12

Motion Field & Optical Flow

A great deal of useful information can be extracted from time-varying images. At first, it might seem foolhardy to consider processing sequences of images, given the difficulty of interpreting even a single image. Curiously, though, some information is easier to obtain from a time sequence. The apparent motion of brightness patterns observed when a camera is moving relative to the objects being imaged is called the optical flow. In this chapter we discover how, given a sequence of images, we can calculate the optical flow. We start by contrasting the optical flow with the motion field, a purely geometric concept. A constraint equation is then uncovered that relates the gradient of brightness to the local flow velocity. We find that it is not possible to recover the optical flow locally. Additional information is required. One way to provide such information is to make assumptions about the shapes of the surfaces being imaged. We derive an iterative scheme for estimating the optical flow under the assumption that it varies smoothly.

The optical flow is a useful concept even when the surfaces being imaged deform, but in the special case of rigid body motion the optical flow is highly constrained. In chapter 17, where we explore passive navigation, we discuss how to recover both the motion of the camera relative to a fixed environment and the shapes of the surfaces being imaged. Some of methods we discuss there use the optical flow as an intermediate result.

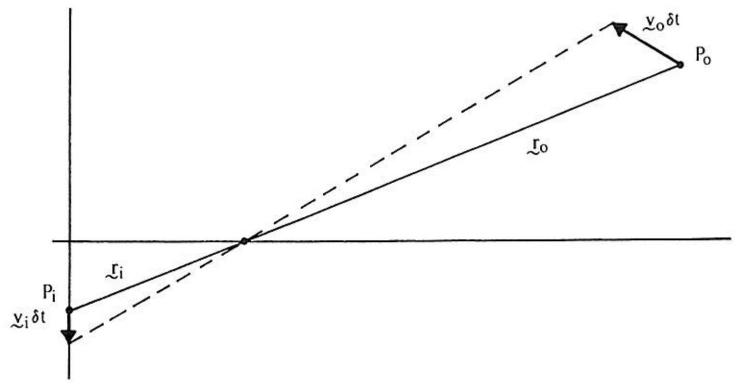

Figure 12-1. Displacement of a point in the environment causes a displacement of the corresponding image point. The relationship between the velocities can be found by differentiating the perspective projection equation.

12.1 Motion Field

When objects move in front of a camera, or when a camera moves through a fixed environment, there are corresponding changes in the image. These changes can be used to recover the relative motions as well as the shapes of the objects.

We first define the *motion field,* which assigns a velocity vector to each point in the image. At a particular instant in time, a point P_i in the image corresponds to some point P_o on the surface of an object. The two are connected by the projection equation. In the case of perspective projection, a ray from the image point through the center of the lens can be extended until it strikes an opaque surface (figure 12-1).

Let the object point P_o have velocity $\mathbf{v}_o$ relative to the camera. This induces a motion $\mathbf{v}_i$ in the corresponding image point P_i. The point P_o moves $\mathbf{v}_o\,\delta t$ in a time interval δt, and its image P_i moves $\mathbf{v}_i\,\delta t$. The velocities are given by

$$\mathbf{v}_o = \frac{d\mathbf{r}_o}{dt} \qquad \text{and} \qquad \mathbf{v}_i = \frac{d\mathbf{r}_i}{dt},$$

where r_o and r_i are related by

$$\frac{1}{f'}\mathbf{r}_i = \frac{1}{\mathbf{r}_o \cdot \hat{\mathbf{z}}}\mathbf{r}_o.$$

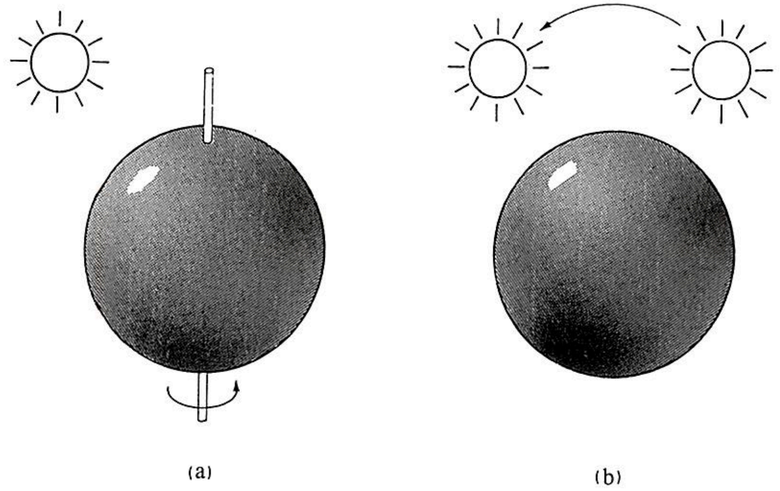

Figure 12-2. The optical flow is not always equal to the motion field. In (a) a smooth sphere is rotating under constant illumination—the image does not change, yet the motion field is nonzero. In (b) a fixed sphere is illuminated by a moving source—the shading in the image changes, yet the motion field is zero.

Differentiation of this perspective projection equation yields

$$\frac{1}{f'}\mathbf{v}_i = \frac{(\mathbf{r}_o \cdot \hat{\mathbf{z}})\mathbf{v}_o - (\mathbf{v}_o \cdot \hat{\mathbf{z}})\mathbf{r}_o}{(\mathbf{r}_o \cdot \hat{\mathbf{z}})^2} = \frac{(\mathbf{r}_o \times \mathbf{v}_o) \times \hat{\mathbf{z}}}{(\mathbf{r}_o \cdot \hat{\mathbf{z}})^2}.$$

We shall not pursue this any further in this chapter. What is important here is that a vector can be assigned in this way to every image point. These vectors constitute the motion field.

Neighboring points on an object have similar velocities. We expect, then, that the induced motion field in the image is also continuous in most places. Exceptions will occur on the silhouettes of the images of the objects, where discontinuities in the motion field can be expected.

12.2 Optical Flow

Brightness patterns in the image move as the objects that give rise to them move. *Optical flow* is the apparent motion of the brightness pattern. Ideally the optical flow will correspond to the motion field, but we show next that this need not always be so.

Consider first a perfectly uniform sphere rotating in front of an imaging system (figure 12-2a). There will be spatial variation of brightness, or

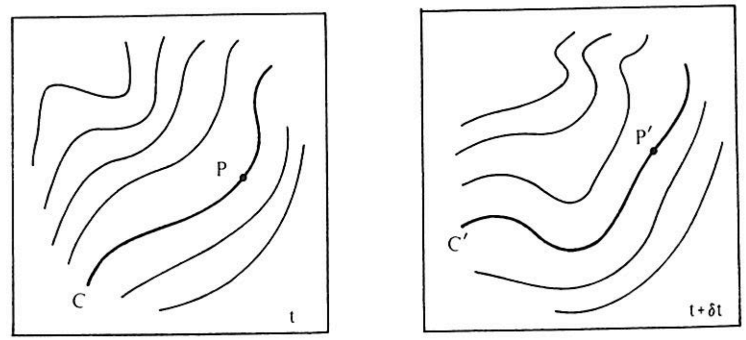

Figure 12-3. The apparent motion of brightness patterns is an awkward concept. It is not easy to decide which point P' on a contour C' of constant brightness in the second image corresponds to a particular point P on the corresponding contour C in the first image.

shading, in the image of the sphere, since the surface is curved. This shading, however, does not move with the surface, and so the image does not change with time. In this case the optical flow is zero everywhere, despite a nonzero motion field. Next, consider a fixed sphere illuminated by a moving light source (figure 12-2b). The shading in the image will change as the source moves. In this case the optical flow is clearly nonzero, while the motion field is zero everywhere. Virtual images and shadows lead us to other cases in which the optical flow is not equal to the motion field.

What is accessible to us is the optical flow, and we shall have to depend on the fact that, except for special situations such as the ones discussed above, the optical flow is not too different from the motion field. This will allow us to estimate relative motion by means of the changing image.

What do we mean by apparent motion of the brightness pattern? Consider a point P in the image with brightness E at time t (figure 12-3). To which point P' in the image at time $t+\delta t$ does it correspond—that is, how did the brightness pattern move in the intervening time? Typically there will be many points near P with the same brightness E. If brightness varies continuously in the part of the image of interest, then the point P will lie on an isobrightness contour C. At time $t + \delta t$ there will be some nearby isobrightness contour C' with the same brightness level. But what is the correspondence between points on C and points on C'? This question is hard to answer since the two contours will typically not even have exactly the same shape.

We thus note that the optical flow is not uniquely determined by local information in the changing image. Another example will make this clear.

Consider a patch of uniform brightness in the image that does not change with time. Perhaps the "most likely" optical flow is one that is zero everywhere. But in fact within the uniform patch we can assign any pattern of vector displacements we like. Presumably, though, we would prefer the simplest explanation of the observed changing (or in this case unchanging) image.

Let $E(x, y, t)$ be the irradiance at time t at the image point (x, y). Then, if $u(x, y)$ and $v(x, y)$ are the x and y components of the optical flow vector at that point, we expect that the irradiance will be the same at time $t + \delta t$ at the point $(x + \delta x, y + \delta y)$, where $\delta x = u\delta t$ and $\delta y = v\delta t$. That is,

$$E(x + u\,\delta t, y + v\,\delta t, t + \delta t) = E(x, y, t)$$

for a small time interval δt. This single constraint is not sufficient to determine both u and v uniquely. It is also clear that we can take advantage of the fact that the motion field is continuous almost everywhere.

If brightness varies smoothly with x, y, and t, we can expand the left-hand side of the equation above in a Taylor series and so obtain

$$E(x, y, t) + \delta x \frac{\partial E}{\partial x} + \delta y \frac{\partial E}{\partial y} + \delta t \frac{\partial E}{\partial t} + e = E(x, y, t),$$

where e contains second- and higher-order terms in δx, δy, and δt. Canceling $E(x, y, t)$, dividing through by δt, and taking the limit as $\delta t \to 0$, we obtain

$$\frac{\partial E}{\partial x}\frac{dx}{dt} + \frac{\partial E}{\partial y}\frac{dy}{dt} + \frac{\partial E}{\partial t} = 0,$$

which is actually just the expansion of the equation

$$\frac{dE}{dt} = 0$$

in the total derivative of E with respect to time. Using the abbreviations

$$u = \frac{dx}{dt}, \quad v = \frac{dy}{dt},$$

$$E_x = \frac{\partial E}{\partial x}, \quad E_y = \frac{\partial E}{\partial y}, \quad E_t = \frac{\partial E}{\partial t},$$

we obtain

$$E_x u + E_y v + E_t = 0.$$

The derivatives E_x, E_y, and E_t are estimated from the image. The above equation is called the *optical flow constraint equation*, since it expresses a constraint on the components u and v of the optical flow.

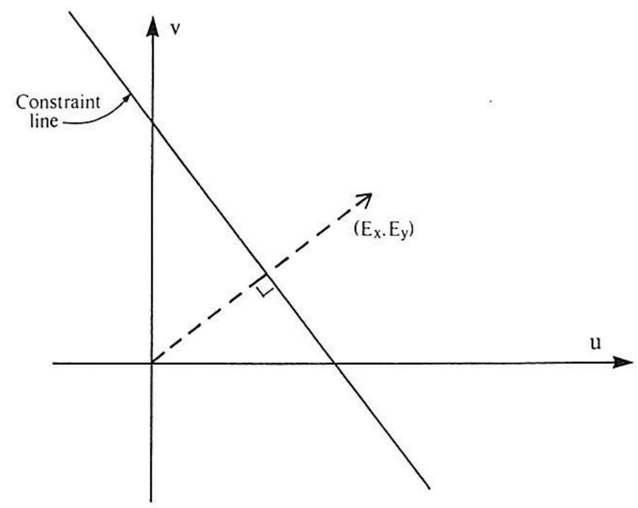

Figure 12-4. Local information on the brightness gradient and the rate of change of brightness with time provides only one constraint on the components of the optical flow vector. The flow velocity has to lie along a straight line perpendicular to the direction of the brightness gradient. We can only determine the component in the direction of the brightness gradient. Nothing is known about the flow component in the direction at right angles.

Consider a two-dimensional space with axes u and v, which we shall call *velocity space* (figure 12-4). Values of (u, v) satisfying the constraint equation lie on a straight line in velocity space. All that a local measurement can do is to identify this constraint line. We can rewrite the constraint equation in the form

$$(E_x, E_y) \cdot (u, v) = -E_t.$$

The component of optical flow in the direction of the brightness gradient $(E_x, E_y)^T$ is thus

$$\frac{E_t}{\sqrt{E_x^2 + E_y^2}}.$$

We cannot, however, determine the component of the optical flow at right angles to this direction, that is, along the isobrightness contour. This ambiguity is also known as the *aperture problem*.

12.3 Smoothness of the Optical Flow

Now it is time to introduce additional constraint. In chapter 17, where we explore passive navigation, we shall make the assumption that we are dealing with rigid bodies. In that case, translation and rotation of the camera relative to the object are the key parameters to recover. The assumption of rigid body motion is very restrictive and provides a powerful constraint on the solution. Here, however, we want to explore less limiting assumptions so that we can deal with more general situations, such as deforming elastic bodies.

Usually the motion field varies smoothly in most parts of the image. We shall try to minimize a measure of departure from smoothness,

$$e_s = \iint \left((u_x^2 + u_y^2) + (v_x^2 + v_y^2)\right) dx\, dy,$$

the integral of the square of the magnitude of the gradient of the optical flow. The error in the optical flow constraint equation,

$$e_c = \iint \left(E_x u + E_y v + E_t\right)^2 dx\, dy,$$

should also be small. Overall, then, we want to minimize $e_s + \lambda e_c$, where λ is a parameter that weights the error in the image motion equation relative to the departure from smoothness. This parameter will be large if brightness measurements are accurate and small if they are noisy. Minimizing an integral of the form

$$\iint F(u, v, u_x, u_y, v_x, v_y)\, dx\, dy$$

is a problem in the calculus of variations (see the appendix). The corresponding Euler equations are

$$F_u - \frac{\partial}{\partial x} F_{u_x} - \frac{\partial}{\partial y} F_{u_y} = 0,$$
$$F_v - \frac{\partial}{\partial x} F_{v_x} - \frac{\partial}{\partial y} F_{v_y} = 0.$$

In our case,

$$F = (u_x^2 + u_y^2) + (v_x^2 + v_y^2) + \lambda \left(E_x u + E_y v + E_t\right)^2,$$

so the Euler equations yield

$$\nabla^2 u = \lambda \left(E_x u + E_y v + E_t\right) E_x,$$
$$\nabla^2 v = \lambda \left(E_x u + E_y v + E_t\right) E_y,$$

where

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$$

is the Laplacian operator. This coupled pair of elliptic second-order partial differential equations can be solved using iterative methods.

12.4 Filling in Optical Flow Information

We can derive some information about the solution method by inspecting these equations. For example, where the brightness gradient is zero, the right-hand sides of the equations are also zero. In these regions, u and v each satisfy Laplace's equation. Thus, in regions of uniform brightness, where the optical flow velocity cannot be found locally, it is interpolated from the optical flow velocities in surrounding areas. Similarly, along an edge, the optical flow component in the direction of the edge itself cannot be found but must be interpolated from values further along the edge. In this case, however, the motion of the edge itself at least provides constraint in the direction of the brightness gradient.

At a corner, the direction of the brightness gradient changes rapidly as we go from one picture cell to another, so that constraint is available in a small neighborhood to determine the optical flow fully. Reliable information is available at *brightness corners*, less so at places where brightness is constant along some direction in the image, and least where brightness does not vary spatially at all. In the case of a rectangular region of uniform brightness moving against a background, information is filled in from the corners and the edges.

12.5 Boundary Conditions

A problem is said to be *well-posed* if a solution exists and if the solution is unique. A partial differential equation typically has an infinite number of solutions unless it is constrained by suitable boundary conditions. Hence, a problem leading to a partial differential equation is well-posed when boundary conditions are given that will ensure a unique solution. In the case of an elliptic linear second-order partial differential equation, such as Poisson's equation, giving the value of the function on a simply closed curve enclosing the region of interest is one way to guarantee a unique solution. Alternatively, we can give the *normal derivative*, that is, the derivative of the unknown function in a direction perpendicular to the boundary curve. Higher-order differential equations typically require additional constraint, such as both the value of the function and the normal derivative.

When we use the calculus of variations, we may have information about the behavior of the solution on the boundary, or we may instead allow the boundary to be free. The latter case leads to *natural boundary conditions* (see the appendix for more details). In minimizing an integral of the form

$$\iint F(u, v, u_x, u_y, v_x, v_y)\, dx\, dy$$

the natural boundary conditions can be shown to be

$$F_{u_x}\frac{dy}{ds} = F_{u_y}\frac{dx}{ds} \quad \text{and} \quad F_{v_x}\frac{dy}{ds} = F_{v_y}\frac{dx}{ds},$$

where s denotes arclength along the boundary curve. Now

$$\hat{\mathbf{n}} = \left(\frac{dy}{ds}, -\frac{dx}{ds}\right)^T$$

is a unit vector perpendicular to the boundary. We can thus rewrite the above conditions in the form

$$\left(F_{u_x}, F_{u_y}\right)^T \cdot \hat{\mathbf{n}} = \mathbf{0} \quad \text{and} \quad \left(F_{v_x}, F_{v_y}\right)^T \cdot \hat{\mathbf{n}} = \mathbf{0}.$$

In our case,

$$(u_x, u_y)^T \cdot \hat{\mathbf{n}} = \mathbf{0} \quad \text{and} \quad (v_x, v_y)^T \cdot \hat{\mathbf{n}} = \mathbf{0},$$

that is, the normal derivatives of u and v must be zero.

12.6 The Discrete Case

We could now approximate the continuous solution using a finite-difference scheme, but readers not yet comfortable with the calculus of variations might prefer the following direct method. Using the values of the optical flow at the grid point (i, j) and its neighbors (figure 12-5) allows us to measure the departure from smoothness as

$$s_{i,j} = \frac{1}{4}\big((u_{i+1,j} - u_{i,j})^2 + (u_{i,j+1} - u_{i,j})^2 + (v_{i+1,j} - v_{i,j})^2 + (v_{i,j+1} - v_{i,j})^2\big),$$

while the error in the optical flow constraint equation is

$$c_{ij} = (E_x u_{ij} + E_y v_{ij} + E_t)^2,$$

where E_x, E_y, and E_t are estimates of the rates of change of brightness with respect to x, y, and t at the point (i, j). We seek a set of values $\{u_{ij}\}$ and $\{v_{ij}\}$ that minimize

$$e = \sum_i \sum_j (s_{ij} + \lambda\, c_{ij}).$$

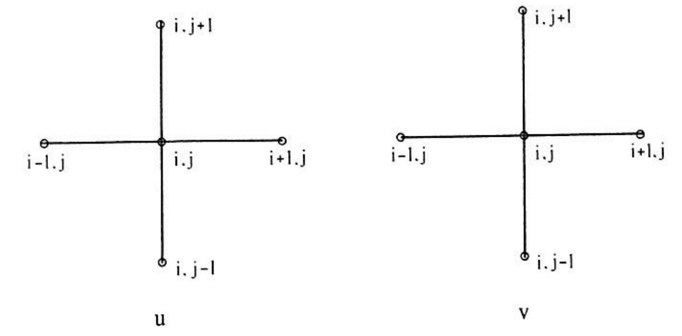

Figure 12-5. The sum of squares of the first partial derivatives of u and v can be estimated using the differences of the optical flow components at neighboring points.

Differentiating e with respect to u_{kl} and v_{kl} yields

$$\frac{\partial e}{\partial u_{kl}} = 2(u_{kl} - \overline{u}_{kl}) + 2\lambda\,(E_x u_{kl} + E_y v_{kl} + E_t)\,E_x,$$

$$\frac{\partial e}{\partial v_{kl}} = 2(v_{kl} - \overline{v}_{kl}) + 2\lambda\,(E_x u_{kl} + E_y v_{kl} + E_t)\,E_y,$$

where $\overline{u}$ and $\overline{v}$ are local averages of u and v. The extremum occurs where the above derivatives of e are zero. The resultant equations can be rewritten in the form

$$(1 + \lambda E_x^2)u_{kl} + \lambda E_x E_y v_{kl} = \overline{u}_{kl} - \lambda E_x E_t,$$

$$\lambda E_y E_x u_{kl} + (1 + \lambda E_y^2)v_{kl} = \overline{v}_{kl} - \lambda E_y E_t.$$

We can think of this as a pair of equations in u_{kl} and v_{kl}. The determinant of the 2×2 coefficient matrix is

$$1 + \lambda(E_x^2 + E_y^2),$$

so that

$$\left(1 + \lambda(E_x^2 + E_y^2)\right) u_{kl} = +(1 + \lambda E_y^2)\overline{u}_{kl} - \lambda E_x E_y \overline{v}_{kl} - \lambda E_x E_t,$$

$$\left(1 + \lambda(E_x^2 + E_y^2)\right) v_{kl} = -\lambda E_y E_x \overline{u}_{kl} + (1 + \lambda E_x^2)\overline{v}_{kl} - \lambda E_y E_t.$$

We can solve these equations now for u_{kl} and v_{kl}. The result immediately

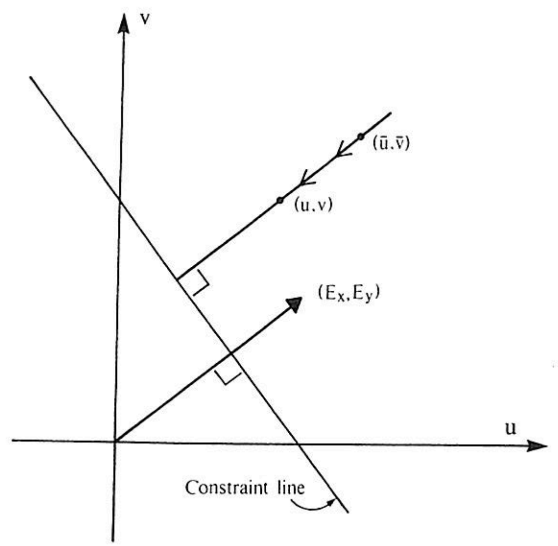

Figure 12-6. In the iterative scheme for estimating the optical flow, the new value (u, v) at a point is the average of the values of the neighbors, $(\overline{u}, \overline{v})$, minus an adjustment in the direction toward the constraint line.

suggests an iterative scheme such as

$$u_{kl}^{n+1} = \overline{u}_{kl}^{n} - \frac{E_x \overline{u}_{kl}^{n} + E_y \overline{v}_{kl}^{n} + E_t}{1 + \lambda(E_x^2 + E_y^2)} E_x,$$

$$v_{kl}^{n+1} = \overline{v}_{kl}^{n} - \frac{E_x \overline{u}_{kl}^{n} + E_y \overline{v}_{kl}^{n} + E_t}{1 + \lambda(E_x^2 + E_y^2)} E_y.$$

There is an interesting geometric interpretation of these equations (figure 12-6). The new value of (u, v) at a point is set equal to the average of the surrounding values, minus an adjustment, which in velocity space is in the direction of the brightness gradient.

To implement this scheme we need to estimate the spatial and time derivatives of brightness. This can be done easily using first differences of values on a grid (figure 12-7). If the indices i, j, and k correspond to x, y, and t, respectively, then consistent estimates of the three first partial derivatives can be obtained as follows:

$$E_x \approx \frac{1}{4\,\delta x}\left(E_{i+1,j,k} + E_{i+1,j,k+1} + E_{i+1,j+1,k} + E_{i+1,j+1,k+1}\right) - \frac{1}{4\delta x}\left(E_{i,j,k} + E_{i,j,k+1} + E_{i,j+1,k} + E_{i,j+1,k+1}\right),$$

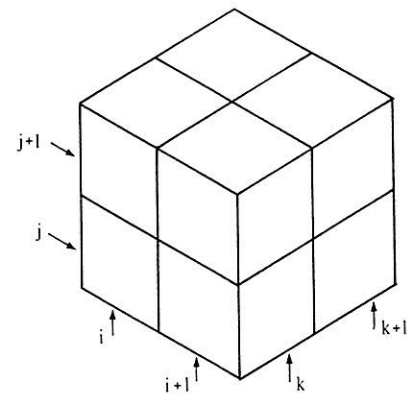

Figure 12-7. The first derivatives required in the iterative scheme can be estimated using first differences in a $2 \times 2 \times 2$ cube of brightness values. The estimates apply to the point where four picture cells meet, at a time halfway between two successive images.

$$E_y \approx \frac{1}{4\,\delta y}\left(E_{i,j+1,k} + E_{i,j+1,k+1} + E_{i+1,j+1,k} + E_{i+1,j+1,k+1}\right) - \frac{1}{4\,\delta y}\left(E_{i,j,k} + E_{i,j,k+1} + E_{i+1,j,k} + E_{i+1,j,k+1}\right),$$

$$E_t \approx \frac{1}{4\,\delta t}\left(E_{i,j,k+1} + E_{i,j+1,k+1} + E_{i+1,j,k+1} + E_{i+1,j+1,k+1}\right) - \frac{1}{4\,\delta t}\left(E_{i,j,k} + E_{i,j+1,k} + E_{i+1,j,k} + E_{i+1,j+1,k}\right).$$

This means that the optical flow velocities will be estimated at points lying between picture cells and between successive frames.

Figure 12-8 shows four successive synthetic images of a rotating sphere covered with a smoothly varying brightness pattern. Spatial and time derivatives of brightness estimated from these images provide the input to the iterative optical flow algorithm described above.

The estimated optical flow after 1, 4, 16, and 64 iterations is shown in figure 12-9. The first guess is influenced strongly by the brightness pattern on the sphere. After a few iterations, the estimated flow vectors converge to the correct solution, except on the silhouette.

The computed motion field is shown for comparison in figure 12-10b. The estimated optical flow shown in figure 12-10a differs only a little from the motion field, except on the silhouette.

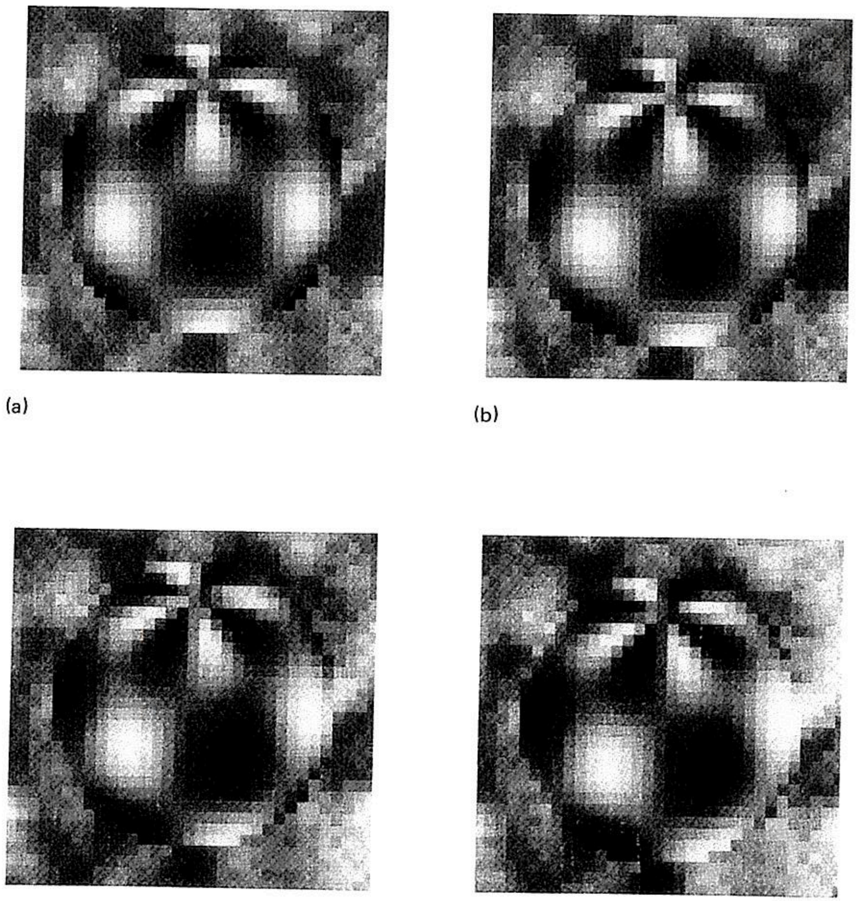

Figure 12-8. Four frames of a synthetic image sequence showing a sphere slowly rotating in front of a randomly patterned background.

12.7 Discontinuities in Optical Flow

There will be discontinuities in the optical flow on the silhouettes, where one object occludes another. We must detect these places if we are to prevent the method presented above from trying to continue the solution smoothly from one region to the other. This seems like a chicken-and-egg problem: If we have a good estimate of the optical flow, we can look for places where it changes very rapidly in order to segment the picture. On

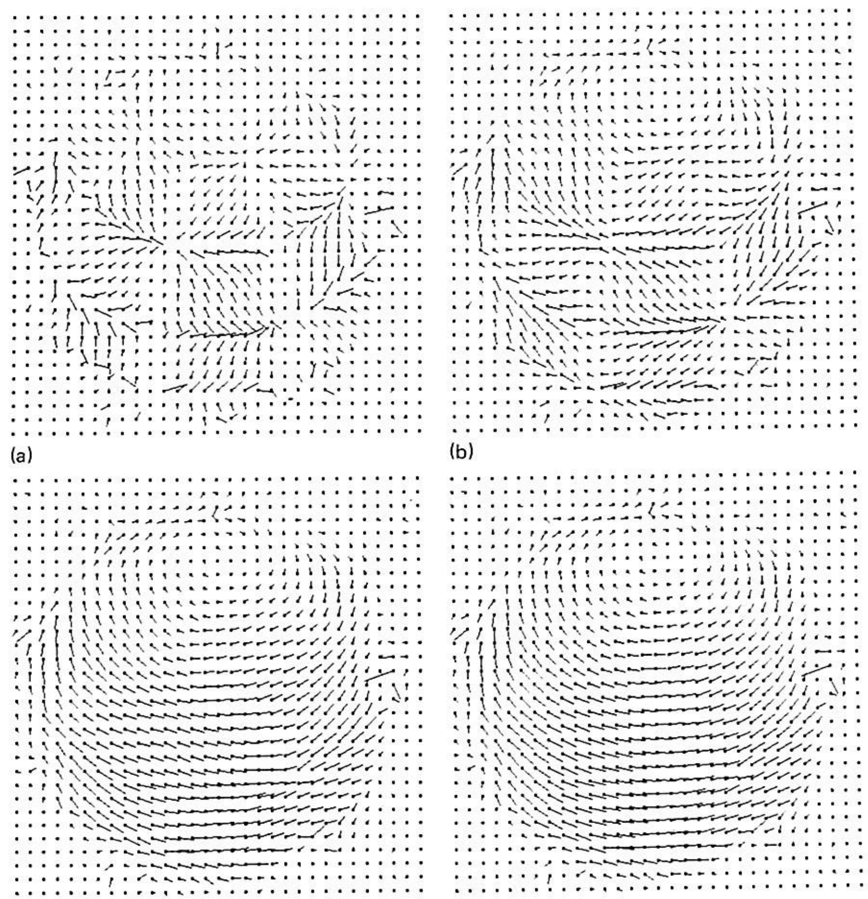

Figure 12-9. Estimates of the optical flow shown in the form of needle diagrams after 1, 4, 16, and 64 iterations of the algorithm.

the other hand, if we could segment the picture well, we would produce a better estimate of the optical flow. The solution to this dilemma is to incorporate the segmentation into the iterative solution for the optical flow. That is, after each iteration we look for places where the flow changes rapidly. At these places we set down marks that inhibit the next iteration from smoothly connecting the solution across the discontinuities. We first set the threshold for this decision very high in order to prevent premature carving up of the image. We reduce the threshold as better and better estimates of the optical flow become available.

We can get some feel for how well our assumption of smoothness holds

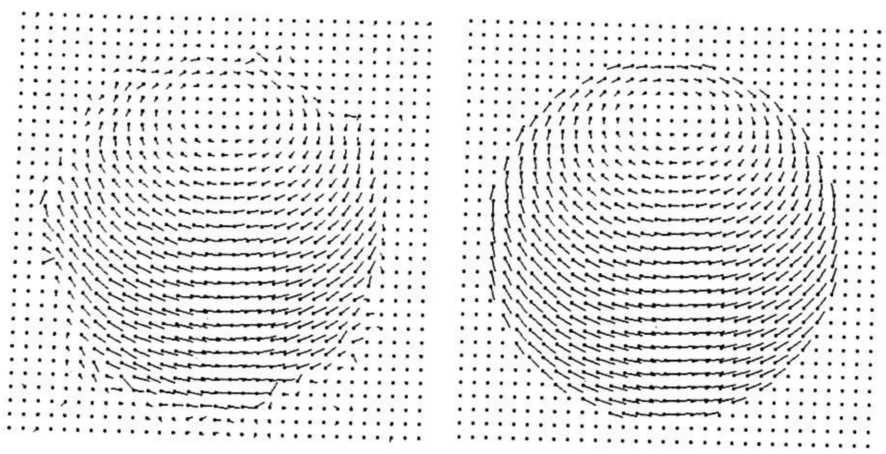

Figure 12-10. (a) The estimated optical flow after several more iterations. (b) The computed motion field.

up in situations of practical interest. A common case is that of an object performing a rigid body motion, a translation, and a rotation. For simplicity, we shall work here with orthographic projection. Translation leads to uniform optical flow in the image region corresponding to the moving object. The flow within that region is (exceptionally) smooth. As expected, there will be discontinuities at the boundaries of the region, where the object occludes the background.

Rotation is a little harder to understand. First of all, we are concerned here with instantaneous quantities only. That is, we are concerned only with derivatives of coordinates at a particular instant in time; we do not follow the motion over time. We show in exercise 12-6 that rotation about an arbitrary axis is equivalent to rotation about a parallel axis through the origin, combined with a translation. This compensating translation is equal to the cross-product of a vector from the origin to the axis and the rotation vector. In any case, we can, without loss of generality, restrict our attention to rotation about an axis through the origin.

Suppose that the rotation of the object is specified by the vector $\omega = (\alpha, \beta, \gamma)^T$, where the length of the vector gives the rate of rotation, while the axis of rotation lies along the vector. The velocity of a point on the object is equal to the cross-product of a vector $\mathbf{r} = (x, y, z)^T$, from the axis of rotation to the point, and the given rotation vector ω. Since we are dealing with orthographic projection, the image coordinates x' and y' are the same as the object coordinates x and y, respectively. The components

u and v of the optical flow are just the x and y components of the velocity of the point on the object:

$$u = \beta z - \gamma y \quad \text{and} \quad v = \gamma x - \alpha z.$$

It is now trivial to determine the departure from smoothness of the optical flow. We have

$$(u_x^2 + u_y^2) + (v_x^2 + v_y^2) = (\alpha^2 + \beta^2)(z_x^2 + z_y^2) - 2\gamma(\alpha z_x + \beta z_y) + 2\gamma^2.$$

The optical flow will not be smooth near silhouettes of smoothly curved objects, since the slope of the surface approaches infinity there. Moreover, rotation about the optical axis increases the measure that we used for the departure from smoothness.

An important application of the method for computing the optical flow is in passive navigation. Here we must determine the path and instantaneous attitude of a vehicle from information gleaned about the environment, but without emitting sampling radiation from the vehicle. This is a subject to which we shall return in chapter 17.

12.8 References

The books by Hildreth, *The Measurement of Visual Motion* [1983], and Ullman, *The Interpretation of Visual Motion* [1979], discuss the determination and interpretation of visual motion. Huang edited a collection of work presented at a NATO conference in *Image Sequence Processing and Dynamic Scene Analysis* [1983].

The calculus of variations is treated in many books, including volume I of *Methods of Mathematical Physics* by Courant & Hilbert [1953] and *Calculus of Variations: With Applications to Physics & Engineering* by Weinstock [1974].

There are numerous books discussing partial differential equations, among them *Partial Differential Equations: Theory and Technique* by Carrier & Pearson [1976], volume II of *Methods of Mathematical Physics* by Courant & Hilbert [1962], *Partial Differential Equations* by John [1971], and a book of the same title by Moon & Spencer [1969]. Regularization is discussed by Tikhonov & Arsenin in *Solutions of Ill-Posed Problems*, [1977].

Some of the earliest work on time-varying imagery had as its motive compression of the highly redundant video signal of a typical slow-changing scene. See, for example, Limb & Murphy [1975], Netravali & Robbins [1979], and Stuller, Netravali, & Robbins [1980]. Other early work was done by Fennema & Thompson [1979].

Gibson coined the term "optical flow." For some of his ideas, see Gibson et al. [1959]. This area has seen tremendous activity in the last five years. Horn & Schunck [1981] developed a simple iterative algorithm for estimating the optical flow. Schunck & Horn [1981] expanded on this and explored a slightly different approach. Edges at various scales contain most of the information in an image, and Hildreth [1983, 1984] has therefore approached the problem from the point of view of determining the motion of curves, such as zero-crossing contours, rather than the brightness pattern directly. Nagel [1982, 1983a] has developed methods for tracking brightness corners. One of his schemes [1983b] is a complex modification of the method developed by Horn & Schunck that allows smoothing effects to propagate only in certain directions.

Another approach to the problem involves tracking distinctive brightness patterns in the image. Edges in particular seem to be suitable targets for such an effort, as shown by Haynes & Jain [1982] and Hildreth [1983, 1984]. Other relevant papers include Thompson [1981], Marr & Ullman [1981], Ullman [1981], Schalkoff & McVey [1982], Paquin & Dubois [1983], Wohn, Davis, & Thrift [1983] and Longuet-Higgins & Prazdny [1980].

12.9 Exercises

12-1 Optical flow cannot be determined locally in image regions where brightness is constant and thus the brightness gradient is zero. The method presented in this chapter fills the flow in from the surround using Laplace's equation. Here we investigate circumstances in which the filling-in operation happens to produce exactly the desired result.

(a) Consider a brightness pattern translating with velocity (u_0, v_0) in the image. Suppose that the optical flow velocity has been determined accurately on the border ∂R of some region R where brightness is constant. Show that the solutions of the Laplace equations $\nabla^2 u = 0$ and $\nabla^2 v = 0$, given the boundary conditions, produce the correct optical flow in the region R. Hint: The expected flow will be the unique solution of the problem if and only if the Laplacian within the region is zero and the flow satisfies the boundary conditions.

(b) Does the filling-in method always produce the correct optical flow in the case of translation of the camera parallel to a planar surface that is perpendicular to the optical axis?

(c) Consider a brightness pattern rotating about the image point (x_0, y_0) with angular velocity ω. What is the optical flow? Suppose that the optical flow velocity has been determined accurately on the border ∂R of some region R in which brightness is constant. Do the solutions of the Laplace equations $\nabla^2 u = 0$ and $\nabla^2 v = 0$, given the boundary conditions, produce the correct optical flow in the region R?

(d) Does the filling-in method always produce the correct optical flow in the case of rotation of the camera about the optical axis?

12-2 Consider a camera moving along its optical axis toward a planar surface at right angles to the optical axis.

(a) Show that the optical flow is given by

$$u = \frac{W}{Z} x \quad \text{and} \quad v = \frac{W}{Z} y,$$

where W is the velocity and Z the distance to the plane. (Note the lack of dependence on the focal length of the lens.)

(b) Is the optical flow stationary (that is, independent of time)?

(c) Is the Laplacian of the optical flow zero?

(d) How could you predict the time to impact? (Note that this can be done despite the fact that we cannot recover the absolute value of either height or velocity.)

12-3 Here we consider some optical flow patterns that do not correspond to rigid body motions.

(a) Suppose you are looking down on the surface of a liquid in a container with a flat bottom. The liquid is running out of a hole in the center of the bottom of the container. If you assume that the depth of the liquid is constant (despite the flow), what is the optical flow? Assume that the optical axis is perpendicular to the surface of the liquid and that it passes through the hole. Hint: The flow of liquid through any one cylindrical surface with axis passing through the hole must equal the flow through any other cylindrical surface with axis passing through the hole.

(b) Find the Laplacian of the flow in the above case. Is there a singularity (that is, a place where the flow becomes infinite)?

(c) A cross section through a flow vortex, around a line in three dimensions, shows circular symmetry with no material moving inward or outward. Furthermore, the angular momentum of a layer of given thickness at one radial distance from the line is the same as that of another layer at a different radial distance. What is the optical flow?

(d) Find the Laplacian of the flow in the above case. Is there a singularity?

12-4 Show that the Laplacian of a second-order polynomial in x and y is zero if the coefficients of x^2 and y^2 are equal in magnitude and opposite in sign. Apply this result to the previous problems.

12-5 Suppose we use as the measure of departure from smoothness the integral

$$e_s = \iint \left(\nabla^2 u\right)^2 + \left(\nabla^2 v\right)^2 \, dx\, dy,$$

instead of

$$e_s = \iint \left((u_x^2 + u_y^2) + (v_x^2 + v_y^2)\right) \, dx\, dy.$$

What would be the appropriate Euler equations?