ feature points over $F$ frames. We then obtain trajectories of image coordinates $\{(u_{fp}, v_{fp}) \mid f = 1, \ldots, F, p = 1, \ldots, P\}$. We write the horizontal feature coordinates $u_{fp}$ into an $F \times P$ matrix $U$: we use one row per frame, and one column per feature point. Similarly, an $F \times P$ matrix $V$ is built from the vertical coordinates $v_{fp}$. The combined matrix of size $2F \times P$

$$W = \left[\frac{U}{\mathrm{V}}\right]$$

is called the *measurement matrix*. The rows of the matrices $U$ and $V$ are then registered by subtracting from each entry the mean of the entries in the same row:

$$\begin{aligned} \tilde{u}_{fp} &= u_{fp} - a_f \\ \tilde{v}_{fp} &= v_{fp} - b_f \,, \end{aligned} \tag{3.1}$$

where

$$\begin{aligned} a_f &= \frac{1}{P}\sum_{p=1}^{P} u_{fp} \\ b_f &= \frac{1}{P}\sum_{p=1}^{P} v_{fp} \,. \end{aligned}$$

This produces two new $F \times P$ matrices $\tilde{U} = [\tilde{u}_{fp}]$ and $\tilde{V} = [\tilde{v}_{fp}]$. The matrix

$$\widetilde{W} = \left[\frac{\tilde{U}}{\tilde{V}}\right]$$

is called the *registered measurement matrix.* This is the input to our factorization method.

## 3.1 The Rank Theorem

We now analyze the relation between camera motion, shape, and the entries of the registered measurement matrix $\widetilde{W}$. This analysis leads to the key result that $\widetilde{W}$ is highly rank-deficient.

Referring to Figure 3.1, suppose we place the origin of the world reference system $x - y - z$ at the centroid of the $P$ points $\mathbf{s}_p = (x_p, y_p, z_p)^T, p = 1, \ldots, P\}$, in space which correspond to the $P$ feature points tracked in the image stream. The orientation of the camera reference system corresponding to frame number $f$ is determined by a pair of unit vectors, $\mathbf{i}_f$and $\mathbf{j}_f$, pointing along the scanlines and the columns of the image respectively, and defined with respect to the world reference system. Under orthography, all projection rays are then parallel to the cross product of $\mathbf{i}_f$and $\mathbf{j}_f$:

$$\mathbf{k}_f = \mathbf{i}_f \times \mathbf{j}_f \ .$$

From Figure 3.1 we see that the projection $(u_{fp}, v_{fp})$, *i.e.*, the image feature position, of point $\mathbf{s}_p = (x_p, y_p, z_p)^T$ onto frame $f$ is given by the equations

$$\begin{aligned} u_{fp} &= \mathbf{i}_f{}^T(\mathbf{s}_p - \mathbf{t}_f) \\ v_{fp} &= \mathbf{j}_f{}^T(\mathbf{s}_p - \mathbf{t}_f) \ , \end{aligned}$$

where $\mathbf{t}_f = (a_f, b_f, c_f)^T$ is the vector from the world origin to the origin of image frame $f$. Here note that since the origin of the world coordinates is placed at the centroid of object points,

$$\frac{1}{P} \sum_{p=1}^{P} \mathbf{s}_p = 0 \ .$$

We can now write expressions for the entries $\tilde{u}_{fp}$ and $\tilde{v}_{fp}$ defined in (3.1) of the registered measurement matrix. For the the registered horizontal image projection we have

$$\tilde{u}_{fp} \quad = \quad u_{fp} - a_f$$

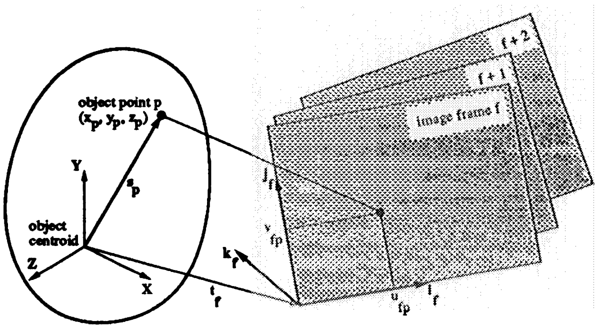


Figure 3.1: The systems of reference used in our problem formulation.

$$
\begin{aligned}
&= \mathbf{i}_f{}^T(\mathbf{s}_p - \mathbf{t}_f) - \frac{1}{P}\sum_{q=1}^{P}\mathbf{i}_f{}^T(\mathbf{s}_q - \mathbf{t}_f) \\
&= \mathbf{i}_f{}^T\left(\mathbf{s}_p - \frac{1}{P}\sum_{q=1}^{P}\mathbf{s}_q\right) \\
&= \mathbf{i}_f{}^T\mathbf{s}_p \ . \qquad (3.2)
\end{aligned}
$$

We can write a similar equation for $\tilde{v}_{fp}$. To summarize,

$$
\begin{aligned}
\tilde{u}_{fp} &= \mathbf{i}_f{}^T\mathbf{s}_p \\
\tilde{v}_{fp} &= \mathbf{j}_f{}^T\mathbf{s}_p \ .
\end{aligned} \qquad (3.3)
$$

Because of the two sets of $F \times P$ equations (3.3), the registered measurement matrix $\widetilde{W}$ can be expressed in a matrix form:

$$\widetilde{W} = RS \qquad (3.4)$$

where

$$R = \begin{bmatrix} \mathbf{i}_1^T \\ \vdots \\ \mathbf{i}_F^T \\ \mathbf{j}_1^T \\ \vdots \\ \mathbf{j}_F^T \end{bmatrix} \qquad (3.5)$$

represents the camera rotation, and

$$S = \begin{bmatrix} \mathbf{s}_1 & \cdots & \mathbf{s}_P \end{bmatrix} \qquad (3.6)$$

is the shape matrix. In fact, the rows of $R$ represent the orientations of the horizontal and vertical camera reference axes throughout the stream, while the columns of $S$ are the coordinates of the $P$ feature points with respect to their centroid.

Since $R$ is $2F \times 3$ and $S$ is $3 \times P$, the equation (3.4) implies the following.

> **Rank Theorem**: *Without noise, the registered measurement matrix $\widetilde{W}$ is at most of rank three.*

The rank theorem expresses the fact that the $2F \times P$ image measurements are highly redundant. Indeed, they could all be described concisely by giving $F$ frame reference systems and $P$ point coordinate vectors, if only these were known.

From the first and the last line of equation (3.2), the original unregistered matrix $W$ can be written as

$$W = RS + \mathbf{t}\mathbf{e}_P^T \ , \tag{3.7}$$

where $\mathbf{t} = (a_1, \ldots, a_F, b_1, \ldots, b_F)^T$ is a $2F$-dimensional vector that collects the projections of camera translation along the image plane (see equation (3.2)), and $\mathbf{e}_P^T = (1, \ldots, 1)$ is a vector of $P$ ones. In scalar form,

$$\begin{aligned} u_{fp} &= \mathbf{i}_f^T \mathbf{s}_p + a_f \\ v_{fp} &= \mathbf{j}_f^T \mathbf{s}_p + b_f \ . \end{aligned} \tag{3.8}$$

Comparing with equations (3.1), we see that the two components of camera translation along the image plane are simply the averages of the rows of $W$.

In the equations above, $\mathbf{i}_f$ and $\mathbf{j}_f$ are mutually orthogonal unit vectors, so they must satisfy the constraints

$$|\mathbf{i}_f| = |\mathbf{j}_f| = 1 \quad \text{and} \quad \mathbf{i}_f^T \mathbf{j}_f = 0 \ . \tag{3.9}$$

Also, the rotation matrix $R$ is unique if the system of reference for the solution is aligned, say, with that of the first camera position, so that:

$$\mathbf{i}_1 = (1, 0, 0)^T \quad \text{and} \quad \mathbf{j}_1 = (0, 1, 0)^T \ . \tag{3.10}$$

The registered measurement matrix $\widetilde{W}$ must be at most of rank three without noise. When noise corrupts the images, however, $\widetilde{W}$ will not be exactly of rank 3. However, the rank theorem can be extended to the case of noisy measurements in a well-defined manner. The next sectionintroduces the notion of approximate rank, using the concept of singular value decomposition [Golub and Reinsch, 1971].

## 3.2 Approximate Rank

Assuming [1] that $2F \geq P$, the matrix $\widetilde{W}$ can be decomposed [Golub and Reinsch, 1971] into a $2F \times P$ matrix $O_1$, a diagonal $P \times P$ matrix $\Sigma$, and a $P \times P$ matrix $O_2$,

$$\widetilde{W} = O_1 \Sigma O_2 \ , \tag{3.11}$$

such that $O_1^T O_1 = O_2^T O_2 = O_2 O_2^T = \mathcal{I}$, where $\mathcal{I}$ is the $P \times P$ identity matrix. $\Sigma$ is a diagonal matrix whose diagonal entries are the *singular values* $\sigma_1 \geq \ldots \geq \sigma_P$ sorted in non-decreasing order. This is the *Singular Value Decomposition* (SVD) of the matrix $\widetilde{W}$.

Suppose that we pay attention only to the first three columns of $O_1$, the first $3 \times 3$ submatrix of $\Sigma$ and the first three rows of $O_2$. If we partition the matrices $O_1$, $\Sigma$, and $O_2$ as follows:

$$\begin{aligned}
O_1 &= \left[ \underbrace{O_1'}_{3} \ \middle| \ \underbrace{O_1''}_{P-3} \right] \}{\scriptstyle 2F} \\
\Sigma &= \begin{array}{c} \left[ \begin{array}{c|c} \Sigma' & 0 \\ \hline 0 & \Sigma'' \end{array} \right] \\ \underbrace{\qquad}_{3} \underbrace{\qquad}_{P-3} \end{array} \begin{array}{l} \}{\scriptstyle 3} \\ \}{\scriptstyle P-3} \\ \ \end{array} \\
O_2 &= \underbrace{\left[ \begin{array}{c} O_2' \\ \hline O_2'' \end{array} \right]}_{P} \begin{array}{l} \}{\scriptstyle 3} \\ \}{\scriptstyle P-3} \end{array} \ ,
\end{aligned} \tag{3.12}$$

we have

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

Let $\widetilde{W}^*$ be the ideal registered measurement matrix, that is, the matrix we would obtain in the absence of noise. Because of the rank theorem, $\widetilde{W}^*$ has at most three non-zero singular values. Since the singular values in $\Sigma$ are sorted in non-increasing order, $\Sigma'$ must contain all the singular values of

[1] This assumption is not crucial: if $2F < P$, everything can be repeated for the transpose of $\widetilde{W}$.

$\widetilde{W}^*$ that exceed the noise level. As a consequence, the term $O_1''\Sigma''O_2''$ must be due entirely to noise, and the best possible rank-3 approximation to the ideal registered measurement matrix $\widetilde{W}^*$ is the product:

$$\hat{W} = O_1'\Sigma'O_2'$$

We can now restate our rank theorem for the case of noisy measurements.

> **Rank Theorem for Noisy Measurements:** *All the shape and rotation information in $\widetilde{W}$ is contained in its three greatest singular values, together with the corresponding left and right eigenvectors.*

Now if we define

$$\begin{aligned} \hat{R} &= O_1'[\Sigma']^{1/2} \\ \hat{S} &= [\Sigma']^{1/2}O_2' , \end{aligned}$$

we can write

$$\hat{W} = \hat{R}\hat{S} . \tag{3.13}$$

The two matrices $\hat{R}$ and $\hat{S}$ are of the same size as the desired rotation and shape matrices $R$ and $S$: $\hat{R}$ is $2F \times 3$, and $\hat{S}$ is $3 \times P$. However, the decomposition (3.13) is not unique. In fact, if $Q$ is *any* invertible $3 \times 3$ matrix, the matrices $\hat{R}Q$ and $Q^{-1}\hat{S}$ are also a valid decomposition of $\hat{W}$, since

$$(\hat{R}Q)(Q^{-1}\hat{S}) = \hat{R}(QQ^{-1})\hat{S} = \hat{R}\hat{S} = \hat{W} .$$

Thus, $\hat{R}$ and $\hat{S}$ are in general different from $R$ and $S$. A striking fact, however, is that except for noise the matrix $\hat{R}$ is a linear transformation of the true rotation matrix $R$, and the matrix $\hat{S}$ is a linear transformation of the true shape matrix $S$. Indeed, in the absence of noise, $R$ and $\hat{R}$ both span the column space of the registered measurement matrix $\widetilde{W} = \widetilde{W}^* = \hat{W}$. Since that column space is three-dimensional because of the rank theorem, $R$ and $\hat{R}$ are different bases for the same space, and there must be a linear transformation between them.

Whether the noise level is low enough that it can be ignored at this juncture depends also on the camera motion and on shape. Notice, however, that the singular value decomposition yields sufficient information to make this decision: the requirement is that the ratio between the third and the fourth largest singular values of $\widetilde{W}$ be sufficiently large.

## 3.3 The Metric Constraints

We have found that the matrix $\hat{R}$ is a linear transformation of the true rotation matrix $R$. Likewise, $\hat{S}$ is a linear transformation of the true shape matrix $S$. More specifically, there exists a $3 \times 3$ matrix $Q$ such that

$$\begin{aligned} R &= \hat{R}Q \\ S &= Q^{-1}\hat{S} \, . \end{aligned} \tag{3.14}$$

In order to find $Q$ we observe that the rows of the true rotation matrix $R$ are unit vectors and the first $F$ are orthogonal to corresponding $F$ in the second half of $R$. These *metric constraints* yield the over-constrained, quadratic system

$$\begin{aligned} \hat{\mathbf{i}}_f^T QQ^T \hat{\mathbf{i}}_f &= 1 \\ \hat{\mathbf{j}}_f^T QQ^T \hat{\mathbf{j}}_f &= 1 \\ \hat{\mathbf{i}}_f^T QQ^T \hat{\mathbf{j}}_f &= 0 \end{aligned} \tag{3.15}$$

in the entries of $Q$. This is a simple data fitting problem which, though nonlinear, can be solved efficiently and reliably. Its solution is determined up to a rotation of the whole reference system, since the orientation of the world reference system was arbitrary. This arbitrariness can be removed by enforcing the constraints (3.10), that is, selecting the $x - y$ axes of the world reference system to be parallel with those of the first frame.

## 3.4 Outline of the Complete Algorithm

Based on the development in the previous chapters, we now have a complete algorithm for the factorization of the registered measurement matrix $\widetilde{W}$ derived from a stream of images into shape $S$ and rotation $R$ as defined in equations (3.4) - (3.6).

1. Compute the singular-value decomposition $\widetilde{W} = O_1 \Sigma O_2$.

2. Define $\hat{R} = O_1'(\Sigma')^{1/2}$ and $\hat{S} = (\Sigma')^{1/2} O_2'$, where the primes refer to the block partitioning defined in (3.12).

3. Compute the matrix $Q$ in equations (3.14) by imposing the metric constraints (equations (3.15)).

4. Compute the rotation matrix $R$ and the shape matrix $S$ as $R = \hat{R}Q$ and $S = Q^{-1}\hat{S}$.

5. If desired, align the first camera reference system with the world reference system by forming the products $RR_0$ and $R_0^T S$, where the orthonormal matrix $R_0 = [\mathbf{i}_1\ \mathbf{j}_1\ \mathbf{k}_1]$ rotates the first camera reference system into the identity matrix.

# Chapter 4

# Experiment

We test the factorization method with two real streams of images: one taken in a controlled laboratory environment with ground-truth motion data, and the other in an outdoor environment with a hand-held camcorder.

## 4.1 "Hotel" Image Stream in a Laboratory

Some frames in this stream are shown in figure 4.1. The images depict a small plastic model of a building. The camera is a Sony CCD camera with a 200 mm lens, and is moved by means of a high-precision positioning platform. Camera pitch, yaw, and roll around the model are all varied as shown by the dashed curves in figure 4.2. The translation of the camera is such as to keep the building within the field of view of the camera.

For feature tracking, we extended the Lucas-Kanade method described in [Lucas and Kanade, 1981] to allow also for the automatic selection of image features. The Lucas-Kanade method of tracking obtains the displacement vector of the window around a feature as the solution of a linear $2 \times 2$ equation system. As good image features we select those points for which the above equation systems are stable. The details are presented in [Tomasi, 1991, Tomasi and Kanade, 1992].

The entire set of 430 features thus selected is displayed in figure 4.3, overlaid on the first frame of the stream. Of these features, 42 were abandoned during tracking because their appearance changed too much. The trajectories of the remaining 388 features are used as the measurement matrix for the computation of shape and motion.

The motion recovery is precise. The plots in figure 4.2 compare the rotation components computed by the factorization method (solid curves) with the values measured mechanically from the mobile platform (dashed curves). The differences are magnified in figure 4.4. The errors are everywhere less than 0.4 degrees and on average 0.2 degrees. The computed motion follows closely also rotations with curved profiles, such as the roll profile between frames 1 and 20 (second plot in figure 4.2), and faithfully preserves all discontinuities in the rotational velocities: the factorization method does not smooth the results.

Between frames 60 and 80, yaw and pitch are nearly constant, and the camera merely rotates about its optical axis. That is, the motion is actually degenerate during this period, but still it has been correctly recovered. This demonstrates that the factorization method can deal without difficulty with streams that contain degenerate substreams, because the information in the stream is used *as a whole* in the method.

The shape results are evaluated qualitatively in figure 4.5, which shows the computed shape viewed from above. The view in figure 4.5 is similar to that in figure 4.6, included for visual comparison. Notice that the walls, the windows on the roof, and the chimneys are recovered in their correct positions.

To evaluate the shape performance quantitatively, we measured some distances on the actual house model with a ruler and compared them with the distances computed from the point coordinates in the shape results. Figure 4.7 shows the selected features. The diagram in figure 4.8 shows the distances between pairs of features measured on the actual model and those computed by the factorization method. The measured distances between the steps along the right side of the roof (7.2 mm) were obtained by measuring five steps and dividing the total distance (36 mm) by five. The differences between computed and measured results are of the order of the resolution of our ruler measurements (one millimeter).

Part of the errors in the results is due to the use of orthography as the projection model. However, it tends to be fairly small for many realistic situations. In fact, it has been shown that errors due to the orthogrphic distortion are approximately about the same percentage as the ratio of the object size in depth to the distance of the object from the camera [Tomasi, 1991].

## 4.2 Outdoor "House" Image Stream

The factorization method has been tested with an image stream of a real building, taken with a hand-held camera. Figure 4.9 shows some of the 180 frames of the building stream. The overall motion covers a relatively small rotation angle, approximately 15 degrees. Outdoor images are harder to process than those produced in a controlled environment of the laboratory, because lighting changes less predictably and the motion of the camera is more difficult to control. As a consequence, features are harder to track: the images are unpredictably blurred by motion, and corrupted by vibrations of the video recorder's head, both during recording and digitization. Furthermore, the camera's jumps and jerks produce a wide range of image disparities.

The features found by the selection algorithm in the first frame are shown in figure 4.10. There are many false features. The reflections in the window partially visible in the top left of the image move non-rigidly. More false features can be found in the lower left corner of the picture, where the vertical bars of the handrail intersect the horizontal edges of the bricks of the wall behind. We masked away these two parts of the image from the analysis.

In total, 376 features were found by the selection algorithm and tracked. Figure 4.11 plots the tracks of some (60) of the features for illustration. Notice the very jagged trajectories due to the vibrating motion of the hand-held camera.

Figures 4.12 and 4.13 show a front and a top view of the building as reconstructed by the factorization method. To render these figures for display, we triangulated the computed 3D points into a set of small surface patches and mapped the pixel values in the first frame onto the resulting surface. The structure of the visible part of the building's three walls has clearly been reconstructed. In these figures, the left wall appears to bend somewhat on the right where it intersects the middle wall. This occurred because the feature selector found features along the shadow of the roof just on the right of the intersection of the two walls, rather than at the intersection itself. Thus, the appearance of a bending wall is an artifact of the triangulation done for rendering.

This experiment with an image stream taken outdoors with the jerky motion produced by a hand-held camera demonstrates that the factorization method does not require a smooth motion assumption. The identification of false features, that is, of features that do not move rigidly with respect of the environment, remains an open problem that must be solved for a fully autonomous system. An initial effort has been seen in [Boult and Brown, 1991].

# Chapter 5

# Occlusions

In reality, as the camera moves, features can appear and disappear from the image, because of occlusions. Also, a feature tracking method will not always succeed in tracking features throughout the image stream. These phenomena are frequent enough to make a shape and motion computation method unrealistic if it cannot deal with them.

Sequences with appearing and disappearing features result in a measurement matrix $W$ which is only partially filled in. The factorization method introduced in chapter3 cannot be applied directly. However, there is usually sufficient information in the stream to determine all the camera positions and all the three-dimensional feature point coordinates. If that is the case, we can not only solve the shape and motion recovery problem from the incomplete measurement matrix $W$, but we can even hallucinate the unknown entries of $W$ by projecting the computed three-dimensional feature coordinates onto the computed camera positions.

## 5.1 Solution for Noise-Free Images

Suppose that a feature point is not visible in a certain frame. If the same feature is seen often enough in other frames, its position in space should be recoverable. Moreover, if the frame in question includes enough other features, the corresponding camera position be recoverable as well. Then from point and camera positions thus recovered, we should also be able to reconstruct the missing image measurement. Formally, we have the following

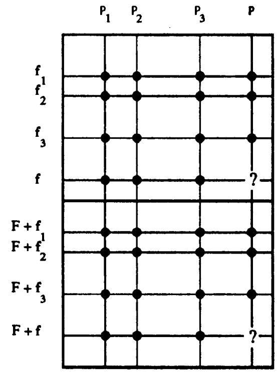


Figure 5.1: The Reconstruction Condition. If the dotted entries of the measurement matrix are known, the two unknown ones (question marks) can be reconstructed.

sufficient condition.

> **Condition for Reconstruction:** In the absence of noise, an unknown image measurement pair $(u_{fp}, v_{fp})$ in frame $f$ can be reconstructed if point $p$ is visible in at least three more frames $f_1, f_2, f_3$, and if there are at least three more points $p_1, p_2, p_3$ that are visible in all the four frames: the original $f$ and the additional $f_1, f_2, f_3$.

Referring to Figure 5.1, this means that the dotted entries must be known to reconstruct the question marks. This is equivalent to Ullman's result [Ullman, 1979] that three views of four points determine structure and motion. In this section, we prove the reconstruction condition in our formalism and develop the reconstruction procedure. To this end, we notice that the rows and columns of the noise-free measurement matrix $W$ can always be permuted so that $f_1 = p_1 = 1$, $f_2 = p_2 = 2$, $f_3 = p_3 = 3$, $f = p = 4$. We can therefore suppose that $u_{44}$ and $v_{44}$ are the only two unknown entries in the

$8 \times 4$ matrix

$$W = \left[\frac{U}{V}\right] = \begin{bmatrix} u_{11} & u_{12} & u_{13} & u_{14} \\ u_{21} & u_{22} & u_{23} & u_{24} \\ u_{31} & u_{32} & u_{33} & u_{34} \\ u_{41} & u_{42} & u_{43} & ? \\ v_{11} & v_{12} & v_{13} & v_{14} \\ v_{21} & v_{22} & v_{23} & v_{24} \\ v_{31} & v_{32} & v_{33} & v_{34} \\ v_{41} & v_{42} & v_{43} & ? \end{bmatrix} .$$

Then, the factorization method can be applied to the first three rows of $U$ and $V$, that is, to the $6 \times 4$ submatrix

$$W_{6\times4} = \begin{bmatrix} u_{11} & u_{12} & u_{13} & u_{14} \\ u_{21} & u_{22} & u_{23} & u_{24} \\ u_{31} & u_{32} & u_{33} & u_{34} \\ v_{11} & v_{12} & v_{13} & v_{14} \\ v_{21} & v_{22} & v_{23} & v_{24} \\ v_{31} & v_{32} & v_{33} & v_{34} \end{bmatrix} \tag{5.1}$$

to produce the partial translation and rotation submatrices

$$\mathbf{t}_{6\times1} = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ b_1 \\ b_2 \\ b_3 \end{bmatrix} \quad \text{and} \quad R_{6\times3} = \begin{bmatrix} \mathbf{i}_1^T \\ \mathbf{i}_2^T \\ \mathbf{i}_3^T \\ \mathbf{j}_1^T \\ \mathbf{j}_2^T \\ \mathbf{j}_3^T \end{bmatrix} \tag{5.2}$$

and the full shape matrix

$$S = \begin{bmatrix} \mathbf{s}_1 & \mathbf{s}_2 & \mathbf{s}_3 & \mathbf{s}_4 \end{bmatrix} \tag{5.3}$$

such that

$$W_{6\times4} = R_{6\times3}S + \mathbf{t}_{6\times1}\mathbf{e}_4^T$$

where $\mathbf{e}_4^T = (1, 1, 1, 1)$.

To complete the rotation solution, we need to compute the vectors $\mathbf{i}_4$ and $\mathbf{j}_4$. However, a registration problem must be solved first. In fact, only three points are visible in the fourth frame, while equation (5.3) yields all four points in space. Since the factorization method computes the space coordinates with respect to the centroid of the points, we have $\mathbf{s}_1 + \mathbf{s}_2 + \mathbf{s}_3 + \mathbf{s}_4 = 0$, while the image coordinates in the fourth frame are measured with respect to the centroid of just three observed points (1, 2, 3). Thus, before we can compute $\mathbf{i}_4$ and $\mathbf{j}_4$ we must make the two origins coincide by referring all coordinates to the centroid

$$\mathbf{c} = \frac{1}{3}(\mathbf{s}_1 + \mathbf{s}_2 + \mathbf{s}_3)$$

of the three points that are visible in all four frames. In the fourth frame, the projection of $\mathbf{c}$ has coordinates

$$\begin{aligned} a'_4 &= \frac{1}{3}(u_{41} + u_{42} + u_{43}) \\ b'_4 &= \frac{1}{3}(v_{41} + v_{42} + v_{43}) \ , \end{aligned}$$

so we can define the new coordinates

$$\mathbf{s}'_p = \mathbf{s}_p - \mathbf{c} \quad \text{for} \quad p = 1, 2, 3$$

in space and

$$\begin{aligned} u'_{4p} &= u_{4p} - a'_4 \\ v'_{4p} &= v_{4p} - b'_4 \end{aligned} \quad \text{for} \quad p = 1, 2, 3$$

in the fourth frame. Then, $\mathbf{i}_4$ and $\mathbf{j}_4$ are the solutions of the two $3 \times 3$ systems

$$\begin{aligned} \left[ \begin{array}{ccc} u'_{41} & u'_{42} & u'_{43} \end{array} \right] &= \mathbf{i}_4^T \left[ \begin{array}{ccc} \mathbf{s}'_1 & \mathbf{s}'_2 & \mathbf{s}'_3 \end{array} \right] \\ \left[ \begin{array}{ccc} v'_{41} & v'_{42} & v'_{43} \end{array} \right] &= \mathbf{j}_4^T \left[ \begin{array}{ccc} \mathbf{s}'_1 & \mathbf{s}'_2 & \mathbf{s}'_3 \end{array} \right] \end{aligned} \tag{5.4}$$

derived from equation (3.4). The second equation in (5.2) and the solution to (5.4) yield the entire rotation matrix $R$, while shape is given by equation (5.3).

The components $a_4$ and $b_4$ of translation in the fourth frame with respect to the centroid of all four points can be computed by postmultiplying equation (3.7) by the vector $\eta_4 = (1, 1, 1, 0)^T$:

$$W\eta_4 = RS\eta_4 + \mathbf{t}\mathbf{e}_4^T\eta_4 \ .$$

Since $\mathbf{e}_4^T \eta_4 = 3$, we obtain

$$\mathbf{t} = \frac{1}{3}(W - RS)\eta_4 \,. \tag{5.5}$$

In particular, rows 4 and 8 of this equation yield $a_4$ and $b_4$. Notice that the unknown entries $u_{44}$ and $v_{44}$ are multiplied by zeros in equation (5.5).

Now that both motion and shape are known, the missing entries $u_{44}$, $v_{44}$ of the measurement matrix $W$ can be found by orthographic projection (equation (3.8)):

$$\begin{aligned} u_{44} &= \mathbf{i}_4^T \mathbf{s}_4 + a_4 \\ v_{44} &= \mathbf{j}_4^T \mathbf{s}_4 + b_4 \,. \end{aligned}$$

The procedure thus completed factors the full $6 \times 4$ submatrix of $W$ and then reasons on the three points that are visible in all the frames to compute motion for the fourth frame. Alternatively, one can first apply factorization to the $8 \times 3$ submatrix

$$W_{8\times3} = \begin{bmatrix} u_{11} & u_{12} & u_{13} \\ u_{21} & u_{22} & u_{23} \\ u_{31} & u_{32} & u_{33} \\ u_{41} & u_{42} & u_{43} \\ v_{11} & v_{12} & v_{13} \\ v_{21} & v_{22} & v_{23} \\ v_{31} & v_{32} & v_{33} \\ v_{41} & v_{42} & v_{43} \end{bmatrix} \tag{5.6}$$

to produce the full translation and rotation submatrices

$$\mathbf{t}' = \begin{bmatrix} a_1' \\ a_2' \\ a_3' \\ a_4' \\ b_1' \\ b_2' \\ b_3' \\ b_4' \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} \mathbf{i}_1^T \\ \mathbf{i}_2^T \\ \mathbf{i}_3^T \\ \mathbf{i}_4^T \\ \mathbf{j}_1^T \\ \mathbf{j}_2^T \\ \mathbf{j}_3^T \\ \mathbf{j}_4^T \end{bmatrix} \tag{5.7}$$

and the partial shape matrix

$$S_{3\times 3} = \begin{bmatrix} \mathbf{s}'_1 & \mathbf{s}'_2 & \mathbf{s}'_3 \end{bmatrix} \tag{5.8}$$

such that

$$W_{8\times 3} = RS'_{3\times 3} + \mathbf{t}'\mathbf{e}_3^T \ .$$

The primes here signal again that coordinates refer to the centroid of only the first three points. Then, this partial solution can be extended to $\mathbf{s}'_4$ by solving the following overconstrained system of six equations in the three unknown entries of $\mathbf{s}'_4$:

$$\begin{bmatrix} \mathbf{i}_1^T \\ \mathbf{i}_2^T \\ \mathbf{i}_3^T \\ \mathbf{j}_1^T \\ \mathbf{j}_2^T \\ \mathbf{j}_3^T \end{bmatrix} \mathbf{s}'_4 + \begin{bmatrix} a'_1 \\ a'_2 \\ a'_3 \\ b'_1 \\ b'_2 \\ b'_3 \end{bmatrix} = \begin{bmatrix} u'_{14} \\ u'_{24} \\ u'_{34} \\ v'_{14} \\ v'_{24} \\ v'_{34} \end{bmatrix} \tag{5.9}$$

where

$$\begin{array}{l} u'_{f4} = u_{f4} - a'_f \\ v'_{f4} = v_{f4} - b'_f \end{array} \quad \text{for} \quad f = 1,2,3 \ .$$

The "primed" shape coordinates can now be registered with respect to their centroid to yield the "unprimed" coordinates:

$$\mathbf{s}_p = \mathbf{s}'_p - \frac{1}{4}S'\mathbf{e}_4 \quad \text{for} \quad p = 1,2,3,4$$

and the "unprimed" translation can again be found from equation (5.5).

In summary, the full motion and shape solution can be found in either of the following ways:

1. factor $W_{6\times 4}$ to find a partial motion and full shape solution, and propagate it to include motion for the remaining frame (equations (5.4)). This will be used for reconstructing the complete $W$ by row-wise extension.

2. factor $W_{8\times 3}$ to find a full motion and partial shape solution, and propagate it to include the remaining feature point (equation (5.9)). This will be used for reconstructing the complete $W$ by column-wise extension.

## 5.2 Solution in the Presence of Noise

The solution propagation method introduced in the previous sectioncan be extended to $2F \times P$ measurement matrices with $F \geq 4$ and $P \geq 4$. In fact, the only difference is that the propagation equations (5.4) and (5.9) now become overconstrained. If the measurement matrix $W$ is noisy, this redundancy is beneficial, since equations (5.4) and (5.9) can be solved in the Least Square Error sense, and the effect of noise is reduced.

In the general case of a noisy $2F \times P$ matrix $W$ the solution propagation method can be summarized as follows. A possibly large, full subblock of $W$ is first decomposed by factorization. Then, this initial solution is grown one row or one column at a time by solving systems analogous to those in (5.4) or (5.9) in the Least Square Error sense.

However, because of noise, the order in which the rows and columns of $W$ are incorporated into the solution can affect the exact values of the final motion and shape solution. Consequently, once the solution has been propagated to the entire measurement matrix $W$, it may be necessary to refine the results with a steepest-descent minimization of the residue

$$||W - RS - \frac{1}{P}\mathbf{t}\mathbf{e}_P^T||$$

(see equation (3.7)).

There remain the two problems of how to choose the initial full subblock to which factorization is applied and in what order to grow the solution. In fact, however, because of the final refinement step, neither choice is critical as long as the initial matrix is large enough to yield a good starting point. We illustrate this point in the next chapterof experiments.

# Chapter 6

# More Experiments

We will first test the propagation method with image streams which include substantial occlusions. We first use an image stream taken in a laboratory. Then, we demonstrate the robustness of the factorization method with another stream taken with a hand-held amateur camera.

## 6.1 "Ping-Pong Ball" Image Stream

A ping-pong ball with black dots marked on its surface is rotated 450 degrees in front of the camera, so features appear and disappear. The rotation between adjacent frames is 2 degrees, so the stream is 226 frames long. Figure 6.14 shows the first frame of the stream, with the automatically selected features overlaid.

Every 30 frames (60 degrees) of rotation, the feature tracker looks for new features. In this way, features that disappear on one side around the ball are replaced by new ones that appear on the other side. Figure 6.15 shows the tracks of 60 features, randomly chosen among the total 829 found by the selector.

If all measurements are collected into the noisy measurement matrix $W$, the $U$ and $V$ parts of $W$ have the same fill pattern: if the $x$ coordinate of a measurement is known, so is the $y$ coordinate. Figure 6.16 shows this *fill matrix* for our experiment. This matrix has the same size as either $U$ or $V$, that is, $F \times P$. A column corresponds to a feature point, and a row to a frame. Shaded regions denote known entries. The fill matrix shown has

$226 \times 829 = 187354$ entries, of which 30185 (about 16 percent) are known.

To start the motion and shape computation, the algorithm finds a large full submatrix by applying simple heuristics based on typical patterns of the fill matrix. The choice of the starting matrix is not critical, as long as it leads to a reliable initialization of the motion and shape matrices. The initial solution is then grown by repeatedly solving overconstrained versions of the linear system corresponding to (5.4) to add new rows, and of the system corresponding to (5.9) to add new columns. The rows and columns to add are selected so as to maximize the redundancy of the linear systems. Eventually, all of the motion and shape values are determined. As a result, the unknown 84 percent of the measurement matrix can be hallucinated from the known 16 percent.

Figure 6.17 shows two views of the final shape results, taken from the top and from the side. The missing features at the bottom of the ball in the side view correspond to the part of the ball that remained always invisible, because it rested on the rotating platform.

To display the motion results, we look at the $\mathbf{i}_f$ and $\mathbf{j}_f$ vectors directly. We recall that these unit vectors point along the rows and columns of the image frames $f$ in $1, \ldots, F$. Because the ping-pong ball rotates around a fixed axis, both $\mathbf{i}_f$ and $\mathbf{j}_f$ should sweep a cone in space, as shown in Figure 6.18. The tips of $\mathbf{i}_f$ and $\mathbf{j}_f$ should describe two circles in space, centered along the axis of rotation. Figure 6.19 shows two views of these vector tips, from the top and from the side. Those trajectories indicate that the motion recovery was done correctly. Notice the double arc in the top part of figure 6.19 corresponding to more than 360 degrees rotation. If the motion reconstruction were perfect, the two arcs would be indistinguishable.

## 6.2 "Cup and Hand" Image Stream

In this sectionwe describe an experiment with a natural scene including occlusion as a dominant phenomenon. A hand holds a cup and rotates it by about ninety degrees in front of the camera mounted on a fixed stand. Figure 6.20 shows four out of the 240 frames of the stream.

An additional need in this experiment is figure/ground segmentation. Since the camera was fixed, however, this problem is easily solved: features that do not move belong to the background. Also, the stream includes some nonrigid motion: as the hand turns, the configuration and relative position of the fingers changes slightly. This effect, however, is small and did not affect the results appreciably.

A total of 207 features was selected. Occlusions were marked by hand in this experiment. The fill matrix of figure 6.22 illustrates the occlusion pattern. Figure 6.21 shows the image trajectory of 60 randomly selected features.

Figures 6.23 and 6.24 show a front and a top view of the cup and the visible fingers as reconstructed by the propagation method. The shape of the cup was recovered, as well as the rough shape of the fingers. These renderings were obtained, as for the "House" image stream in section4.1, by triangulating the tracked feature points and mapping pixel values onto the resulting surface.

# Chapter 7

# Conclusion

The rank theorem, which is the basis of the factorization method, is both surprising and powerful. Surprising because it states that the correlation among measurements made in an image stream has a simple expression *no matter what the camera motion is and no matter what the shape of an object is*, thus making motion or surface assumptions (such as smooth, constant, linear, planar and quadratic) fundamentally superfluous. Powerful because the rank theorem leads to factorization of the measurement matrix into shape and motion in a well-behaved and stable manner.

The factorization method exploits the redundancy of the measurement matrix to counter the noise sensitivity of structure-from-motion and allows using very short inter-frame camera motion to simplify feature tracking. The structural insight into shape-from-motion afforded by the rank theorem led to a systematic procedure to solve the occlusion problem within the factorization method. The experiments in the lab demonstrate the high accuracy of the method, and the outdoor experiments show its robustness.

The rank theorem is strongly related to Ullman's twelve year old result that three pictures of four points determine structure and motion under orthography. Thus, in a sense, the theoretical foundation of our result has been around for a long time. The factorization method evolves the applicability of that foundation from mathematical images to actual noisy image streams.

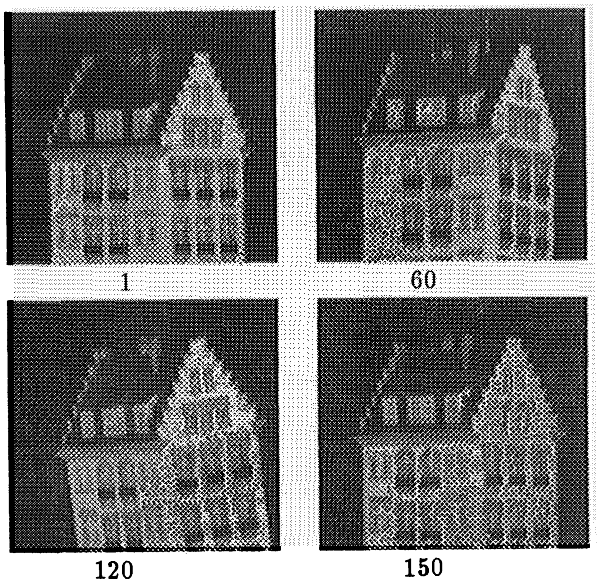


Figure 4.1: Some frames in the sequence. The whole sequence is 150 frames.

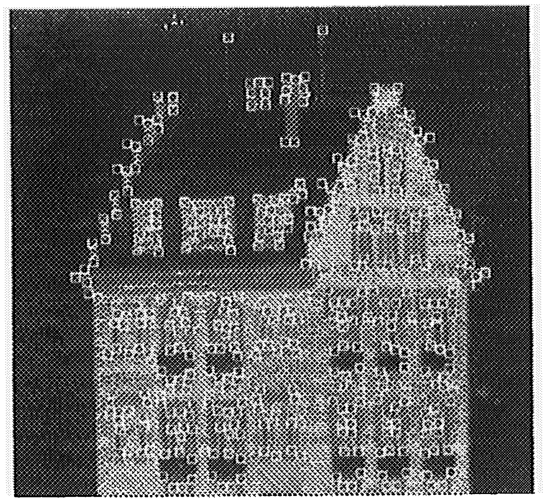

Figure 4.3: The 430 features selected by the automatic detection method.

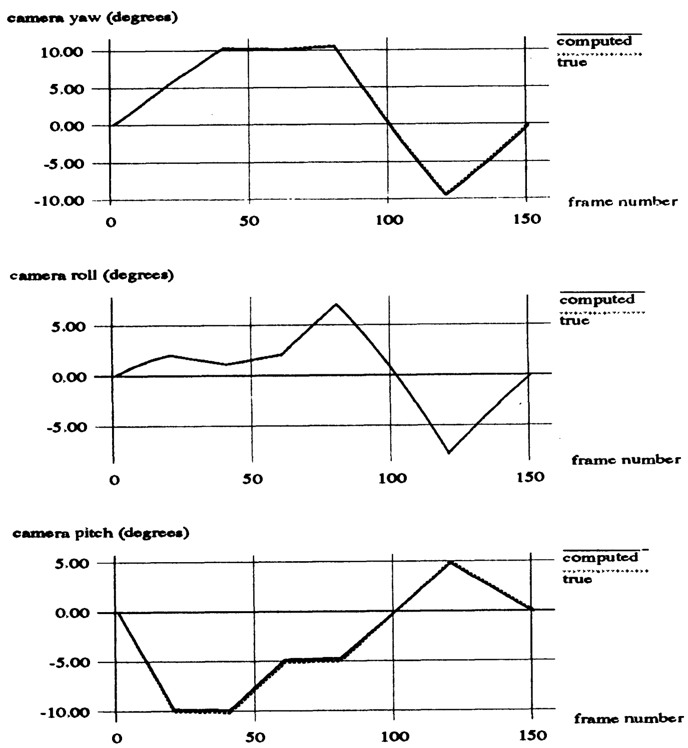


Figure 4.2: True and computed camera yaw, roll, pitch.

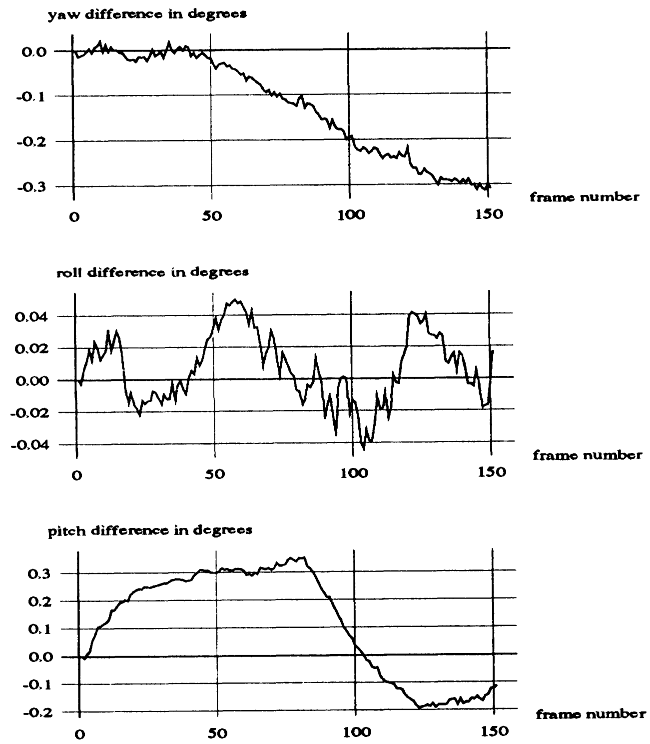


Figure 4.4: Blow-up of the errors in figure 4.2.

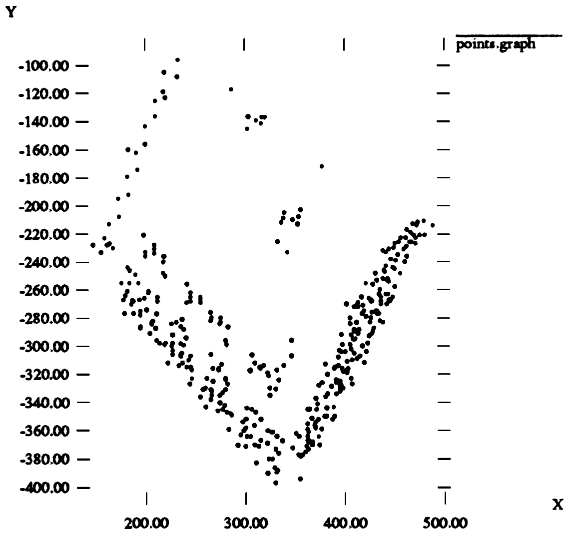


Figure 4.5: A view of the computed shape from approximately above the building (compare with figure 4.6).

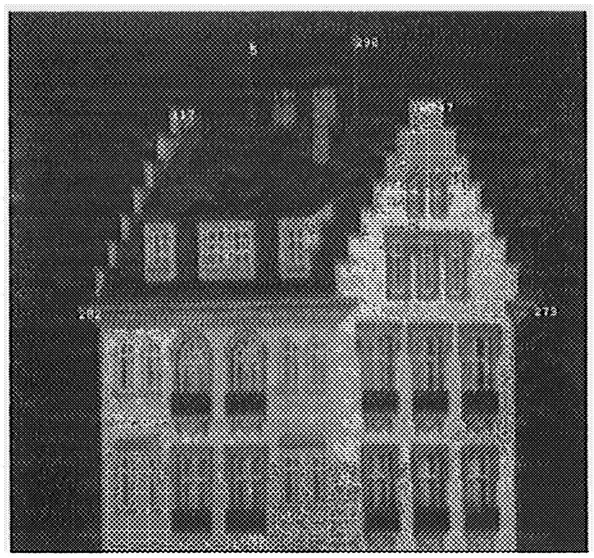

Figure 4.7: For a quantitative evaluation, distances between the features shown in the picture were measured on the actual model, and compared with the computed results. The comparison is shown in figure 4.8.

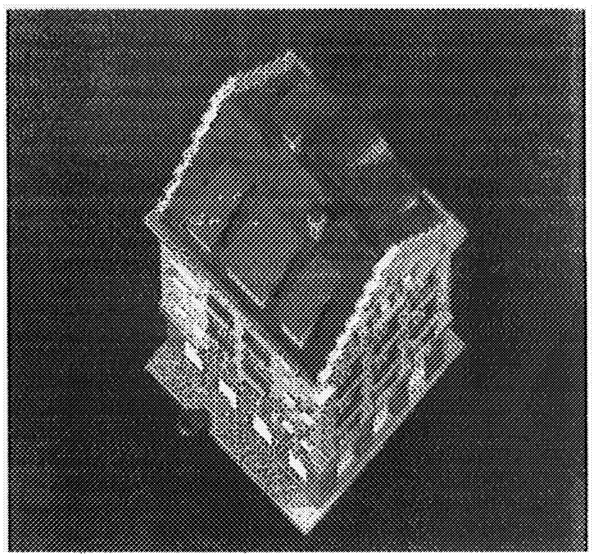

Figure 4.6: A real picture from above the building, similar to figure 4.5.

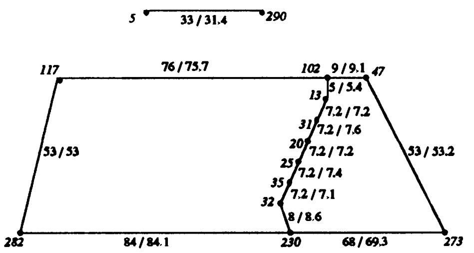


Figure 4.8: Comparison between measured and computed distances for the features in figure 4.7. The number before the slash is the measured distance, the one after is the computed distance. Lengths are in millimeters. Computed distances were scaled so that the computed distance between features 117 and 282 is the same as the measured distance.

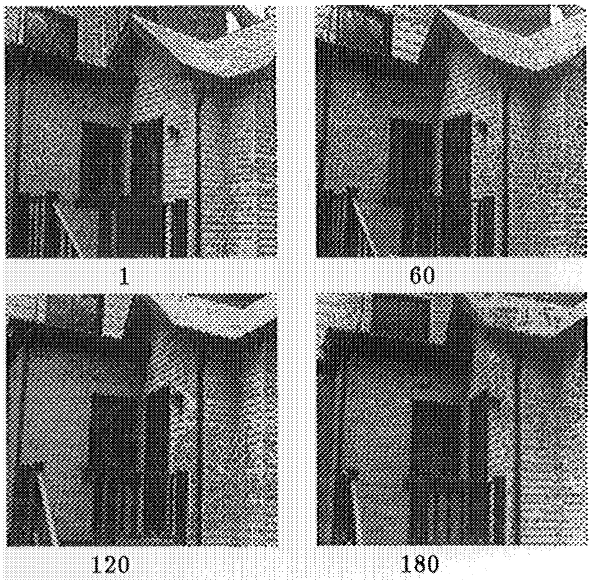


Figure 4.9: Four out of the 180 frames of the real house image stream.

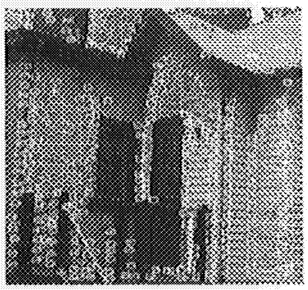

Figure 4.10: The features selected in the first frame of the real house stream (figure 4.9)
.

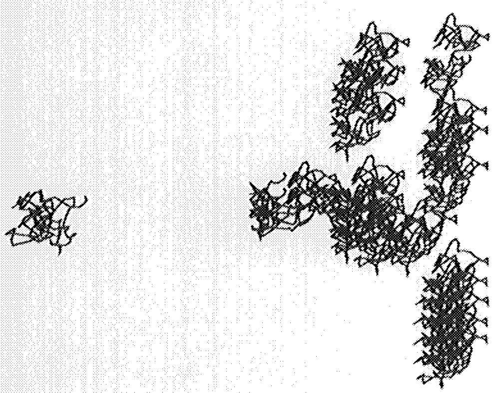

**Figure 4.11: Tracks of 60 randomly selected features from the real house stream (figure 4.9.)**

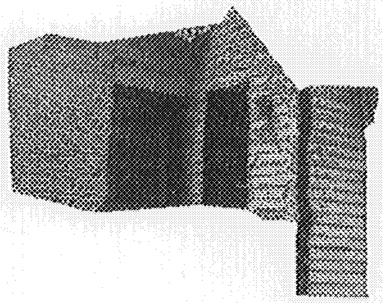

**Figure 4.12:** A front view of the three **reconstructed walls, with** the original image intensities **mapped onto the resulting** surface.

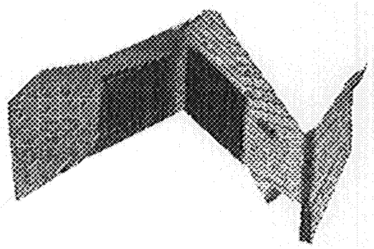

Figure **4.13: A view from above of the three reconstructed walls, with image intensities mapped onto the surface.**

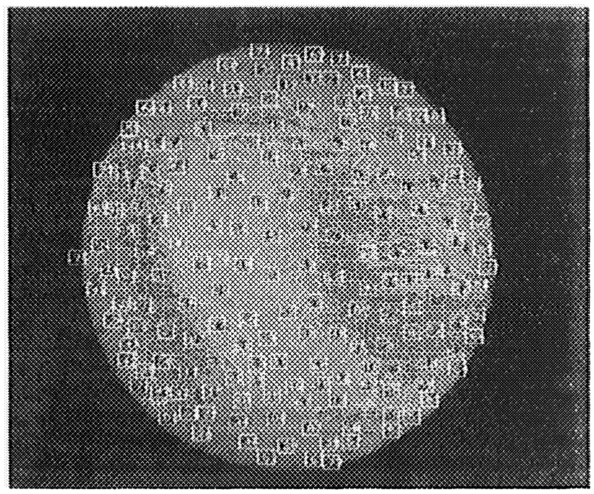

Figure 6.14: The first frame of the ping-pong stream, with overlaid features.

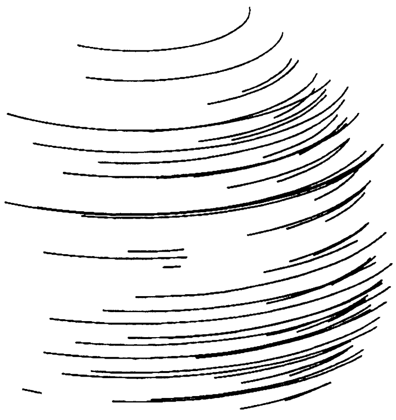

Figure 6.15: Tracks of 60 randomly selected features from the stream of figure 6.14.

Figure 6.16: The fill matrix for the ping-pong ball experiment. Shaded entries are known.

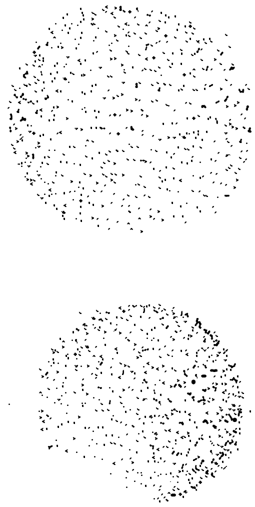

Figure 6.17: Top and side views of the reconstructed ping-pong ball.

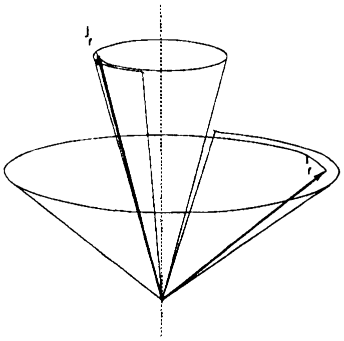

Figure 6.18: Rotational component of the camera motion for the ping-pong stream. Because rotation occurs around a fixed axis, the two mutually orthogonal unit vectors $\mathbf{i}_f$ and $\mathbf{j}_f$, pointing along rows and columns of the image sensor, sweep two 450-degree cones in space.

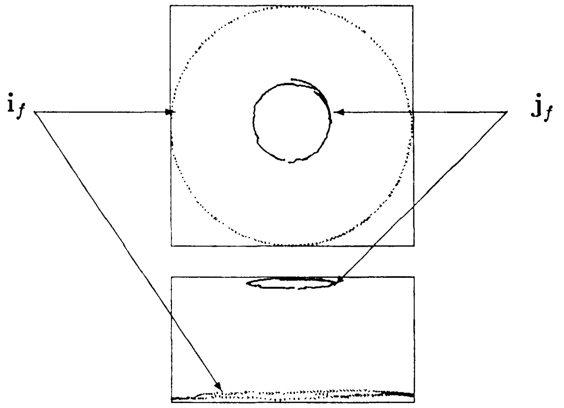


Figure 6.19: Top and side views of the $\mathbf{i}_f$ and $\mathbf{j}_f$ vectors identifying the camera rotation. See Figure 6.18.

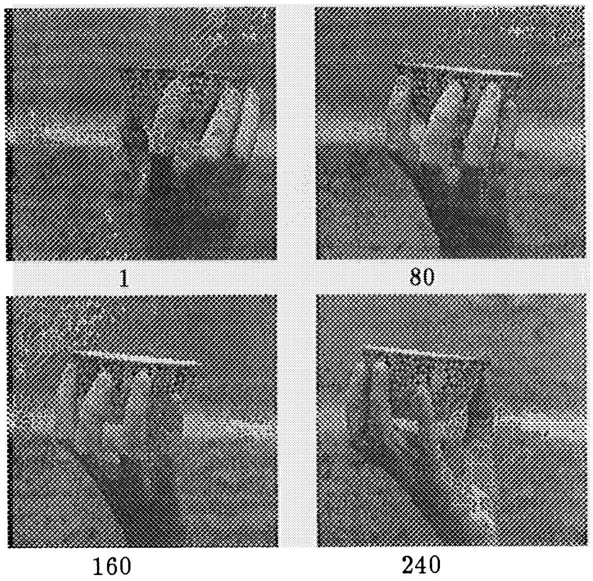


Figure 6.20: Four out of the 240 frames of the cup image stream.

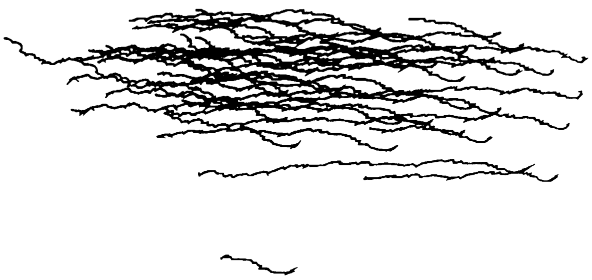

Figure 6.21: Tracks of 60 randomly selected features from the cup stream.

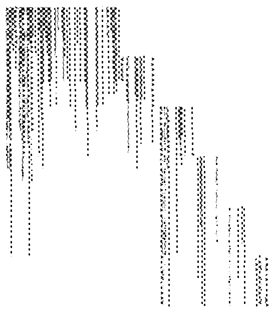

Figure 6.22: The $240 \times 207$ fill matrix for the cup stream (figure 6.20). Shaded entries are known.

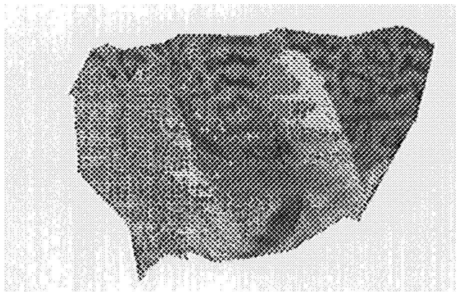

Figure 6.23: A front view of the cup and fingers, with the original image intensities mapped onto the resulting surface.

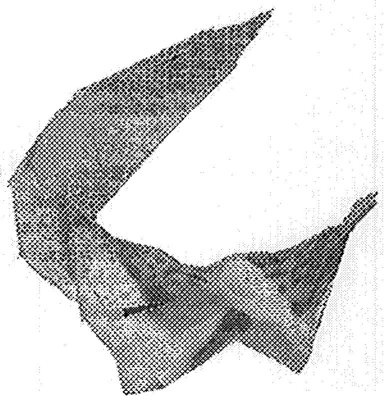

Figure 6.24: A view from above of the cup and fingers with image intensities mapped onto the surface.

[Ullman, 1979]
S. ULLMAN. *The Interpretation of Visual Motion.* The MIT Press, Cambridge, Ma, 1979.

[Waxman and Wohn, 1985]
A. M. WAXMAN AND K. WOHN. Contour evolution, neighborhood deformation, and global image flow: planar surfaces in motion. *International Journal of Robotics Research*, 4:95–108, 1985.

moving in space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-1(2):127–135, April 1979.

[Spetsakis and Aloimonos, 1989]
M. E. SPETSAKIS AND J. Y. ALOIMONOS. Optimal motion estimation. In *Proceedings of the IEEE Workshop on Visual Motion*, pages 229–237, Irvine, California, March 1989.

[Tomasi and Kanade, 1990a]
C. TOMASI AND T. KANADE. Shape and motion without depth. In *Proceedings of the Third International Conference in Computer Vision (ICCV)*, Osaka, Japan, December 1990.

[Tomasi and Kanade, 1990b]
C. TOMASI AND T. KANADE. Shape and motion without depth. In *Proceedings of the DARPA Image Understanding Workshop*, pages 258–270, Pittsburgh, Pa, September 1990.

[Tomasi and Kanade, 1991]
C. TOMASI AND T. KANADE. Shape and motion from image streams: a factorization method - 2. point features in 3d motion. Technical Report CMU-CS-91-105, Carnegie Mellon University, Pittsburgh, PA, January 1991.

[Tomasi and Kanade, 1992]
C. TOMASI AND T. KANADE. Selecting and tracking features for image sequence analysis. *submitted to Robotics and Automation*, 1992.

[Tomasi, 1991]
C. TOMASI. *Shape and Motion from Image Streams: a Factorization Method.* PhD thesis, CMU, September 1991.

[Tsai and Huang, 1984]
R. Y. TSAI AND T. S. HUANG. Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(1):13–27, January 1984.

[Debrunner and Ahuja, 1991]
C. H. DEBRUNNER AND N. AHUJA. Motion and structure factorization and segmentation of long multiple motion image sequences. Technical Report UI-BI-CV-5-91, CSL, Univ. of Illinois, Urbana-Champaign, IL, 1991.

[Golub and Reinsch, 1971]
G. H. GOLUB AND C. REINSCH. Singular value decomposition and least squares solutions, In *Handbook for Automatic Computation*, volume 2, chapter I/10, pages 134–151. Springer Verlag, New York, NY, 1971.

[Heeger and Jepson, 1989]
D. J. HEEGER AND A. JEPSON. Visual perception of three-dimensional motion. Technical Report 124, MIT Media Laboratory, Cambridge, Ma, December 1989.

[Heel, 1989]
J. HEEL. Dynamic motion vision. In *Proceedings of the DARPA Image Understanding Workshop*, pages 702–713, Palo Alto, Ca, May 23-26 1989.

[Horn *et al.*, 1988]
B. K. P. HORN, H. M. HILDEN, AND S. NEGAHDARIPOUR. Closed-form solution of absolute orientation using orthonormal matrices. *Journal of the Optical Society of America A*, 5(7):1127–1135, July 1988.

[Lucas and Kanade, 1981]
B. D. LUCAS AND T. KANADE. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 1981.

[Matthies *et al.*, 1989]
L. MATTHIES, T. KANADE, AND R. SZELISKI. Kalman filter-based algorithms for estimating depth from image sequences. *International Journal of Computer Vision*, 3(3):209–236, September 1989.

[Prazdny, 1980]
K. PRAZDNY. Egomotion and relative depth from optical flow. *Biological Cybernetics*, 102:87–102, 1980.

[Roach and Aggarwal, 1979]
J. W. ROACH AND J. K. AGGARWAL. Computer tracking of objects

# Bibliography


[Adiv, 1985]
G. ADIV. Determining three-dimensional motion and structure from optical flow generated by several moving objects. *IEEE Pattern Analysis and Machine Intelligence*, 7:384–401, 1985.

[Bolles *et al.*, 1987]
R. C. BOLLES, H. H. BAKER, AND D. H. MARIMONT. Epipolar-plane image analysis: An approach to determining structure from motion. *International Journal of Computer Vision*, 1(1):7–55, 1987.

[Boult and Brown, 1991]
T. E. BOULT AND L. G. BROWN. Factorization-based segmentation of motions. In *Proceedings of the IEEE Workshop on Visual Motion*, pages 179–186, October 1991.

[Broida *et al.*, 1990]
T. BROIDA, S. CHANDRASHEKHAR, AND R. CHELLAPPA. Recursive 3-d motion estimation from a monocular image sequence. *IEEE Transactions on Aerospace and Electronic Systems*, 26(4):639–656, July 1990.

[Bruss and Horn, 1983]
A. R. BRUSS AND B. K. P. HORN. Passive navigation. *Computer Vision, Graphics, and Image Processing*, 21:3–20, 1983.

[Debrunner and Ahuja, 1990]
C. H. DEBRUNNER AND N. AHUJA. A direct data approximation based motion estimation algorithm. In *Proceedings of the 10th International Conference on Pattern Recognition*, pages 384–389, Atlantic City, NJ, June 1990.