Direct Methods for Recovering Motion

Berthold K. P. Horn E. J. Weldon Jr.[†]

In this chapter we consider the problem of determining the motion of a monocular observer moving with respect to a rigid, unknown world. We use a "least squares," as opposed to a discrete method of solving for the motion parameters; our method uses all of the points in a two-image sequence and does not attempt to establish correspondence between the images. Hence the method is relatively robust with respect to quantization error, noise, illumination gradients, and other effects.

So far, we can determine the observer motion in two special cases:

- When the motion is pure rotation.
- When the motion is pure translation or when the rotational component of the motion is known.

At this writing we have not developed a direct method that is applicable to arbitrary motion.

Earlier work

In the continuous or least squares approach to motion vision, motion parameters are found that are consistent with the observed motion of the entire image. Bruss and Horn [1983] use this approach to calculate motion parameters assuming that the optical flow is known at each point. Adiv [1984] uses the approach of Bruss and Horn to segment the scene into independently moving planar objects; he shows that given the optical

[†]Department of Electrical Engineering, University of Hawaii at Manoa, Honolulu, Hawaii 96822.

flow, segmentation can be performed and the motion calculated. Negahdaripour and Horn [1987] eschew the use of optical flow and calculate the observer's motion directly from the spatial and temporal derivatives of the image brightness, assuming a planar world. The advantage of this direct approach, which we also use here, is that certain computational difficulties inherent in the calculation of optical flow are avoided. In particular, it is not necessary to make the usual assumption that the optical flow field is smooth; an assumption that is violated near object boundaries, necessitating flow segmentation.

Waxman and Ullman [1985] and Waxman and Wohn [1985] also avoid the discrete approach to motion vision; their techniques make use of first and second derivatives of the optical flow to compute both the motion parameters and the structure of the imaged world. In the interests of developing methods which can be implemented, the techniques presented in this chapter avoid the use of second and higher-order derivatives.

Summary of the chapter

One of our approaches to the motion vision problem can be summarized as follows: Given the observer motion and the spatial brightness function of the image one can predict the time derivative of brightness at each point in the image. We find the motion that minimizes the integral of the square of the difference between this predicted value and the observed time derivative. The integral is taken over the image region of interest, which, in the discussion here, is usually taken to be the whole image.

We use auxiliary vectors derived from the derivatives of brightness and the image position that occur in the basic brightness change constraint equation. Study of the distribution of the directions of these vectors on the unit sphere suggests specific algorithms and also helps uncover relationships between accuracy and parameters of the imaging situation.

We have developed a simple robust algorithm for recovering the angular velocity vector in the case of pure rotation. This algorithm involves solving three linear equations in the three unknown components of the rotation vector. The coefficients of the equations are moments of components of one of the auxiliary vectors over the given image region. We show that the accuracy of the recovered component of rotation about the direction towards the image region is poor relative to the other components, unless the image region subtends a substantial solid angle.

We have developed several algorithms for recovering the translational velocity in the case of pure translation. These algorithms exploit the constraint that objects have to be in front of the camera in order to be imaged. This constraint leads to a non-linear constrained optimization problem. The performance of these algorithms depends on a number of factors including:

- Angle subtended by the image, that is, the field of view.
- Direction of motion relative to the optical axis.
- Depth range.
- Distribution of brightness gradients.
- Noise in the estimated time derivative of brightness.
- Noise in the estimated spatial gradient of brightness.
- Number of picture cells considered.

We have not yet been able to select a "best" algorithm from the set developed, because one may be more accurate under one set of circumstances while another is better in a different situation. Also, the better algorithms tend to require more computation, and some do not lend themselves to parallel implementation. Further study using real image data will be needed to determine the range of applicability of each algorithm.

We found a strong dependence of the accuracy of recovery of certain components of the motion on the size of the field of view. This is in concert with other reports describing difficulties with small fields of view, such as Adiv [1984] and Waxman and Wohn [1985].

Comments on sampling, filtering, and aliasing

Work with real image data has demonstrated the need to take care in filtering and sampling. The estimates of spatial gradient and time derivatives are sensitive to aliasing effects resulting from inadequate low-pass filtering before sampling. This is easily overlooked, particularly in the time direction. It is usually a mistake, for example, to simply pick every n^{th} frame out of an image sequence. At the very least, n consecutive frames should be averaged before sampling in order to reduce the high frequency components. One may object to the "smearing" introduced by this technique, but a series of widely separated snap-shots typically do not obey the conditions of the sampling theorem, and as a result the estimates of the derivatives may contain large errors.

This, of course, is nothing new, because the same considerations apply when one tries to estimate the optical flow using first derivatives of image brightness [Horn & Schunck 1981]. It is important to remember that the filtering must be applied before sampling—once the data has been sampled, the damage has been done.

The Brightness-Change Constraint Equation

Following Longuet-Higgins and Prazdny [1980] and Bruss and Horn [1983] we use a viewer-based coordinate system. Figure 1 depicts the system

under consideration. A world point

$$\mathbf{R} = (X, Y, Z)^T \tag{1}$$

is imaged at

$$\mathbf{r} = (x, y, 1)^T \,. \tag{2}$$

That is, the image plane has equation Z = 1. The origin is at the projection center and the Z-axis runs along the optical axis. The X- and Y-axes are parallel to the x- and y-axes of the image plane. Image coordinates are measured relative to the principal point, the point $(0,0,1)^T$ where the optical axis pierces the image plane. The points **r** and **R** are related by the perspective projection equation

$$\mathbf{r} = (x, y, 1)^T = \left(\frac{X}{Z}, \frac{Y}{Z}, \frac{Z}{Z}\right)^T = \frac{\mathbf{R}}{\mathbf{R} \cdot \hat{\mathbf{z}}}, \qquad (3)$$

with

$$Z = \mathbf{R} \cdot \hat{\mathbf{z}} \,, \tag{4}$$

and where $\hat{\mathbf{z}}$ denotes the unit vector in the Z direction.



Figure 1. The viewer-centered coordinate system. The translational velocity of the camera is $\mathbf{t} = (U, V, W)^T$, while the rotational component is $\boldsymbol{\omega} = (A, B, C)^T$.

Suppose the observer moves with instantaneous translational velocity $\mathbf{t} = (U, V, W)^T$ and instantaneous rotational velocity $\boldsymbol{\omega} = (A, B, C)^T$ relative to a fixed environment, then the time derivative of the vector \mathbf{R} can be written as

$$\mathbf{R}_t = -\mathbf{t} - \boldsymbol{\omega} \times \mathbf{R} \,. \tag{5}$$

440

The motion of the world point \mathbf{R} results in motion of the corresponding image point; the value of this *motion field* is given by

$$\mathbf{r}_t = \frac{d\mathbf{r}}{dt} = \frac{d}{dt} \left(\frac{\mathbf{R}}{\mathbf{R} \cdot \hat{\mathbf{z}}} \right) = \frac{\mathbf{R}_t (\mathbf{R} \cdot \hat{\mathbf{z}}) - (\mathbf{R}_t \cdot \hat{\mathbf{z}}) \mathbf{R}}{(\mathbf{R} \cdot \hat{\mathbf{z}})^2} \,. \tag{6}$$

This can also be expressed as

$$\mathbf{r}_t = \frac{\hat{\mathbf{z}} \times (\mathbf{R}_t \times \mathbf{r})}{\mathbf{R} \cdot \hat{\mathbf{z}}}, \qquad (7)$$

since $\mathbf{a} \times (\mathbf{b} \times \mathbf{c}) = (\mathbf{c} \cdot \mathbf{a})\mathbf{b} - (\mathbf{a} \cdot \mathbf{b})\mathbf{c}$. Substituting equation (5) into this result (see Negahdaripour and Horn [1987]) gives:

$$\mathbf{r}_{t} = -\hat{\mathbf{z}} \times \left(\mathbf{r} \times (\mathbf{r} \times \boldsymbol{\omega} - \frac{\mathbf{t}}{\mathbf{R} \cdot \hat{\mathbf{z}}}) \right).$$
(8)

In component form this can be expressed as

$$\mathbf{r}_{t} = \begin{pmatrix} x_{t} \\ y_{t} \\ 0 \end{pmatrix} = \begin{pmatrix} \frac{-U+xW}{Z} + Axy - B(x^{2}+1) + Cy \\ \frac{-V+yW}{Z} - Bxy + A(y^{2}+1) - Cx \\ 0 \end{pmatrix}, \qquad (9)$$

a result first obtained by Longuet-Higgins and Prazdny [1980].

This shows how, given the world motion, the motion field can be calculated for every image point. If we assume that the brightness of a small surface patch is not changed by motion, then expansion of the total derivative of brightness E leads to

$$\frac{\partial E}{\partial x}\frac{dx}{dt} + \frac{\partial E}{\partial y}\frac{dy}{dt} + \frac{\partial E}{\partial t} = 0.$$
(10)

(The applicability of the constant brightness assumption is discussed in Appendix A.) Denoting the vector $(\partial E/\partial x, \partial E/\partial y, 0)^T$ by $E_{\mathbf{r}}$ and $\partial E/\partial t$ by E_t , permits us to express this result more compactly in the form

$$E_{\mathbf{r}} \cdot \mathbf{r}_t + E_t = 0. \tag{11}$$

Substituting equation (8) into this result and rearranging gives

$$E_t - \left(\left((E_{\mathbf{r}} \times \hat{\mathbf{z}}) \times \mathbf{r} \right) \times \mathbf{r} \right) \cdot \boldsymbol{\omega} + \frac{\left((E_{\mathbf{r}} \times \hat{\mathbf{z}}) \times \mathbf{r} \right) \cdot \mathbf{t}}{\mathbf{R} \cdot \hat{\mathbf{z}}} = 0.$$
(12)

To simplify this expression we let

$$\mathbf{s} = (E_{\mathbf{r}} \times \hat{\mathbf{z}}) \times \mathbf{r} \tag{13}$$

 and

$$\mathbf{v} = -\mathbf{s} \times \mathbf{r} \,, \tag{14}$$

so equation (12) reduces to the brightness change constraint equation of Negahdaripour and Horn [1987], namely

$$\mathbf{v} \cdot \boldsymbol{\omega} + \frac{\mathbf{s} \cdot \mathbf{t}}{\mathbf{R} \cdot \hat{\mathbf{z}}} = -E_t \,. \tag{15}$$

The vectors \mathbf{s} and \mathbf{v} can be expressed in component form as

$$\mathbf{s} = \begin{pmatrix} -E_x \\ -E_y \\ xE_x + yE_y \end{pmatrix} \quad \text{and} \quad \mathbf{v} = \begin{pmatrix} +E_y + y(xE_x + yE_y) \\ -E_x - x(xE_x + yE_y) \\ yE_x - xE_y \end{pmatrix}.$$
(16)

Note that $\mathbf{s} \cdot \mathbf{r} = 0$, $\mathbf{v} \cdot \mathbf{r} = 0$, and $\mathbf{s} \cdot \mathbf{v} = 0$. These three vectors thus form an orthogonal triad. The vectors \mathbf{s} and \mathbf{v} are inherent properties of the image. Note that the projection of \mathbf{s} into the image plane is just the (negative) gradient of the image. Also, the quantity \mathbf{s} indicates the directions in which translation of a given magnitude will contribute maximally to the temporal brightness change of a given picture cell. The quantity \mathbf{v} plays a similar role for rotation.

Solving the Brightness Change Constraint Equation

Equation (15) relates observer motion $(\mathbf{t}, \boldsymbol{\omega})$, the depth of the world $\mathbf{R} \cdot \hat{\mathbf{z}} = Z(x, y)$, and certain measurable quantities of the image (\mathbf{s}, \mathbf{v}) . In general, it is not possible to solve for the first two of these given the last. Some interesting special cases are addressed in this chapter and in Negahdaripour and Horn [1987]; these are:¹

- 1 Known depth: We show that given Z, s and v, the quantities, t and ω can be calculated in closed form using a least-squares method.
- 2 **Pure rotation** $(||\mathbf{t}|| = 0)$: We show that given \mathbf{v} , the rotation vector $\boldsymbol{\omega}$ can be calculated in closed form.
- 3 **Pure translation or known rotation:** We present a least-squares method for determining **t**. Once **t** is known, the brightness change constraint equation can be used to find the depth at each picture cell:

$$Z = \mathbf{R} \cdot \hat{\mathbf{z}} = -\frac{\mathbf{s} \cdot \mathbf{t}}{E_t + \mathbf{v} \cdot \boldsymbol{\omega}} \,. \tag{17}$$

- 4 **Planar world:** Negahdaripour and Horn [1987] present a closed form solution for \mathbf{t} , $\boldsymbol{\omega}$ and the normal \mathbf{n} of the world plane.
- 5 **Quadratic patches:** Negahdaripour [1986] gives a closed-form solution in the case that a portion of the world can be represented as a quadratic form.

In this chapter we consider various integrals over an image region thought to correspond to a single rigid object in motion relative to the viewer. In the simplest case, the observer is moving relative to a static environment and the whole image can be used. The size of the field of view has a strong effect on the accuracy of the determination of the components of motion along the optical axis. When we need to estimate this accuracy, we will, for convenience, assume a circular image of radius r_v . This corresponds to a conical field of view with half angle θ_v , where $r_v = \tan \theta_v$, because we have assumed that the focal length equals one (we assume that $0 < \theta_v < \pi/2$).

We will show that the field of view should be large. Although orthographic projection usually simplifies machine vision problems, this is one case in which we welcome the effects of perspective "distortion"!

¹We do not discuss here related methods using optical flow, such as those of Bruss and Horn [1983].

Depth known

When depth is known, it is straightforward to recover the motion. (Depth may have been obtained using a binocular stereo system or some kind of range finder.) We cannot, in general, find a motion to satisfy the brightness change constraint equation at every picture cell, because of noise in the measurements. Instead we minimize

$$\iint \left[E_t + \mathbf{v} \cdot \boldsymbol{\omega} + (1/Z)\mathbf{s} \cdot \mathbf{t} \right]^2 \, dx \, dy \,. \tag{18}$$

Differentiating with respect to $\boldsymbol{\omega}$ and \mathbf{t} and setting the results equal to zero leads to the pair of vector equations

$$\left[\iint (1/Z)^2 \mathbf{s} \mathbf{s}^T \, dx \, dy\right] \mathbf{t} + \left[\iint (1/Z) \mathbf{s} \mathbf{v}^T \, dx \, dy\right] \boldsymbol{\omega}$$
$$= -\iint E_t (1/Z) \mathbf{s} \, dx \, dy$$
$$\left[\iint (1/Z) \mathbf{v} \mathbf{s}^T \, dx \, dy\right] \mathbf{t} + \left[\iint \mathbf{v} \mathbf{v}^T \, dx \, dy\right] \boldsymbol{\omega} = -\iint E_t \mathbf{v} \, dx \, dy \,. \tag{19}$$

This is a set of six linear equations in six unknowns with a symmetric coefficient matrix. (The equations can be solved by partitioning in order to reduce the computational effort.) The coefficients are all integrals of products of components of (1/Z)s and v. It may be useful to note that

$$\operatorname{Trace}(\mathbf{s}\mathbf{v}^{T}) = \operatorname{Trace}(\mathbf{v}\mathbf{s}^{T}) = \mathbf{s} \cdot \mathbf{v} = 0.$$
(20)

We could have obtained slightly different equations for $\boldsymbol{\omega}$ and \mathbf{t} if we had chosen to weight the integrand in equation (18) differently. We study the special case in which $\|\mathbf{t}\| = 0$ and the special case in which $\|\boldsymbol{\omega}\| = 0$ later.

One application of the above result is to "dynamic stereo." A binocular stereo system can provide disparity estimates from which 1/Z can be calculated. The above equations can then be used to solve for the motion, provided estimates of the derivatives of image brightness are also supplied. The correspondence problem of binocular stereo has, unfortunately, been found to be a difficult one. It would represent the major computational burden in a "dynamic stereo" system. We hope that motion vision research will eventually lead to simpler methods for recovering depth than those used for binocular stereo—although they are likely to be relatively inaccurate when based only on instantaneous translational and rotational velocity estimates.

Pure rotation

When $\|\mathbf{t}\| = 0$, the brightness change constraint equation reduces to

$$E_t + \mathbf{v} \cdot \boldsymbol{\omega} = 0. \tag{21}$$

We wish to find the value of $\boldsymbol{\omega}$ that minimizes the sum of the squares of the errors in the time derivative of brightness, that is, we want to minimize

$$\iint \left[E_t + \mathbf{v} \cdot \boldsymbol{\omega} \right]^2 \, dx \, dy \,. \tag{22}$$

Differentiating with respect to $\boldsymbol{\omega}$ and setting the result equal to zero gives us

$$2\iint [E_t + \mathbf{v} \cdot \boldsymbol{\omega}] \, \mathbf{v} \, dx \, dy = 0 \,. \tag{23}$$

Since $(\mathbf{v} \cdot \boldsymbol{\omega})\mathbf{v} = \mathbf{v}(\mathbf{v} \cdot \boldsymbol{\omega}) = (\mathbf{v}\mathbf{v}^T)\boldsymbol{\omega}$, we can write this in the form

$$\left[\iint \mathbf{v}\mathbf{v}^T \, dx \, dy\right] \boldsymbol{\omega} = -\iint E_t \mathbf{v} \, dx \, dy \,, \tag{24}$$

which is just a special case of equation (19). This is a set of three linear equations in the three unknown components of $\boldsymbol{\omega}$, namely, A, B, and C. The coefficient matrix is symmetric and only the right-hand side depends on the time derivative of brightness.

One can tell whether one is dealing with a case of pure rotation or not. In the presence of a translational component, equation (21) will not be a good approximation and so the integral in formula (22) will not be small. Experiments show that this simple method of determining rotation is robust and easy to implement. Slight variations are possible by weighting the intergrand differently. This method is reminiscent of the optical flow based method of Bruss and Horn [1983] and very similar to a method developed by Aloimonos and Brown [1984], to which our attention was drawn after we wrote this chapter.

Shown in figure 2 is every tenth frame out of a 40 frame sequence taken with a tripod-mounted CCD camera rotated manually about its vertical axis. The vertical component of the computed rotational velocity is shown in figure 3 as a function of the frame number. The units along the vertical axis are picture cells per time step in the center of the image (rather than say radians/second). After the initial acceleration, image components near the center of the image move by about 7 to 8 picture cells between successive frames. Three curves are given for varying amounts of image low-pass filtering and sub-sampling. The lowest curve (A) corresponds to the raw image data, and shows that for this particular scene at least, a motion of 7 to 8 picture cells is too much for accurate recovery of the angular velocity. The computed velocity appears to "saturate" at around 4 picture cells per time step. The next higher curve (B) corresponds to image compression by low-pass filtering and sub-sampling by a factor of two in each direction. In the compressed image sequence, the motion is in effect only about 3 to 4 picture cells per time step. The top curve (C) was obtained using images that were low-pass filtered and sub-sampled a second time to reduce them by a total of a factor of four in each direction. In this doubly compressed sequence, motion in the center of the image amounts to only about 1.5



Figure 2. Shown here are the $(a)10^{th}$, $(b)20^{th}$, $(c)30^{th}$, and $(d)40^{th}$ image out of a 40-image sequence obtained when a CCD camera mounted on a tripod was (manually) rotated about its vertical axis. After the initial acceleration, the image motion in the center is between 7 and 8 picture cells between successive frames. Image motion between frames is slightly larger in the corners of the image.



Figure 3. Recovered vertical component of the angular velocity vector as a function of the frame number. The angular velocity is given in picture cells of image displacement at the center of the image per time step. Curve A was obtained using the raw image data, curve B from a low-pass filtered and subsampled image sequence, and curve C from an image sequence that was low-pass filtered and sampled twice. Further low-pass filtering and subsampling produces essentially the same curve.

to 2 picture cells per time step, and the angular velocity is accurately recovered. Further filtering and sub-sampling leads to velocity estimates that are essentially the same as the ones obtained with this sequence.

\diamond Distribution of the directions of v

To understand the properties of the algorithm for recovering the instantaneous rotational velocity, one needs to study the matrix obtained by integrating \mathbf{vv}^T . We can think of the direction of \mathbf{v} as identifying a point on the unit sphere and of $||\mathbf{v}||$ as the mass of a particle placed there. The collection of vectors corresponding to an image region then can be thought of as a set of particles on the unit sphere. The integral of \mathbf{vv}^T is the symmetric 3×3 matrix whose elements are integrals of the nine pair-wise products of components of \mathbf{v} . This matrix is related to the inertia matrix of this set of particles. If the particles were spread uniformly over the surface of the sphere, this matrix would be the total mass times the identity matrix. As we show next, the particles are confined to a band, so this matrix, while diagonal on average, is not a multiple of the identity matrix.



Figure 4. A cross-section through the v-sphere defined by the image point r.

We know that $\mathbf{v} \cdot \mathbf{r} = 0$ and that the possible directions of \mathbf{r} lie within a cone defined by the field of view. For a particular value of \mathbf{r} , the equation $\mathbf{v} \cdot \mathbf{r} = 0$ defines a plane that cuts the unit sphere in a great circle (see figure 4). The vector \mathbf{v} must point in a direction corresponding to a point on this great circle. Since \mathbf{r} lies inside a cone of directions with half-angle θ_v , these great circles have axes that lie in this cone also. The collection of great circles lies in a band around the unit sphere of width equal to the total width of the visual field.

We can obtain the same result algebraically as follows: Let $\hat{\mathbf{r}}$, $\hat{\mathbf{v}}$, and $\hat{\mathbf{s}}$ be unit vectors in the directions \mathbf{r} , \mathbf{v} , and \mathbf{s} . Then, because \mathbf{r} , \mathbf{s} , and \mathbf{v} are mutually orthogonal,

$$(\hat{\mathbf{z}} \cdot \hat{\mathbf{r}})^2 + (\hat{\mathbf{z}} \cdot \hat{\mathbf{v}})^2 + (\hat{\mathbf{z}} \cdot \hat{\mathbf{s}})^2 = 1, \qquad (25)$$

while

$$(\hat{\mathbf{x}}\cdot\hat{\mathbf{v}})^2 + (\hat{\mathbf{y}}\cdot\hat{\mathbf{v}})^2 + (\hat{\mathbf{z}}\cdot\hat{\mathbf{v}})^2 = 1, \qquad (26)$$

where $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$, and $\hat{\mathbf{z}}$ are unit vectors in the X, Y, and Z directions. Subtracting the two equalities we obtain

$$(\hat{\mathbf{x}}\cdot\hat{\mathbf{v}})^2 + (\hat{\mathbf{y}}\cdot\hat{\mathbf{v}})^2 = (\hat{\mathbf{z}}\cdot\hat{\mathbf{r}})^2 + (\hat{\mathbf{z}}\cdot\hat{\mathbf{s}})^2, \qquad (27)$$

which, since

$$(\hat{\mathbf{z}} \cdot \hat{\mathbf{r}})^2 \ge \cos^2 \theta_v$$
 and $(\hat{\mathbf{z}} \cdot \hat{\mathbf{s}})^2 \ge 0$, (28)

tells us that

$$(\hat{\mathbf{x}} \cdot \hat{\mathbf{v}})^2 + (\hat{\mathbf{y}} \cdot \hat{\mathbf{v}})^2 \ge \cos^2 \theta_v \,. \tag{29}$$

Thus the directions of \mathbf{v} lie within an angle θ_v of the "equator" of the unit sphere. We call this band (shown in figure 5) the *permissible band*.



Figure 5. The permissible band on the v-sphere (front and rear views).

\diamond Estimate of the condition number

It is important to determine under what circumstances the recovery of the rotational velocity is ill conditioned, and whether the different components of the rotation vector are equally affected by noise in the data. To study these issues, one needs to estimate some of the properties of the coefficient matrix of the equations. We can get a rough idea of what the integral of \mathbf{vv}^T looks like by assuming that the particles corresponding to \mathbf{v} are spread *uniformly* within the permissible band. In Appendix B we show that the area of this band is

$$K = 4\pi \sin \theta_v \,. \tag{30}$$

The mixed inertia terms (such as the integral of the product of x and y) are all zero because of the symmetry of the band. The resulting 3×3 matrix is thus diagonal. Furthermore, if we let I_{xx} , I_{yy} , and I_{zz} be the diagonal terms then it can be shown (in Appendix B) that

$$I_{xx} = I_{yy} = (K/3) \left[1 + (1/2) \cos^2 \theta_v \right] , \qquad (31)$$

while

$$I_{zz} = (K/3)\sin^2\theta_v \,. \tag{32}$$

Thus I_{xx}/K and I_{yy}/K vary little with θ_v , while I_{zz}/K changes dramatically. The (L^2) condition number of this matrix—the ratio of the largest to the smallest eigenvalue—is just

$$\frac{I_{xx}}{I_{zz}} = \frac{1 + (1/2)\cos^2\theta_v}{\sin^2\theta_v} = \frac{3/2}{\sin^2\theta_v} - \frac{1}{2}.$$
(33)

This is very large when the field of view is small. When the condition number is very large, small errors in the right hand sides of the equations, or in the coefficients of the matrix itself, can lead to large errors in the solution.

In fact, the particles are *not* spread uniformly within the permissible band and do not have mass independent of position, so the above is only an estimate. We obtain the exact result later.

\diamond Stability of the solution method

The numerical stability of the solution for $\boldsymbol{\omega}$ is reduced when the condition number is large. In practice, the elements of the matrix of equation (24) will be corrupted by noise in the measurements, as will the right-hand side vector in this equation. The estimate of the third component of $\boldsymbol{\omega}$ will be affected more by these perturbations than the other two. Experiments confirm that the component of rotation about the optical axis is disturbed more by noise than the others. The ratio of the errors grows roughly as the inverse of the size of the field of view. This is not a peculiarity of our method, but applies in general.

It is intuitively obvious why this should be. Rotations about the xand y-axes produce motion fields that vary but little over the image. A small field of view can be used to estimate these components with almost the same accuracy as can a large field of view (provided the same number of picture cells used.) Rotation about the z-axis, on the other hand, produces a motion field that varies directly with distance from the principal point. Thus the maximum velocity depends on the size of the field of view. With a small field of view, the maximum velocity in the image will be small and relative errors in measurements correspondingly larger.

If an image region is used that is smaller than the whole field of view and perhaps off center, the analysis becomes more complex. In this case, the component of rotation about the direction towards the center of the region is less accurately known; the accuracy again decreasing with the size of the image region. This shows the futility of approaches based on data from small image patches, or higher derivatives of brightness at one point in the image. When working with very small image regions, the best one can do is to estimate the optical flow—there is no point in trying to recover the "rotation" about the center of the region.

\diamond Ensemble average of the integral of $\mathbf{v}\mathbf{v}^T$

The integral of \mathbf{vv}^T varies from image to image. It has already been suggested, however, that it will be approximately diagonal. We can obtain a more precise answer by averaging over an ensemble of images with all possible directions for the brightness gradient at each image point. We assume that different directions for the brightness gradient are equally likely. The result so obtained can be viewed in another way: it is the integral obtained in the limit from a textured image as the scale of the texture is made smaller and smaller. In this case we can arrange for every direction of the brightness gradient to be found in any small patch of the image. By suitable choice of the texture we can arrange that no directions of the brightness gradient is favored—all directions occur with equal frequency.

If we take into account the distribution of directions of \mathbf{v} and the weights $\|\mathbf{v}\|$, we find (in Appendix B) that

$$\overline{\iint \mathbf{v}\mathbf{v}^T \, dx \, dy} = k_v \begin{pmatrix} 1 + r_v^2/2 + r_v^4/6 & 0 & 0\\ 0 & 1 + r_v^2/2 + r_v^4/6 & 0\\ 0 & 0 & r_v^2/2 \end{pmatrix}, \quad (34)$$

where r_v is the diameter of the image and the constant k_v depends on the size of the field of view as well as the distribution of magnitudes of the brightness gradient. In practice we can easily find k_v since

Trace
$$\left(\iint \mathbf{v}\mathbf{v}^T \, dx \, dy\right) = \iint \operatorname{Trace}(\mathbf{v}\mathbf{v}^T) \, dx \, dy = \iint \mathbf{v} \cdot \mathbf{v} \, dx \, dy \,, \quad (35)$$

so

$$2k_v(1+3r_v^2/4+r_v^4/6) = \overline{\iint \mathbf{v} \cdot \mathbf{v} \, dx \, dy} \,. \tag{36}$$

Note that the condition number is

$$\frac{1+r_v^2/2+r_v^4/6}{r_v^2/2} = \frac{2}{r_v^2} + 1 + \frac{r_v^2}{3}.$$
(37)

It attains a minimum of $1+2\sqrt{2/3} = 2.633...$ when $r_v = \sqrt{\sqrt{6}} = 1.565...$. Thus the component of rotation about the optical axis is not recovered as accurately as the other two components, no matter how large the field of view. Also, as far as rotation is concerned, there is little advantage to making the field of view wider than a half-angle $\theta_v = \tan^{-1}\sqrt{\sqrt{6}} =$ 57.42... degrees, because the condition number reaches its minimum there.

Some simplifications of the method for recovering the rotational velocity based on the above analysis are discussed by us in Horn and Weldon [1987].

\diamond The v-Bar projection

We know that the directions of the vectors \mathbf{v} lie in the permissible band. But what about the vectors

$$\overline{\mathbf{v}} = -E_t \mathbf{v} \tag{38}$$

occurring in the integral on the right-hand side of equation (24)? We know that in the case of pure rotation $E_t = -\mathbf{v} \cdot \boldsymbol{\omega}$, so $\overline{\mathbf{v}} = (\mathbf{v} \cdot \boldsymbol{\omega})\mathbf{v}$. We conclude that

$$\overline{\mathbf{v}} \cdot \boldsymbol{\omega} = (\mathbf{v} \cdot \boldsymbol{\omega})^2 \ge 0.$$
(39)

Thus the directions of the vectors $\overline{\mathbf{v}}$ are confined to a hemisphere with $\boldsymbol{\omega}$ at its pole or "navel." We call this the *compatible hemisphere* for the case of pure rotation.

If the vectors $\overline{\mathbf{v}}$ covered this hemisphere uniformly, we could easily estimate $\boldsymbol{\omega}$ by finding the center of mass of the particles on the unit sphere corresponding to the values of $\overline{\mathbf{v}}$. The center of mass of a hemisphere of uniform density is at a point midway between the center of the sphere and the navel of the hemisphere, so we could use the formula

$$\boldsymbol{\omega} \approx \frac{2 \iint \boldsymbol{\overline{\nabla}} \, dx \, dy}{\iint \|\boldsymbol{\overline{\nabla}}\| \, dx \, dy} \,. \tag{40}$$

Unfortunately, the vectors $\overline{\mathbf{v}}$ do not cover the whole compatible hemisphere, because they are confined to the permissible band, in the same way as vectors \mathbf{v} . In fact, the vectors $\overline{\mathbf{v}}$ lie in the intersection of the permissible band and the compatible hemisphere, as shown in figure 6.



Figure 6. The v-sphere showing the intersection of the permissible band and the compatible hemisphere, in the case when the field of view is wide (front and rear views).



Figure 7. The intersection of the permissible band with three different compatible hemispheres when the field of view is narrow.

We can now see in another way why a small field of view reduces the accuracy with which we can estimate the component of rotation about the optical axis. If the field of view is small, the permissible band will be narrow, a mere ring. Our task is to guess which hemisphere cut the ring in half. This is easy when we are dealing with a band that covers almost the whole sphere—when it is very narrow, however, there is some uncertainty. A hemisphere claimed to provide a solution can easily be rotated about the line connecting the ends of the cut ring without significantly changing the intersection of the hemisphere and the ring as illustrated in figure 7. Thus the z-component of the direction to the navel of the hemisphere is uncertain in the presence of noise, or, when measurements from only a small number of picture cells are available.

We could use the geometric insight that the vectors \mathbf{v} all lie in the intersection of the permissible band and the compatible hemisphere to construct algorithms for recovering the direction of the vector $\boldsymbol{\omega}$. Some other method would then have to be found to estimate the magnitude of $\boldsymbol{\omega}$. However, we do not need to approach the problem in this way, in light of the leastsquares solution presented above. This geometric approach, however, *will* be fruitful in the case of pure translation where we find a similar geometric constraint and have no need to find the magnitude of the motion vector.

If there is a translational component to the motion, the points $\overline{\mathbf{v}}$ will not be confined to a hemisphere. This provides a convenient test to see whether the method presented above can be applied or not.

Rotation known or pure translation

If the rotation is known, perhaps measured by some other instrument, but depth is not, then the general problem reduces to the problem of pure translation. We can write

$$E'_t + (1/Z)\mathbf{s} \cdot \mathbf{t} = 0, \qquad (41)$$

where

$$E'_t = E_t + \mathbf{v} \cdot \boldsymbol{\omega} \,. \tag{42}$$

In the remainder of this section we do not distinguish between E_t and E'_t .

Note that equation (41) is not altered if we replace Z by kZ and t by kt. Thus we can recover motion and depth only up to a scale factor. In the sequel we will set ||t|| = 1 when convenient.

First, note that unlike the case in which depth is known, we cannot obtain a useful result by simply minimizing

$$\iint \left[E_t + (1/Z)\mathbf{s} \cdot \mathbf{t} \right]^2 \, dx \, dy \,, \tag{43}$$

because the integrand can be trivially made equal to zero at each point by the choice

$$Z = -(\mathbf{s} \cdot \mathbf{t})/E_t \,. \tag{44}$$

(However, this may produce negative values for Z, a fact that we exploit later.) Given the correct value of \mathbf{t} , the above equation provides a means for recovering depth, as already mentioned.

Equation (44) and the fact that depth must be positive leads to a simple upper bound on the depth at a particular point even when the direction of translational velocity is not known. Because Z > 0 we can write

$$Z = |Z| = |\mathbf{s} \cdot \mathbf{t}| / |E_t| , \qquad (45)$$

and so

$$Z \le \|\mathbf{s}\| \|\mathbf{t}\| / |E_t| . \tag{46}$$

The right-hand side here is the depth computed on the assumption that \mathbf{s} is parallel to \mathbf{t} . Of course this is only an upper bound, because Z will be much smaller if \mathbf{s} happens to be nearly orthogonal to \mathbf{t} . The bound is particularly poor, as a result, where \mathbf{r} is nearly parallel to \mathbf{t} , that is, near the focus of expansion (or the focus of compression) in the image.

\diamond Depth known—the case of pure translation

If we know the depth, as earlier, we can minimize the total error in the time derivative of brightness:

$$\iint \left[E_t + (1/Z)\mathbf{s} \cdot \mathbf{t} \right]^2 \, dx \, dy \tag{47}$$

by differentiating with respect to t. Setting the result equal to zero gives

$$\left[\iint (1/Z)^2 \mathbf{s} \mathbf{s}^T\right] \mathbf{t} \, dx \, dy = -\iint (1/Z) E_t \mathbf{s} \, dx \, dy \,, \tag{48}$$

which is just equation (19) with $\|\boldsymbol{\omega}\| = 0$. This is a set of three linear equations in the three components of \mathbf{t} (U, V, and W). The coefficient matrix is symmetric and only the right-hand side depends on the time derivative of brightness. Note that in equation (48) we attach less weight to information from points where Z is large.

The method is accurate if the correct values of depth are given. If estimates are used, the quality of the result will depend on the quality of the estimates. The accuracy of the result also depends on the size of the field of view, as we show later.

We get slightly different results if we weight the integrand in equation (47) differently. Multiplying by Z, for example, gives $[ZE_t + \mathbf{s} \cdot \mathbf{t}]$ for the integrand and

$$\left[\iint \mathbf{s}\mathbf{s}^T \, dx \, dy\right] \mathbf{t} = -\iint Z E_t \mathbf{s} \, dx \, dy \,, \tag{49}$$

for the solution. Alternatively multiplying by Z/E_t gives $[Z+(1/E_t){\bf s}\cdot{\bf t}]$ for the integrand and

$$\left[\iint (1/E_t)^2 \mathbf{s} \mathbf{s}^T \, dx \, dy\right] \mathbf{t} = -\iint (Z/E_t) \mathbf{s} \, dx \, dy \,, \tag{50}$$

for the solution. In this case we are minimizing the error in depth, rather than the error in the time derivative of brightness, as in equation (48).

The two alternate solutions given in Eqs. (49) and (50) have the advantage that the depth Z does not appear in the integrals on the left-hand side. This means that they are particularly well suited for iterative schemes where Z is re-estimated on each cycle. The solution of equation (49) has the further advantage that neither Z nor E_t appear on the left-hand side. This makes it easy to compute an ensemble average for this integral.

\diamond Distribution of the directions of s

To understand the properties of the above algorithms for recovering \mathbf{t} , we must examine the matrix obtained by integrating multiples of \mathbf{ss}^T . Once again, we can think of the direction of \mathbf{s} as identifying a point on the unit sphere and of a multiple of $||\mathbf{s}||$ as the mass of a particle placed there. The integral considered is related to the inertia matrix of the set of particles on the unit sphere.

Now just as the directions of \mathbf{v} lay in a band of width equal to the width of the field of view, because $\mathbf{v} \cdot \mathbf{r} = 0$, so do the directions of \mathbf{s} , because $\mathbf{s} \cdot \mathbf{r} = 0$. The distribution of points within the band is not quite the same, but we will ignore such details for now. First, assuming again a

uniform distribution within the permissible band, we get the same estimate of the condition number as previously, namely

$$\frac{1 + (1/2)\cos^2\theta_v}{\sin^2\theta_v} = \frac{3/2}{\sin^2\theta_v} - \frac{1}{2}.$$

Accuracy in the determination of W, the Z component of \mathbf{t} , will be reduced relative to that of the other two components when the field of view is small. Experiments confirm that for small fields of view, the estimate of the component of translation along the optical axis is disturbed more by noise than the other two. Hence a wide field of view is called for.

The integral of \mathbf{ss}^T varies from image to image. In order to better understand the matrix defined by \mathbf{ss}^T , we would like to examine a "typical" image. Because it is difficult to define such an image, instead, we can take an average over an ensemble of images containing all possible directions for the brightness gradient at each image point. If we take into account the distribution of directions of \mathbf{s} and the weights $\|\mathbf{s}\|$, we find (in Appendix B) that

$$\overline{\iint \mathbf{ss}^T \, dx \, dy} = k_s \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & r_v^2/2 \end{pmatrix} , \tag{51}$$

where r_v is the radius of the image and the constant k_s depends on the size of the field of view as well as the distribution of magnitudes of the brightness gradient. In practice we can find k_s by noting that

Trace
$$\left(\iint \mathbf{ss}^T \, dx \, dy\right) = \iint \operatorname{Trace}(\mathbf{ss}^T) \, dx \, dy = \iint \mathbf{s} \cdot \mathbf{s} \, dx \, dy \,, \quad (52)$$

 \mathbf{SO}

$$2k_s\left(1+r_v^2/4\right) = \overline{\iint \mathbf{s} \cdot \mathbf{s} \, dx \, dy} \,. \tag{53}$$

Note that the condition number is just $\min(r_v^2/2, 2/r_v^2)$ which reaches a minimum of 1 when $r_v = \sqrt{2}$. In the case of pure translation, the component of translation along the optical axis is found with more accuracy than the other two when the field of view has a half-angle wider than $\theta_v = \tan^{-1}\sqrt{2} = 54.74...$ degrees, because $r_v^2/2$ then is larger than one.

Some simplifications of the method for recovering the translational velocity based on the above analysis are discussed by us in Horn and Weldon [1987].

Translation with rotation known

In this section we deal with the problem of determining the direction of translation and depth Z(x, y) given the rotation vector $\boldsymbol{\omega}$.



Figure 8. The intersection of the permissible band on the s-shape and the compatible hemisphere for two cases. (a) Focus of expansion within field of view. (b) Focus of expansion outside field of view.

\diamond The importance of a wide field of view

In the general case, the need for a wide field of view is very clear. In a small image region near the center of the image, for example, rotation about the y-axis looks the same as translation along the x-axis, while rotation about the x-axis looks the same as translation along the (negative) y-axis. As is well known in stereo-photogrammetry, a large field of view is needed to separate these components of the transformation between two camera positions [Wolf 1983; Gosh 1972].

If we take note of this ambiguity, and the uncertainty with which the components of rotation and translation along the optical axis can be determined, we see that locally, out of six parameters, only two combinations can be estimated. These two quantities are just the components of the motion field. The same argument can be made for points at some distance from the principal point of the image.

There is a difference between the case when the motion is predominantly along the optical axis and the case where it is predominantly parallel to the image plane. The transition between the two situations occurs when the direction of the vector **t** moves outside the cone of directions of the field of view, that is, when the focus of expansion (or compression) moves outside the image. When the focus of expansion is inside the image, then the great circle defined by $\mathbf{s} \cdot \mathbf{t} = 0$ lies entirely inside the permissible band on the unit sphere (figure 8a). The measured values of **s** then provide constraint all the way around the great circle. Conversely, when the focus of expansion lies outside the image, the great circle cuts the permissible band (see figure 8b).

In this case the known values of s provide constraint only along two segments of the great circle. These segments get shorter and shorter as the vector t becomes more and more parallel of the image plane. It should be clear that in this case the direction of the vector t can be determined with somewhat lower accuracy than when the focus of expansion is near the principal point.

\diamond The s-Bar projection

The integrals on the right hand side of the equations for \mathbf{t} developed earlier contain positive multiples of the vector

$$\bar{\mathbf{s}} = -\mathrm{sign}(E_t)\mathbf{s}\,.\tag{54}$$

(Here we only care about the directions of the vectors, so we ignore scale factors.) Now in the case of translation with known rotation, we have (from equation (41)) $E_t = -(1/Z) \mathbf{s} \cdot \mathbf{t} ,$

and

$$\overline{\mathbf{s}} \cdot \mathbf{t} = (1/Z)\operatorname{sign}(\mathbf{s} \cdot \mathbf{t}) \,\mathbf{s} \cdot \mathbf{t} = (1/Z) \,|\mathbf{s} \cdot \mathbf{t}| \ge 0 \,, \tag{55}$$

because Z > 0. We are only interested at this point in the sign of $\overline{\mathbf{s}} \cdot \mathbf{t}$, so we can use any convenient positive multiple of $\overline{\mathbf{s}}$ such as

$$-(1/E_t)\mathbf{s},$$
 $-\operatorname{sign}(E_t)\mathbf{s},$ or $-E_t\mathbf{s}$

in the discussion that follows.

Equation (55) states that $\bar{\mathbf{s}}$ can only lie in the hemisphere that has \mathbf{t} as its navel. We call this the *compatible hemisphere* in the case of translation with known rotation. Because $\bar{\mathbf{s}}$ is a multiple of \mathbf{s} , it must also lie in the permissible band. Thus $\bar{\mathbf{s}}$ can only lie in the intersection of the permissible band and the compatible hemisphere. We will exploit this geometric insight shortly.

Our task can be viewed as that of finding the hemisphere that contains all of the directions specified by the vectors $\overline{\mathbf{s}}$ derived from the image. Note that the solution may not be unique and that there may not be any solution. Later we will modify the problem definition somewhat to deal with these possibilities.

If there is a rotational component to the motion, then the points $\overline{\mathbf{s}}$ will not be confined to a hemisphere. This provides a convenient test to see whether the methods presented here can be applied or not.

\Diamond Motion determination as a linear programming problem

We wish to find a vector **t** that makes $\overline{\mathbf{s}} \cdot \mathbf{t} \ge 0$ at all image points. We can think of this as a gigantic linear programming problem.² There are three unknowns and one inequality for every picture cell. (Actually, because we do not care about the magnitude of **t**, there are only two degrees of freedom.)

Because we do not have a criterion function to be extremized, we will have an infinite number of solutions—if there are any solutions at all. All of these solutions will lie in a convex polygon on the unit sphere. The sides of this polygon are portions of great circles corresponding to constraints which we will call *critical constraints* (see figure 9). With data from a large number of cells we expect this solution polygon to be small. We may choose its center as the "best" solution.



Figure 9. The s-sphere showing great circles for many image points. Circles for critical points (emphasized) constrain location of t.

²We do not wish to suggest, by the way, that linear programming algorithms could be fruitfully applied to this problem.

Typically, the solution polygon will have relatively few sides. Thus data from a small number of *critical picture cells* fully constrains the solution. First, note that each side of this polygon corresponds to an equality of the form $\mathbf{\bar{s}} \cdot \mathbf{t} = 0$ for some picture cell. From the brightness change constraint equation we know that $E_t = 0$ when $\mathbf{\bar{s}} \cdot \mathbf{t} = 0$. Thus the critical constraints are provided by picture cells where E_t is small (and $\mathbf{\bar{s}}$ is not). This is an important observation, which can be used to reduce the size of the linear programming problem; we simply disregard the inequalities arising from picture cells where E_t is large.

There is a class of points for which $\overline{\mathbf{s}} \cdot \mathbf{t}$ is arbitrary, even though E_t is small and $\overline{\mathbf{s}}$ is not; these are image points for which Z is large. Such points provide "false" constraints on \mathbf{t} . For a practical system, some means must be found for identifying these points. One way of doing this, for images with large depth range, is based on the following observation. In a real image, regions for which Z is large, that is, the background, tend to encompass a significant area with all points in the area having $E_t \approx 0$. On the other hand, points with $\mathbf{s} \cdot \mathbf{t} = 0$ and $\overline{\mathbf{s}}$ large are usually isolated and surrounded by regions for which $E_t \neq 0$. The above difficulty appears in all of the methods of determining motion; it is harder to determine \mathbf{t} when the depth range is large.

We observe in passing that the points that are most useful in constraining the translational motion vector are the very same points where it is difficult to calculate depth accurately! One may also make the observation that a method for segmenting the scene into foreground and background regions would be very useful in the case of general motion, because the background regions can then be used to recover the rotationally component. The known rotational velocity can then, in turn, be used to recover the translational component from the foreground regions.

The linear programming method of determining t discussed above uses relatively little of the image data. In fact, only points at the edge of the compatible hemisphere influence the solution at all. While this is a sensible procedure if the data is trustworthy, it will be quite sensitive to noise. For noisy, that is, real images, it may be worthwhile to consider other points, such as those in a band for which E_t is less than a small cutoff value.

\diamond The perceptron "learning" algorithm

One way of finding the solution of a large number of homogeneous inequalities is by means of the iterative perceptron "learning" algorithm [Minsky & Papert 1969; Duda & Hart 1973]. Given a set of vectors $\{\bar{\mathbf{s}}_i\}$, this procedure is guaranteed to find a vector \mathbf{t} that satisfies $\bar{\mathbf{s}}_i \cdot \mathbf{t} \geq 0$, if such a vector exists. It even does this in a finite number of steps, provided there exists some ϵ such that $\overline{\mathbf{s}}_i \cdot \mathbf{t} > \epsilon$ for all $\overline{\mathbf{s}}_i$ in the given set (which almost always happens when the set is finite).

The idea is to start with some non-zero vector \mathbf{t}^0 and to test whether the inequalities are satisfied. (A reasonable choice for \mathbf{t}^0 is one of the vectors $\bar{\mathbf{s}}_i$.) If the inequality is not satisfied for a particular vector in the set, then the smallest adjustment is made to make the dot-product zero. (Note that this may disturb inequalities that have passed the test already). Suppose that the present estimate for the direction of the translation vector is the direction of \mathbf{t}^n . We now test the dot-product $\bar{\mathbf{s}}_i \cdot \mathbf{t}^n$. If it is negative, we adjust our estimate of the vector \mathbf{t} according to the rule

$$\mathbf{t}^{n+1} = \mathbf{t}^n + \delta \mathbf{t}^n \,, \tag{56}$$

where

$$\delta \mathbf{t}^n = -\frac{\mathbf{s}_i \cdot \mathbf{t}^n}{\mathbf{s}_i \cdot \mathbf{s}_i} \mathbf{s}_i \,. \tag{57}$$

Note that $\overline{\mathbf{s}}_i \cdot \mathbf{t}^{n+1} = 0$ and that the magnitude of \mathbf{s}_i does not matter. (Also, the test above can be replaced with a test that checks whether $-\mathbf{s}_i \cdot \mathbf{t}^n$ has the same sign as E_t .)

If the inequalities are inconsistent, that is, if the $\bar{\mathbf{s}}_i$ are not confined to a hemisphere (or nearly so), as will happen in practice due to noise, the algorithm will not converge. Furthermore, there is no guarantee that the guess at any stage is particularly good. We discuss several simple refinements that can help in this case in Horn and Weldon [1987].

The vector \mathbf{t}^n in the perceptron "learning" algorithm is obviously a linear combination of vectors drawn from the set $\{\bar{\mathbf{s}}_i\}$. Vectors in this set have directions that correspond to points in the permissible band. Now suppose that this band is very narrow. Then, to build a vector with a significant z-component one has to add many of these vectors. In order to keep the x- and y-components small, these vectors must almost come in pairs from opposite ends of the narrow band. Not surprisingly, the algorithm performs rather poorly in this situation; it is much happier with vectors sprinkled uniformly in direction over a full hemisphere.

It should also be noted that in a "real-time" application, we do not expect the velocity estimates to change rapidly. Thus the previous value of the velocity is likely to be an excellent first estimate for the current value of t. This means that very few iterations will be needed to get an acceptable new value. A considerable amount of computation can be saved this way, just as it can in the computation of the optical flow [Horn & Schunck 1981]. We discuss a "parallel" perceptron algorithm in Horn and Weldon [1987].

\diamond Minimizing the integral of Z^2

In this section we assume that the depth range Z_{max}/Z_{min} is finite. This will generally be the case in robotic applications. The method discussed in

this section can also be applied to images in which the background has very large Z, if, as discussed earlier, these regions are excised from the image before the motion vector is calculated.

We have seen that we can compute depth when the motion \mathbf{t} is known using equation (44)

$$Z = -(1/E_t)\mathbf{s} \cdot \mathbf{t} \,.$$

Now if we use the *wrong* value \mathbf{t}' in this formula, we get the wrong depth value:

$$Z' = -(1/E_t)\mathbf{s} \cdot \mathbf{t}' = Z(\mathbf{s} \cdot \mathbf{t}')/(\mathbf{s} \cdot \mathbf{t}).$$
(58)

We expect only positive values for Z, but this formula may give us negative values, because $(\mathbf{s} \cdot \mathbf{t}')$ may be negative where $(\mathbf{s} \cdot \mathbf{t})$ is positive and vice versa. More interestingly, we may obtain very large values for Z (both positive and negative), because $(\mathbf{s} \cdot \mathbf{t})$ may be almost zero while $(\mathbf{s} \cdot \mathbf{t}')$ is not. That is, the magnitude of Z will often be very large near points where $E_t \approx 0$. We may conclude that we could determine the correct value for \mathbf{t} by minimizing the integral of Z^2 over the image, that is by minimizing the quadratic form

$$\iint (1/E_t)^2 (\mathbf{s} \cdot \mathbf{t})^2 \, dx \, dy = \mathbf{t}^T \left[\iint (1/E_t^2) \mathbf{s} \mathbf{s}^T \, dx \, dy \right] \mathbf{t} \,, \tag{59}$$

subject to the constraint $||\mathbf{t}|| = 1$. The solution is the eigenvector of the real symmetric 3×3 matrix

$$M = \iint (1/E_t^2) \mathbf{s} \mathbf{s}^T \, dx \, dy \tag{60}$$

associated with the smallest eigenvalue. We can prove this by minimizing the sum

$$S = \mathbf{t}^T M \mathbf{t} + \lambda (1 - \mathbf{t}^T \mathbf{t}), \qquad (61)$$

where λ is a Lagrangian multiplier. Then

$$\frac{\partial S}{\partial \mathbf{t}} = 2M\mathbf{t} - 2\lambda\mathbf{t} = 0, \qquad (62)$$

which yields

$$M\mathbf{t} = \lambda \mathbf{t} \,. \tag{63}$$

Thus λ is an eigenvalue of M, and **t** is the corresponding eigenvector. Substituting equation (63) into equation (61) gives the results $S = \lambda$. Thus $\mathbf{t}^T M \mathbf{t}$ is minimized by taking the smallest of the three eigenvalues of M for λ .

To minimize problems due to noise, we can add a small positive constant to E_t^2 commensurate with the expected noise in E_t^2 . That is, we take as our solution the eigenvector of

$$M' = \iint \frac{1}{E_t^2 + n^2} \mathbf{s} \mathbf{s}^T \, dx \, dy \tag{64}$$

associated with the smallest eigenvalue.

If e is an eigenvector, so is -e. But we want Z to be positive. Rather than test this condition at a single point, we compute an average like

$$\overline{\mathbf{s}}_0 = -\iint (1/E_t) \mathbf{s} \, dx \, dy \qquad \text{or} \qquad \overline{\mathbf{s}}_0 = -\iint \frac{E_t}{E_t^2 + n^2} \mathbf{s} \, dx \, dy \qquad (65)$$

and check whether

$$\overline{\mathbf{s}}_0 \cdot \mathbf{t} > 0. \tag{66}$$

If it is not, we simply change the sign of the solution \mathbf{t} .

As before, we may choose to weight the integral of equation (59) according to some measure of how trustworthy the data from each picture cell is.

The method presented in this section produces an estimate of the translation vector \mathbf{t} in closed form and with high accuracy. Of course, a cubic must be solved to obtain the eigenvalues—but there is an analytic method for doing that. The corresponding eigenvectors can then be found by taking cross products of two rows of a 3×3 matrix.

The preceding method of calculating **t** has another justification that some readers might find more persuasive. From equation (41) we know that $\mathbf{s} \cdot \mathbf{t} \approx 0$ for points with $E_t \approx 0$ (again ignoring background points). Thus we are basically looking for a vector **t** that makes $\mathbf{s} \cdot \mathbf{t} \approx 0$ whenever $E_t \approx 0$. The points where the time derivatives are small provide most constraint, as already discussed. We could try to minimize something like

$$\iint_C (\mathbf{s} \cdot \mathbf{t})^2 \,, \tag{67}$$

where C is the set of image points where $E_t \approx 0$. Rather than use a strict cut-off on E_t , we may consider a weighting scheme in an integral like

$$\iint w(\mathbf{s} \cdot \mathbf{t})^2 \, dx \, dy \,, \tag{68}$$

over the whole image where the weighting function w is chosen to emphasize points where $E_t \approx 0$. A reasonable choice, $w = 1/(E_t^2 + n^2)$, leads to integral given in equation (64). The eigenvector corresponding to the smallest eigenvalue is a normal of the plane that best fits the weighted set of points (see figure 10).

If there is a rotational component to the motion then the vectors \bar{s} where E_t is small will not lie near a great circle. In this case the smallest eigenvalue will not be small. This provides a convenient test. We discuss a method that avoids the need to find eigenvalues and eigenvectors in Horn and Weldon [1987]. Related methods for finding the focus of expansion are discussed in Negahdaripour and Horn [1987].

Figure 11 shows a scatter plot of positions on the unit sphere for \mathbf{t} recovered from noisy synthetic data. Each estimate is based on brightness gradient at 200 picture cells with 1% noise in the derivatives of brightness. Note the elongation of the cluster of points in a direction parallel to the



Figure 10. The great circle corresponding to the motion t which best fits the points on the s-sphere for which $E_t \approx 0$.

optical axis. When tens of thousands of picture cells are used, instead of hundreds, the algorithm can tolerate considerably more noise. While further experimentation is called for, we found that this algorithm behaves at least as well, if not better than the others we have investigated.

Conclusions

We have developed methods for recovering motion directly from the first derivatives of brightness in an image region in the cases of pure rotation and pure translation (and general motion when the rotational component is known). We have tested these methods on synthetic image data and, to a limited extent, on some kinds of real image sequences. In the case of pure rotation we give an exact simple solution to the obvious least-squares problem. In the case of pure translation we give several methods with different trade-offs between accuracy, noise-sensitivity and computational expense. While we have preliminary ideas about the relative merits of these methods, detailed conclusions will have to await further careful experimentation with real images. Some further results on both synthetic and real image



Figure 11. Plot of several noisy estimates of the translation vector \mathbf{t} on the s-sphere (200 pixels/estimate, 1% noise in brightness measurements).

data are reported in Horn and Weldon [1987] and Negahdaripour and Horn [1987].

We show that it is trivial to recover depth when the motion is known and that it is trivial to recover the motion when depth is known. We emphasize the importance of a large field of view and point out difficulties arising in the pure translation case when there is a very large depth range. We also note that image points where the brightness derivative is small provide most constraint on the translation vector while the depth at these points is hard to recover. We also show that it is difficult to recover the translational motion towards, and rotational motion about, the line connecting the projection center to the image region of interest, when that region is small. We emphasize the need for adequate low-pass filtering in both spatial and time dimension before sampling in order to ensure that estimates of derivatives are accurate.

The discussion is facilitated by introduction of the auxiliary vectors \mathbf{s} and \mathbf{v} . The directions of these vectors have been shown to be constrained to lie in the intersection of a "permissible band" and a "compatible hemisphere" on the unit sphere. These geometric concepts help lend intuitive support to the algebraic results.

References

- Adiv, G. [1984], "Determining 3-D Motion and Structure from Optical Flow Generated by Several Moving Objects," COINS TR 84-07, Computer and Information Science, University of Massachusetts, Amherst, MA.
- Aloimonos, Y., and C. Brown [1984], "Direct Processing of Curvilinear Motion from a Sequence of Perspective Images," Proceedings of the Workshop on Computer Vision, Representation and Control, Annapolis, Maryland.
- Bruss, A. R., and B. K. P. Horn [1983], "Passive Navigation," Computer Vision, Graphics, and Image Processing, vol. 21, pp. 3-20.
- Duda, R. O., and P. E. Hart [1973], Pattern Classification and Scene Analysis, John Wiley, New York.
- Gosh, S. K. [1972], *Theory of Stereophotogrammetry*, Ohio State University Bookstores.
- Horn, B. K. P., and B. G. Schunck [1981], "Determining Optical Flow," Artificial Intelligence, vol. 17, pp. 185–203.
- Horn, B. K. P., and E. J. Weldon Jr. [1987], "Computationally efficient methods of recovering translational motion," *International Conference on Computer* Vision, London, England, pp. 8-11.
- Longuet-Higgins, H. C., and K. Prazdny [1980], "The Interpretation of a Moving Retinal Image," Proc. of Royal Society of London, Series B, vol. 208, pp. 385-397.
- Minsky, M., and S. Papert [1969], Perceptrons: An Introduction to Computational Geometry, MIT Press, Cambridge, MA.
- Negahdaripour, S. [1986], "Direct Methods for Structure from Motion," Ph.D. Thesis, Mechanical Engineering, Massachusetts Institute of Technology.
- Negahdaripour, S., and B. K. P. Horn [1987], "A Direct Method for Locating the Focus of Expansion," Report AIM-939 Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge.
- Negahdaripour, S., and B. K. P. Horn [1987], "Direct Passive Navigation," IEEE Trans. Pattern Analysis and Machine Intelligence, vol. 9, no. 1, pp. 168–176.
- Waxman, A. M., and K. Wohn [1985], "Contour Evolution, Neighborhood Deformation, and Global Image Flow: Planar Surfaces in Motion," *The International Journal of Robotics Research*, vol. 4, no. 3, pp. 95–108.
- Waxman, A. M., and S. Ullman [1985], "Surface Structure and 3-D Motion from Image Flow: A Kinematic Analysis." The International Journal of Robotics Research, vol. 4, no. 3, pp. 72–94.

Wolf, P. R. [1983], Elements of Photogrammetry, McGraw-Hill, New York.

The research in this chapter was supported by the National Science Foundation under Grant No. DMC85-11966. Additional support was provided by NASA (Grant No. GSFC 5-1162) and by the Veteran's Administration.

The authors would like to thank the other members of the Waikiki Surfing and Computer Vision Society, and especially Shahriar Negahdaripour, for their contributions to this chapter.

Appendix A: The Brightness Change Constraint Equation

The brightness change constraint equation is based on the assumption that the brightness of the image of a patch on the surface does not change as the observer moves relative to the surface. Expansion of the total derivative in the equation dE/dt = 0 by means of the chain-rule leads to the constraint equation

$$\frac{\partial E}{\partial x}\frac{dx}{dt} + \frac{\partial E}{\partial y}\frac{dy}{dt} + E_t = 0, \qquad (A1)$$

or

$$uE_x + vE_y + E_t = 0, \qquad (A2)$$

where E_x , E_y , and E_t are the partial derivatives of brightness with respect to x, y, and t, while u and v are the time derivatives of x and y.

In practice, the brightness of a patch rarely remains exactly the same. The brightness change constraint equation is nevertheless a useful approximation, as long as the change in brightness at an image point due to the motion is much larger than the change in brightness due to other effects, such as change in viewing direction or illumination. This will be the case as long as there is good contrast at high spatial frequencies, as will be shown next.

Suppose that the brightness of a patch does in fact change due to changes in viewing direction or changes in illumination. In most cases the rate of change of brightness will be relatively small. Let us say that $dE/dt = \epsilon$, and so

$$E_t = -(uE_x + vE_y) + \epsilon. \tag{A3}$$

Consider now a simple grating pattern in the image that, at a particular time, is described by the equation

$$E = E_0 (1 + \sin(ax + by)). \tag{A4}$$

Then the components of the brightness gradients are

$$E_x = aE_0\cos(ax + by)$$

$$E_y = bE_0\cos(ax + by).$$
(A5)

Consequently

$$uE_x + vE_y = (au + bv)E_0\cos(ax + by).$$
(A6)

It is clear that the error in E_t , the rate of change of brightness at a point in the image, resulting from changes in the brightness of the surface, is relatively small, as long as $(au + bv)E_0$ is large compared to ϵ . (This term, $(au + bv)E_0$, will be zero when the image motion happens to be parallel to the ridges of the grating. In practice, however, surface markings will contain many spatial frequency components and most of these will not be aligned in this special way). We conclude that the relative error in the rate of change of brightness with time is small, as long as there is significant contrast at the higher spatial frequencies.

Chapter 40 Direct Methods for Recovering Motion

The approximation breaks down when the surface markings are weak and the changes of brightness due to changes in viewing direction or illumination are rapid. This happens, for example, in the rare situation where a specular surface momentarily lines up exactly to reflect the light from a point source towards the viewer. It also happens when an object moving relative to a point source enters a cast shadow.

A number of additional factors help keep the relative error in E_t small. First of all, some surfaces have the property that they appear equally bright from all viewing directions. A Lambertian surface is a very special case of this, where the brightness varies as the cosine of the incident angle. The image brightness is not affected at all by the motion of the observer when a surface of this type is fixed relative to the light source. While most real surfaces do not appear equally bright from all viewing directions, brightness typically varies slowly with changes in observer position (slowly enough that we are usually not aware of any such changes).

Brightness variations resulting from changes in surface orientation are most severe when there is a single point source. These variations are reduced when there are multiple sources or an extended source. In the extreme case of a scene illuminated from all sides, for example, image brightness does not depend on surface orientation at all, even if the surface is specular!

Similarly, a lens occupying a large solid angle, as seen from the object, will smooth out changes in brightness resulting from changes in viewer position. One can see this easily in the extreme case of a glossy reflection, which will be seen only over a small range of positions if a small lens or pin-hole is used. A large entrance aperture on the other hand will smear out the highlight effect over a larger range of viewing positions. This is not a big help in many imaging situations, however, because objects are far from the sensor relative to the size of the sensor, except in the case of microscopy.

To summarize: The brightness of the image of a patch may change somewhat as the observer moves relative to the surface. The brightness change constraint equation nevertheless provides a good way of estimating the rate of change of brightness with respect to time at a point in the image. The relative error in this estimate will be small when there is significant contrast in the surface markings at higher spatial frequencies. (There will be no error at all when the surface appears equally bright from all viewing directions and the object does not move relative to the light source).

Appendix B: Ensemble Averages

Some of the integrals that appear in this chapter, while functions of the scene content, tend to lie close to "average" values when evaluated over sufficiently large textured regions. These average values can be useful in two different enterprises:

- Analyzing the relative stability of the components of the solution.
- Developing simplified methods for recovering the solution (as shown in Horn and Weldon [1987]).

In order to compute these averages we have to make some assumptions about the probability distribution of brightness gradients. We assume here that this distribution is rotationally symmetric and independent of image position. That is, on average we see the same brightness gradients at every image point and all directions of the brightness gradient are equally likely. The distribution of magnitude of this gradient is left arbitrary, however, because it does not directly affect the main results.

Moment integrals for the uniform band

Before we start, let us quickly obtain the equivalent results under the assumption that data points (s or v) are uniformly distributed over the permissible band. Let η denote the latitude and ξ the longitude on the unit sphere. We see that the area of the band is just

$$K = \int_{-\pi}^{\pi} \int_{-\theta_v}^{\theta_v} \cos\eta \, d\eta \, d\xi = 4\pi \sin\theta_v \,. \tag{B1}$$

The Cartesian coordinates of a point on the unit sphere are given by $\mathbf{p}=(x,y,z)^T$ where

$$\begin{aligned} x &= \cos \eta \cos \xi \,, \\ y &= \cos \eta \sin \xi \,, \\ z &= \sin \eta \,. \end{aligned}$$
 (B2)

Let the integral of $\mathbf{p}\mathbf{p}^T$ be

$$\int_{-\pi}^{\pi} \int_{-\theta_v}^{\theta_v} \mathbf{p} \mathbf{p}^T \, d\eta \, d\xi = \begin{pmatrix} I_{xx} & I_{xy} & I_{xz} \\ I_{yx} & I_{yy} & I_{yz} \\ I_{zx} & I_{zy} & I_{zz} \end{pmatrix} \,. \tag{B3}$$

Then

$$I_{xx} = \int_{-\pi}^{\pi} \int_{-\theta_{v}}^{\theta_{v}} x^{2} \cos \eta \, d\eta \, d\xi = \int_{-\pi}^{\pi} \cos^{2} \xi \, d\xi \int_{-\theta_{v}}^{\theta_{v}} \cos^{3} \eta \, d\eta \,, \qquad (B4)$$

that is,

$$I_{xx} = \frac{\pi}{3} \left[3\sin\theta_v + (1/3)\sin 3\theta_v \right] = \frac{4\pi}{3}\sin\theta_v \left[1 + (1/2)\cos^2\theta_v \right] , \quad (B5)$$

while

$$I_{zz} = \int_{-\pi}^{\pi} \int_{-\theta_v}^{\theta_v} z^2 \cos \eta \, d\eta \, d\xi = \int_{-\pi}^{\pi} d\xi \int_{-\theta_v}^{\theta_v} \sin^2 \eta \cos \eta \, d\eta \,, \qquad (B6)$$

that is,

$$I_{zz} = \frac{4\pi}{3}\sin^3\theta_v \,. \tag{B6}$$

Also, by symmetry, $I_{yy} = I_{xx}$ and the off-diagonal terms, I_{xy} , I_{yz} , and I_{zx} , are all zero. We have

$$I_{xx} = (K/3) \left[1 + (1/2) \cos^2 \theta_v \right] = I_{yy}, \qquad (B7)$$

 and

$$I_{zz} = (K/3)\sin^2\theta_v , \qquad (B8)$$

 \mathbf{so}

$$I_{xx} + I_{yy} + I_{zz} = K. (B9)$$

The moment matrix is diagonal and so I_{xx} , I_{yy} and I_{zz} are the three eigenvalues. The condition number is the ratio of the largest to the smallest or

$$\frac{1+(1/2)\cos^2\theta_v}{\sin^2\theta_v}\,.\tag{B10}$$

These results give us a quick estimate of the ensemble averages of the integrals if \mathbf{ss}^T and \mathbf{vv}^T . To do better, we have to take into account the actual distribution of \mathbf{s} and \mathbf{v} in the permissible band.

Ensemble average of the integral of ss^T

It is convenient to use polar coordinates in the case of a circular image. We have

$$\begin{aligned} x &= r \cos \theta & -\pi \le \theta \le +\pi \,, \\ y &= r \sin \theta & 0 \le r \le r_v \,. \end{aligned}$$
 (B11)

Similarly, we may use polar coordinates for the brightness gradient

$$E_x = \rho \cos \phi \qquad -\pi \le \phi \le +\pi ,$$

$$E_y = \rho \sin \phi \qquad 0 \le \rho .$$
(B12)

Let the probability of seeing a brightness gradient with magnitude lying between ρ and $\rho + \delta \rho$ be $2\pi \rho P(\rho) \delta \rho$. Now

$$\mathbf{s} = \begin{pmatrix} -E_x \\ -E_y \\ xE_x + yE_y \end{pmatrix} = \rho \begin{pmatrix} -\cos\phi \\ -\sin\phi \\ r\cos(\theta - \phi) \end{pmatrix}.$$
(B13)

Consequently

$$\mathbf{s} \cdot \mathbf{s} = \rho^2 \left(1 + r^2 \cos^2(\theta - \phi) \right) \,. \tag{B14}$$

Consider first the integral of $\mathbf{s} \cdot \mathbf{s}$:

$$\int_{-\pi}^{\pi} \int_{0}^{r_{v}} (\mathbf{s} \cdot \mathbf{s}) \, r \, dr \, d\theta \,. \tag{B15}$$

To obtain the desired ensemble average we integrate over ρ and ϕ as follows:

$$\int_{-\pi}^{\pi} \int_{0}^{\infty} P(\rho) \left[\int_{-\pi}^{\pi} \int_{0}^{r_{\upsilon}} (\mathbf{s} \cdot \mathbf{s}) r \, dr \, d\theta \right] \rho \, d\rho \, d\phi \,. \tag{B16}$$

This integral can be split into two parts:

$$\int_0^\infty \rho^3 P(\rho) \, d\rho \int_{-\pi}^{\pi} d\phi \int_0^{r_v} r \, dr \int_{-\pi}^{\pi} d\theta = 2\pi^2 P_2 \, r_v^2 \,, \qquad (B17)$$

and

$$\int_{0}^{\infty} \rho^{3} P(\rho) \, d\rho \int_{-\pi}^{\pi} d\phi \int_{0}^{r_{v}} r^{3} \, dr \int_{-\pi}^{\pi} \cos^{2} \theta' \, d\theta' = \frac{\pi^{2}}{2} P_{2} \, r_{v}^{4} \,, \qquad (B18)$$

where $\theta' = (\theta - \phi)$ and

$$P_2 = \int_0^\infty \rho^3 P(\rho) \, d\rho \,. \tag{B19}$$

We note in passing that P_2 is a measure of the average squared magnitude of the brightness gradient. Combining the two parts above we find that

$$\iint \mathbf{s} \cdot \mathbf{s} \, dx \, dy = 2\pi^2 P_2 \, r_v^2 (1 + r_v^2/4) \,. \tag{B20}$$

Similarly we obtain

$$I_{xx} = \int_{-\pi}^{\pi} \int_{0}^{\infty} P(\rho) \left[\int_{-\pi}^{\pi} \int_{0}^{r_{v}} \rho^{2} \cos^{2} \phi r \, dr \, d\theta \right] \rho \, d\rho \, d\phi$$

=
$$\int_{0}^{\infty} \rho^{3} P(\rho) \, d\rho \int_{-\pi}^{\pi} \cos^{2} \phi \, d\phi \int_{0}^{r_{v}} r \, dr \int_{-\pi}^{\pi} d\theta = \pi^{2} P_{2} \, r_{v}^{2} \,, \qquad (B21)$$

and

$$I_{zz} = \int_{-\pi}^{\pi} \int_{0}^{\infty} P(\rho) \left[\int_{-\pi}^{\pi} \int_{0}^{r_{v}} \rho^{2} r^{2} \cos^{2}(\theta - \phi) r \, dr \, d\theta \right] \rho \, d\rho \, d\phi$$

= $\int_{0}^{\infty} \rho^{3} P(\rho) \, d\rho \int_{-\pi}^{\pi} d\phi \int_{0}^{r_{v}} r^{3} \, dr \int_{-\pi}^{\pi} \cos^{2} \theta' \, d\theta' = \frac{\pi^{2}}{4} P_{2} r_{v}^{4} ,$ (22)
ile $I_{w} = I_{rre}$

while I_{yy}

The moment matrix is diagonal, so I_{xx} , I_{yy} and I_{zz} are the eigenvalues. The condition number is the ratio of the largest to the smallest or

$$\frac{2}{r_v^2} \quad \text{for} \quad r_v^2 \le 2 \qquad \text{and} \qquad \frac{r_v^2}{2} \quad \text{for} \quad r_v^2 \ge 2 \,. \tag{B23}$$

This result does not depend on P_2 , as stated earlier.

Ensemble average of the integral of vv^T

Here we proceed much as in the previous section with

$$\mathbf{v} = \begin{pmatrix} +E_y + y(xE_x + yE_y) \\ -E_x - x(xE_x + yE_y) \\ yE_x - xE_y \end{pmatrix} = \rho \begin{pmatrix} +\sin\phi + r^2\sin\theta\cos(\theta - \phi) \\ -\cos\phi - r^2\cos\theta\cos(\theta - \phi) \\ r\sin(\theta - \phi) \end{pmatrix},$$
(B24)

and

$$\mathbf{v} \cdot \mathbf{v} = \rho^2 (1+r^2) \left(1 + r^2 \cos^2(\theta - \phi) \right) ,$$
 (B25)

which follows from $\mathbf{v} = \mathbf{s} \times \mathbf{r}$, $\mathbf{s} \cdot \mathbf{r} = 0$, $\mathbf{r} \cdot \mathbf{r} = 1 + r^2$ and

$$\mathbf{s} \cdot \mathbf{s} = \rho^2 \left(1 + r^2 \cos^2(\theta - \phi) \right) \,. \tag{B26}$$

After some tedious manipulations, similar to those in the last section, we find

$$\overline{\iint \mathbf{v} \cdot \mathbf{v} \, dx \, dy} = \frac{\pi^2}{2} P_2 \, r_v^2 \left(1 + 3r_v^2 / 4 + r_v^4 / 6 \right) \,. \tag{B27}$$

 and

$$I_{xx} = \pi^2 P_2 r_v^2 \left(1 + r_v^2 / 2 + r_v^4 / 6 \right) = I_{yy} , \qquad (B28)$$

and

$$I_{zz} = \pi^2 P_2 \left(r_v^4 / 2 \right). \tag{B29}$$

Again, the matrix is diagonal and so I_{xx} , I_{yy} and I_{zz} are the eigenvalues. The condition number is just

$$\frac{2}{r_v^2} + 1 + \frac{r_v^2}{3}, \qquad (B30)$$

which is independent of P_2 once more.

40

Imagine that you are a terrestrial, aquatic, or flying robot and that you want to keep track of where you are, relative to some coordinate system, given only that you have an eye or a camera to look through. You have no gyroscope, no compass, no navigation radio, no form of radar or sonar, and no map.

You might think that the only hope would be to understand what you see at least to the point of identifying some features or objects to keep track of as you move. Perhaps surprisingly, it seems better, in most cases, to work directly with large numbers of image points, noting only their first derivatives of brightness with respect to time and space. Although a general solution is not yet known, Horn and Weldon show how to deal with special cases in which the motion is pure rotation, pure translation, or arbitrary motion when the rotation is known. Their analysis demonstrates the critical importance of sufficient field of view and of adequate, aliasing-avoiding spatial and temporal filtering.

ARTIFICIAL INTELLIGENCE AT MIT EXPANDING FRONTIERS

edited by Patrick Henry Winston with Sarah Alexandra Shellard

Volume 2

