

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 734

September, 1983

THE COMPUTATION OF THE VELOCITY FIELD

Ellen C. Hildreth

ABSTRACT: The organization of movement in the changing retinal image provides a valuable source of information for analyzing the environment in terms of objects, their motion in space, and their three-dimensional structure. A description of this movement is not provided to our visual system directly, however; it must be inferred from the pattern of changing intensity that reaches the eye. This paper examines the problem of motion measurement, which we formulate as the computation of an instantaneous two-dimensional velocity field from the changing image. Initial measurements of motion take place at the location of significant intensity changes, as suggested by Marr and Ullman (1981). These measurements provide only one component of local velocity, and must be integrated to compute the two-dimensional velocity field. A fundamental problem for this integration stage is that the velocity field is not determined uniquely from information available in the changing image. We formulate an additional constraint of smoothness of the velocity field, based on the physical assumption that surfaces are generally smooth, which allows the computation of a unique velocity field. A theoretical analysis of the conditions under which this computation yields the correct velocity field suggests that the solution is physically plausible. Empirical studies show the predictions of this computation to be consistent with human motion perception.

© Massachusetts Institute of Technology, 1983

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for this research was provided by the Air Force Office of Scientific Research under contract F49620-83-C-0135.


1. INTRODUCTION

The organization of movement in the changing retinal image provides a valuable source of information for analyzing the environment in terms of objects, their motion in space, and their three-dimensional structure. A description of this movement is not provided to our visual system directly, however; it must be inferred from the pattern of changing intensity that reaches the eye. This paper presents a computational study of early motion analysis, whose focus is the measurement of motion in the changing image. The motivation of this work is to understand the nature of the computations that underly motion measurement in the human visual system.


The changing retinal image may be represented by a two-dimensional array of time-dependent light intensities, $I(x, y, t)$. Motion in the image may be described by a two-dimensional vector field $V(x, y, t)$ that specifies the direction and magnitude of velocity at points with coordinates (x, y) at time t . The measurement of visual motion may then be formulated as the computation of $V(x, y, t)$ from $I(x, y, t)$. For some visual tasks, it may be sufficient to compute only certain properties of the velocity field; for example, to respond quickly to a moving object, motion must be detected, but not necessarily measured. Other tasks, such as the recovery of three-dimensional structure, require a more complete and accurate computation of the velocity field (Ullman 1980, 1983a,b; Prazdny 1980; Clocksin 1980; Longuet-Higgins and Prazdny 1981; Longuet-Higgins 1981).

The measurement of motion poses significant theoretical problems for a computational study. First, local motion measurements, obtained directly from the changing image, in general only provide one component of local velocity. This is a consequence of the *aperture problem*, illustrated in Figure 1a (Wallach 1976; Fennema and Thompson 1979; Marr and Ullman 1981; Horn and Schunck 1981, Adelson and Movshon 1982). If the motion of the edge E is to be measured by a local motion detector that examines only an area A that is small compared to the overall extent of the edge, the only motion that can be extracted is the component c in the direction perpendicular to the local orientation of the edge. A local detector cannot distinguish between motions in the directions b, c and d in Figure 1a. In Figures 1b and 1c, a circle and square undergo pure translation in the directions given by the vectors at the center of the figures. The vectors along the contours represent the local perpendicular components of velocity that can be obtained directly from the changing image. To compute the true motion of the figure, a second stage of analysis is required, that combines these local measurements.


This combination stage faces a deeper theoretical problem, however; the movement of elements in the image is not determined uniquely by the pattern of changing intensities. Thus, the true velocity field is not determined uniquely from the initial local motion measurements.



a.



b.



c.

Figure 1. The Aperture Problem. (a) An operation that views the moving edge E through the local aperture A can compute only the component of motion c in the direction perpendicular to the edge. The true motion of the edge is ambiguous. (b) and (c) The perpendicular components of velocity for a translating circle and square.

Two factors contribute to this ambiguity of motion. The first is the loss of information due to the projection of the three-dimensional world onto a two-dimensional image; multiple surfaces, undergoing different motions in space, may project to the same two-dimensional image. The second factor is the loss of information due to the projection into a pattern of changing intensity. The image that a surface projects onto the eye may not be sufficient to determine its movement in space. As an extreme example, a matte white sphere, rotating about a central axis, cannot be determined as such, on the basis of its projected image.

Figure 2 presents two simple examples that illustrate the ambiguity of the velocity field. In Figure 2a, the solid and dotted lines represent the image of a moving circle, at different instants of time. In the first frame (solid line), the circle lies parallel to the image plane, while in the




Figure 2. Ambiguity of the Velocity Field. (a) The arrows represent two possible velocity fields that are consistent with the changing image. (b) The curve C_1 rotates, translates and deforms over time to yield the curve C_2 . The velocity of the point p is ambiguous.

second frame (dotted line), the circle is slanted in depth. One velocity field that is consistent with the two frames is derived from pure rotation of the circle about the central vertical axis, as shown to the left in Figure 2a. (The arrows represent a sample of the local velocity vectors along the circle.) There could also be a component of rotation in the plane of the circle, about its center, as shown to the right in Figure 2a. In addition, this changing image might represent the projection of a different three-dimensional curve that is deforming over time, giving rise to yet another projected velocity field. This ambiguity is not peculiar to symmetric figures such as circles; it is a fundamental problem that is always present. In Figure 2b, the curve C_1 rotates, translates and deforms over time, to yield the curve C_2 . The motion of points from C_1 to C_2 is again ambiguous (consider, for example, different possible velocities for the point p). In general, there are infinitely many two-dimensional velocity fields that are consistent with the changing image.

To compute motion uniquely, additional constraint is therefore required, in the form of basic assumptions about the physical world that generally hold true. The main focus of this paper is the derivation of a particular constraint, the *smoothness constraint*, that allows the computation of a unique two-dimensional velocity field from the changing image of three-dimensional surfaces undergoing general motion in space. The constraint is developed formally in Section 3. In Section 4, we briefly present a velocity field algorithm that embodies this constraint, and show some results of the algorithm for a number of motion sequences. The predictions of the algorithm on known perceptual demonstrations suggest that a constraint such as smoothness may underly motion measurement in the human visual system.

2. BACKGROUND

In this section, we present some essential background material regarding three aspects of motion measurement that are relevant to the present study: (1) the possible existence of short range and long range motion processes in human vision, (2) the nature of the initial motion measurements from the changing image, and (3) additional assumptions for motion measurement that have been suggested previously.

2.1 Two Motion Processes in Human Vision

The computation of an instantaneous velocity field requires that movement in the image be roughly continuous. There are alternate representations of visual motion that are not so restricted. For example, motion can be described by an explicit *correspondence* over time, between elements in the image that represent the same physical feature under motion (Ullman 1979a). Motion measurement in this case requires locating identifiable elements in the changing image, and matching them over time. The input for a correspondence scheme may consist of a set of discrete frames, with large spatial separations between corresponding elements. The perception of motion by the human visual system also does not require that a stimulus move continuously across the visual field. With appropriate spatial and temporal presentation parameters, a stimulus presented sequentially can produce the impression of smooth, uninterrupted motion, as in motion pictures (Wertheimer 1912). The visual system can "fill-in" the gaps in the discrete presentation even when the stimuli are separated by up to several degrees of visual arc (Zeeman and Roelofs 1953) and by long temporal intervals (up to 400 milliseconds (Neuhaus 1930)). Why, then, have we chosen a formulation of the motion measurement problem that relies on roughly continuous motion?

More recent psychophysical investigation has suggested that in the human visual system, motion may be analyzed by two different systems, termed short range and long range processes by Braddick (1974, 1980). It has been proposed that the short range process analyzes continuous motion, or motion that is discrete, but with spatial displacements at most 10' — 15' of visual arc, and temporal intervals up to 60 — 100 milliseconds. The long range process would then analyze motion over larger spatial and temporal intervals. The influence of spatial and temporal parameters on the smoothness of perceived motion was investigated by many researchers (for example, Wertheimer 1912; Korte 1915; Holvand 1935; Neff 1936; Corbin 1942; Graham 1965; Kahneman and Wolman 1970; Kolers 1972; Attneave and Block 1973). A distinction was made between the analysis of short range and long range motion because there exists visual patterns that yield qualitatively different perceptions of motion, depending on the range of spatial and temporal displacements between frames. Examples include the Ternus configuration (Ternus 1926; Pantle




Figure 3. Initial Processing of an Image. (a) The original image, containing 320×320 picture elements. (b) The convolution of the image with a $\nabla^2 G$ operator. (c) The resulting zero-crossing contours.

and Picciano 1976), patterns yielding reversed phi (Anstis 1970, 1980; Anstis and Rogers 1975), and dynamic random dot patterns (Regan and Spekreijse 1970; Julesz 1971; Braddick 1974).

The long range motion phenomena illustrate the ability of the human visual system to derive a correspondence between elements in the changing image, over considerable distances and temporal intervals. Under these conditions, there is no continuous motion of elements across the image to be measured directly. Thus, it is likely that a correspondence computation underlies the long range process. Short range motion is roughly continuous, however; we propose that the measurement of motion by the short range process may therefore be appropriately formulated as the computation of an instantaneous velocity field.

2.2 The Initial Motion Measurements

In order to detect movement in a changing image, there must be a variation of intensity over space and time; the combination of the two variations can be used to measure the direction and magnitude of velocity. The explicit comparison of spatial and temporal derivatives of intensity forms the basis of a class of motion measurement schemes referred to as *gradient* schemes (Fennema and Thompson 1979; Horn and Schunck 1981; Marr and Ullman 1981). Other examples of motion detection mechanisms that utilize spatial and temporal intensity changes include those suggested by Hassenstein and Reichardt (1956) and Barlow and Levick (1965).

In principle, motion measurements may be obtained wherever there is a variation of intensity over space and time. Marr and Ullman (1981) proposed, however, that initial motion measurements in the human system are made only at the locations of significant intensity changes. To detect these intensity changes, Marr and Hildreth (1980) proposed that a powerful operator for the

initial filtering of an image is the Laplacian of a Gaussian, $\nabla^2 G$ (approximated in shape by the difference of two Gaussian functions). The elements in the output of an image convolved with $\nabla^2 G$, which correspond to the locations of intensity changes, are the zero-crossings (Marr and Poggio 1979). Figure 3 shows an image that has been convolved with a $\nabla^2 G$ operator, and the resulting zero-crossing contours. Marr and Ullman (1981) proposed that initial motion measurements take place at the locations of these zero-crossings, using a mechanism that combines spatial and temporal gradients of the filtered image (a second mechanism was proposed by Poggio (1983)).

From a computational standpoint, restricting the measurement of motion to the location of significant intensity changes, that give rise to zero-crossings in the filtered image, has two advantages over schemes that base the initial measurements of motion on variations in the original image intensities, wherever the image intensity gradient is nonzero. First, motion measurements are more reliable where the intensity gradient is steeper. The zero-crossings of $\nabla^2 G * I$ correspond to points in the image at which the intensity gradient is locally maximum, yielding the most reliable local velocity measurements. Second, the zero-crossings are tied more closely to physical features; if they move, it is more likely to be the consequence of movement of an underlying physical surface. There are confounding factors, such as changing illumination, that cause intensity to change locally. A scheme that infers motion directly from changing intensity is clearly susceptible to incorrectly inferring motion from changes caused by these confounding factors. While the zero-crossing contours may also be influenced by factors such as changing illumination, their positions are generally more stable.

In two dimensions, the initial measurements face the aperture problem. For the case of contours, local motion measurements provide only the component of motion in the direction perpendicular to the orientation of the contour. The component of velocity along the contour remains undetected. More formally, the two-dimensional velocity field along a contour may be described by the vector function $V(s)$, where s denotes arclength. $V(s)$ can be decomposed into components tangent and perpendicular to the contour, as illustrated in Figure 4. $u^T(s)$ and $u^\perp(s)$ are unit vectors in the directions tangent and perpendicular to the curve, and $v^T(s)$ and $v^\perp(s)$ denote the two components:

$$V(s) = v^T(s)u^T(s) + v^\perp(s)u^\perp(s)$$

The component $v^\perp(s)$, and direction vectors $u^T(s)$ and $u^\perp(s)$, are given directly by the initial measurements from the changing image. The component $v^T(s)$ is not, and must be recovered, to compute $V(s)$. Intuitively, the set of measurements given by $v^\perp(s)$ over an extended contour




Figure 4. Decomposition of Velocity. The local velocity vector $V(s)$ is decomposed into components perpendicular and tangent to the curve. $u^\perp(s)$ and $u^T(s)$ are unit direction vectors, and $v^\perp(s)$ and $v^T(s)$ are the two velocity components.

should provide considerable constraint on the motion of the contour. Additional constraint is still required, however, to determine this motion uniquely. Thus the computation of $V(s)$ requires the integration of the constraints provided by $v^\perp(s)$ along the contour, together with additional constraints necessary to compute $V(s)$ uniquely.

2.3 Additional Assumptions for Motion Measurement

Much of the previous work in motion analysis assumes pure translation of objects in the image plane. Most gradient schemes, for example, assume that velocity is constant over an area of the image (Fennema and Thompson 1979; Marr and Ullman 1981). For gradient schemes, the constraint on velocity imposed by a single measurement of $v^\perp(s)$ can be illustrated graphically in *velocity space*, in which the x and y axes represent the x and y components of velocity, which we denote by V_x and V_y , shown in Figure 5. When mapped to velocity space, the velocity vector at a point on the contour must terminate along the line l perpendicular to the vector $v^\perp(s)u^\perp(s)$; examples are shown by the dotted arrows. For the case of uniform translation, the lines of constraint formed by the measurement of $v^\perp(s)$ along a contour intersect at a single point in velocity space. Some schemes for motion measurement make explicit use of this intersection point (Fennema and Thompson 1979; Thompson and Barnard 1981; Adelson and Movshon 1982). Marr and Ullman (1981) proposed a zero-crossing based scheme, in which each local motion measurement restricts the true direction of velocity of a patch to lie within a 180° range of directions to one side of the zero-crossing contour. A set of measurements taken at different orientations along the contour further restrict the allowable velocity directions, until a single direction is obtained, which is consistent with all the local measurements.




Figure 5. Velocity Constraints in Velocity Space. The x and y axes represent the x and y components of velocity, V_x and V_y . $v^\perp(s)$ is the perpendicular component of velocity, and $u^\perp(s)$ is the unit perpendicular direction vector, at a point p on the image curve. The velocity vector at p must project to the line L ; examples are shown with dotted lines.

Motion measurement schemes based on the cross-correlation of intensity, used both in computer vision (Smith and Phillips 1972; Leese, Novak and Taylor 1970; Lillestrand 1972; Wolferts 1974) and in modelling biological vision systems (Reichardt 1961; Anstis 1970, 1980; Bell and Lappin 1973; Lappin and Bell 1976; Pantle and Picciano 1976; Petersik, Hicks and Pantle 1978), also rely on pure translation. In addition, several correspondence schemes assume local translation of features in the image (for example, Potter 1975, 1977; Lawton 1983).

Some motion measurement schemes allow objects to undergo rigid rotation and translation in the image plane (for example, Davis, Wu and Sun 1982; Nagel 1982; Ullman and Hildreth 1983). For the case of contours moving rigidly in the image, if the direction of velocity is known at two points on the contour, then the direction of velocity may be obtained everywhere, using a simple geometric construction (Ullman and Hildreth 1983). If, in addition, $v^\perp(s)$ is known along the contour, the full velocity field $V(s)$ may be computed.

Methods for motion measurement that assume rigid motion in the image plane may be useful for the initial detection and rough measurement of motion in the periphery, the analysis of motion during smooth pursuit eye movements, or the recovery of observer motion from optical flow (Prazdny 1980; Longuet-Higgins and Prazdny 1981; Bruss and Horn 1983; Lawton 1983). Analysis of the projected motion of arbitrary surfaces undergoing rigid and nonrigid motion in space requires a more general assumption.

3. THE SMOOTHNESS CONSTRAINT

In this section we derive a more general constraint on the velocity field, that allows the computation of the projected motion of three-dimensional surfaces that move freely in space, and deform over time. We rely on the physical assumption that the real world consists predominantly of solid objects, whose surfaces are generally smooth compared with their distance from the viewer. A smooth surface in motion usually generates a smoothly varying velocity field. Thus, intuitively, we seek a velocity field that is consistent with the motion measurements derived from the changing image, and which varies smoothly. Unfortunately, there is an infinity of velocity fields that satisfy these two properties. Horn and Schunck (1981), in their work on the optical flow computation, suggest that a single solution may be obtained by finding the velocity field that varies as little as possible. In the remainder of this section, we show how this constraint may be formulated more precisely, in a way that guarantees a velocity field solution that is mathematically unique and physically plausible.

3.1 Measuring Variation in Velocity

To find the velocity field that varies the least, some means of measuring the variation in velocity along a contour is required. This can be accomplished in many ways. For example, we could measure the change in direction of velocity as we trace along the contour. Total variation in velocity could then be defined as the total change in direction over the entire contour. A second possibility is to measure the change in magnitude of velocity along the contour. Third, the change in the full velocity vector could be measured, incorporating both the direction and magnitude of velocity. Other measures are also possible. The goal of the computation is to find a velocity field that is consistent with the changing image, and minimizes one of these measures of variation in velocity.

A measure of variation may be described more formally by defining a mathematical functional, Θ , which maps the space of all possible vector fields (along the contour), V , into the real numbers: $\Theta:V \mapsto \mathbb{R}$. This functional should be such that the smaller the variation in the velocity field, the smaller the real number assigned to it. Two candidate velocity fields may then be compared, by comparing their corresponding real numbers. This formulation allows the development of an explicit method for computing the velocity field of least variation.

A set of functionals can now be derived, based on the measures of variation that were previously mentioned informally: (1) variation in the full velocity vector, $V(s)$, (2) variation in the direction of velocity, and (3) variation in the magnitude of velocity, all with respect to the curve.




Figure 6. Measuring Variation in Velocity. (a) The velocity vectors $V(s)$ are displayed at two nearby points on the image curve C . (b) The velocity vectors drawn in velocity space, where $\frac{\partial V}{\partial s}$ is indicated by the dotted arrow. (c) The direction of velocity for points on the contour is represented by the angle ϕ . (d) The velocity vectors of (c) are drawn in velocity space, where $\frac{\partial \phi}{\partial s}$ is shown.

1. Variation in $V(s)$

The local variation of $V(s)$ with respect to the contour is given by $\frac{\partial V}{\partial s}$. A scalar measure may be obtained by taking its magnitude: $\left| \frac{\partial V}{\partial s} \right|$. In Figure 6a, two nearby velocity vectors along the image contour C are shown. The vectors are translated to a common origin in velocity space in Figure 6b, where the vector $\frac{\partial V}{\partial s}$ is shown as a dotted arrow. A measure of the total variation in the velocity field over an entire contour may be derived by integrating this local measure, suggesting a functional such as:

$$\Theta(\mathbf{V}) = \int \left| \frac{\partial \mathbf{V}}{\partial s} \right| ds$$

Variations on this functional may also be considered, involving higher order derivatives, or higher powers, such as:

$$\Theta(\mathbf{V}) = \int \left| \frac{\partial^2 \mathbf{V}}{\partial s^2} \right| ds \quad \text{or} \quad \Theta(\mathbf{V}) = \int \left| \frac{\partial \mathbf{V}}{\partial s} \right|^2 ds$$

2. Variation in Direction

Let the direction of velocity be described by the angle φ , measured in the counterclockwise direction from the horizontal, as shown in Figure 6c. In Figure 6d, the local change in direction for two nearby velocity vectors along the image contour, given by $\frac{\partial \varphi}{\partial s}$, is shown in velocity space. Total variation in direction along the contour may again be obtained by integrating this local measure, leading to functionals such as the following:

$$\Theta(\mathbf{V}) = \int \left| \frac{\partial \varphi}{\partial s} \right| ds$$

or variations involving higher order derivatives, or higher powers.

3. Variation in Magnitude

Finally, the total change in magnitude of velocity alone could be measured, using functionals such as:

$$\Theta(\mathbf{V}) = \int \frac{\partial |\mathbf{V}|}{\partial s} ds$$

Again, we could also consider variations on this measure.

The functional that is used to measure variation may also incorporate a measure of the velocity field itself, rather than strictly utilizing changes in the velocity field along the contour, by incorporating a term which is a function of $|\mathbf{V}|$. This might be useful if we sought a velocity field that also exhibits the least total motion. In addition, the functional could become arbitrarily complex in its combination of $\left| \frac{\partial \mathbf{V}}{\partial s} \right|$, $\left| \frac{\partial \varphi}{\partial s} \right|$, $\frac{\partial |\mathbf{V}|}{\partial s}$, or higher order derivatives.

Given that there are many possible measures of variation, what criteria can be used to choose a single measure? First, from a mathematical point of view, there should exist a unique velocity field that minimizes the particular measure of variation; this requirement imposes a set of mathematical constraints on the functional. Second, the velocity field computation should yield solutions that are

physically plausible. Third, if we suggest that such a constraint underlies the motion computation in the human visual system, the minimization of this measure of variation should yield a velocity field that is consistent with human motion perception. These three criteria are important for any additional assumption that is proposed for the motion measurement computation.

3.2 Mathematical Uniqueness of the Velocity Field

An examination of these measures of variation from a mathematical viewpoint suggests that a measure incorporating the change in the full velocity vector is necessary for the velocity field computation. The use of functionals that incorporate only a measure of direction or magnitude of velocity, for example, does not in general lead to a unique velocity field solution (Hildreth 1983). It can be shown, however, that given a simple condition on the constraints that are derived from the image, there exists a unique velocity field that satisfies these constraints, and minimizes the particular measure of variation given by: $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$. To obtain this result, we take advantage of the analysis used by Grimson (1981) for evaluating possible functionals for performing surface interpolation from sparse stereo data. The basic mathematical question is, what conditions on the form of the functional, and the structure of the space of velocity fields, are needed to guarantee the existence of a unique solution? These conditions are captured by the following theorem from functional analysis (see also (Rudin 1973)):

Theorem: *Suppose there exists a complete semi-norm Θ on a space of functions H , and that Θ satisfies the parallelogram law. Then, every nonempty closed convex set $E \subset H$ contains a unique element v of minimal norm, up to an element of the null space. Thus, the family of minimal functions is*

$$\{v + s \mid s \in S\}$$

where

$$S = \{w \mid v + w \in E\} \cap \mathcal{N}$$

and \mathcal{N} is the null space of the functional

$$\mathcal{N} = \{u \mid \Theta(u) = 0\}.$$

In Appendix I, we first show that the functional $\Theta = \{\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds\}^{\frac{1}{2}}$ is a complete semi-norm that satisfies the parallelogram law. Second, we show that the space of possible velocity fields that satisfy the constraints derived from the changing image, is convex. It then follows from the above theorem that this space contains a unique element of minimal norm, up to possibly an element of the null space. The smoothness measure is non-negative, so that minimizing $\{\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds\}^{\frac{1}{2}}$ is equivalent to minimizing $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$.




Figure 7. Uniqueness of the Velocity Field. (a) The constraint imposed by a single measurement of $v^\perp(s)$ on the contour C . Only a uniform translation along the line L_1 can be added to the velocity field. (b) The constraint imposed by two measurements of $v^\perp(s)$ on C .

The null space in this case is the set of constant velocity fields, because the condition that $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds = 0$ implies that $|\frac{\partial \mathbf{V}}{\partial s}| = 0$ everywhere, which implies that $\mathbf{V}(s)$ is constant. Suppose there is a point $(x(s_i), y(s_i))$ on the contour, where $v^\perp(s_i)$ is known. This measurement constrains the velocity $\mathbf{V}(s_i)$ to lie along the line L_1 parallel to the tangent of the contour at this point, as shown in Figure 7a. Suppose there is a velocity field that is consistent with $v^\perp(s_j)$. A uniform translation component can then be added only along the direction of L_1 , in order to obtain another velocity field that is still consistent with this local measure. If a second measurement $v^\perp(s_j)$ is known at a point $(x(s_j), y(s_j))$, for which the direction of the tangent is different (see Figure 7b), then a uniform translation component can be added only along this second direction, in order to obtain a velocity field that is still consistent with $v^\perp(s_j)$. A uniform translation cannot be added to the entire velocity field, which is consistent with both local measurements. Thus, we conclude the following:

If $v^\perp(s)$ is known everywhere along the contour, and there exists at least two points at which the local orientation of the contour is different, then there exists a unique velocity field that satisfies the known velocity constraints and minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$.

An extended straight line does not yield measurements at two different orientations, but in all other cases, there is sufficient information along a contour to guarantee a unique velocity field

solution. The smoothness constraint can be used to compute a projected two-dimensional velocity field for any three-dimensional surface, whether rigid or nonrigid, undergoing general motion in space. While it is not yet clear whether this general formulation of the smoothness constraint, or the particular measure $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$, is the most appropriate for the motion computation, it is important that this measure satisfies certain essential mathematical requirements, that the other measures do not. It is essential that the computation underlying the measurement of motion be mathematically well-founded.

3.3 Physical Plausibility of the Velocity Field Solution

The second criterion for evaluating a particular measure of variation in velocity is the physical plausibility of the resulting solution. One question that can be asked is, under what conditions will the velocity field that minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$ be the correct physical velocity field? If we assume orthographic projection of the scene onto the image, there are at least two classes of motion for which this is true. The first consists of arbitrary rigid objects undergoing pure translation. In this case, $\frac{\partial \mathbf{V}}{\partial s} = 0$ everywhere along contours in the image, and hence $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds = 0$. Since zero is the smallest value that the measure can obtain, it follows that if there exists a valid solution that is consistent with pure translation, then this solution minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$. Consequently, motion measurement schemes that rely on pure translation (such as Fennema and Thompson 1979; Marr and Ullman 1981; Adelson and Movshon 1982) address a special case of this more general method.

The second class of motions includes rigid polyhedra, undergoing general motion in space. In Appendix II, we show the following:

Suppose that a rigid three-dimensional object, consisting of straight lines intersecting in space, projects onto the image plane, using orthographic projection, in such a way that line intersections are preserved (that is, two lines intersect in the image if and only if their generators intersect in space). Further, suppose that this object undergoes a general displacement in space. Then the two-dimensional velocity field that satisfies $v^\perp(s)$ measured along lines in the image, and minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$, is the correct projected two-dimensional velocity field.

The proof of this result takes advantage of the fact that the velocity field varies linearly along a straight line that is moving rigidly in space. The rigid motion of simple polygons in the image plane is a special case of this class of motions, for which an algorithm that minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$ is guaranteed to compute the correct velocity field.

Recently, Yuille (1983) derived a general condition under which the velocity field that minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$ is the correct velocity field. Let $V(s)$ denote the true projected two-dimensional

velocity field for a curve in motion, and let $T(s)$ denote the tangent vector along the curve. If the following relationship holds at every point on the curve:

$$T \cdot \frac{\partial^2 V'}{\partial s^2} = 0$$

then the velocity field $V(s)$ that satisfies the constraints imposed by $v^\perp(s)$ and minimizes $\int |\frac{\partial V}{\partial s}|^2 ds$ is the true velocity field $V'(s)$. The two classes of motion mentioned above correspond to cases for which $\frac{\partial^2 V'}{\partial s^2} = 0$ along the curve, so that this general condition holds trivially.

For the class of smooth curves, moving arbitrarily in space, the velocity field of least variation is, in general, not the physically correct one. Empirical studies suggest, however, that it is often qualitatively similar (Hildreth 1983). Where the two velocity fields differ significantly, it appears that the smoothest velocity field may be more consistent with human motion perception. This claim is investigated further in the next section.

A possible constraint on the velocity field, that is not considered explicitly, is the rigidity of the underlying surface. The computation of a *smoothest* velocity field does not necessarily seek a solution that corresponds to rigid motion, in either two or three dimensions. This may at first seem physically implausible. When a three-dimensional curve rotates in space, however, its two-dimensional projection may undergo significant distortion in the image. Without knowing the three-dimensional structure of the curve, it is very difficult, if not impossible, to find a two-dimensional velocity field that corresponds to a single rigid motion in three dimensions. It is also the case that some of the motion that we encounter arises from surfaces that are nonrigid. If the analysis of motion is a two-stage process, with the measurement of two-dimensional motion preceding the derivation of three-dimensional structure from motion, a constraint such as smoothness may be the most restrictive type of constraint that may be used, which yields a unique solution, and still allows the analysis of general motion.

4. AN ALGORITHM AND EXAMPLES

The velocity field computation has been formulated as an optimization problem. We seek a solution that satisfies the constraints imposed by $v^\perp(s)$, and minimizes the measure of variation given by: $\int |\frac{\partial V}{\partial s}|^2 ds$. The computation may also be described as seeking a solution for which neighboring velocity vectors are as similar as possible. To further test the adequacy of this approach, it is necessary to specify an algorithm, and examine its results for a number of motion sequences. An important aspect of this formulation is that it lends itself naturally to algorithms that are biologically feasible, in that they involve simple, local, parallel operations (Ullman 1979b; Grimson




Figure 8. Pure Translation. (a) and (c) The arrows represent a sampling of the true velocity field for a translating polygon and ellipse. (b) and (d) The initial perpendicular velocity vectors.

1981; Marr 1982). To implement the computation, we chose a standard iterative algorithm from mathematical programming, known as the *conjugate gradient algorithm* (Luenberger 1973). This particular algorithm is certainly not appropriate as a model for human vision. Our aim is to test the basic idea of computing the velocity field of least variation. If the results of the algorithm support the feasibility of this idea, from a physical and perceptual viewpoint, we can then explore alternative algorithms to implement the theory, that are more appropriate for the human system.

A detailed account of the application of the conjugate gradient algorithm to the velocity field computation is given in Hildreth (1983). Because the image is discrete, an image contour consists of a set of n discrete points. The input to the algorithm is the set of n perpendicular components of velocity along the contour. The output of the algorithm is the set of x and y components of velocity at the n points. The algorithm computes the velocity components that minimize a discrete correlate to the continuous functional, subject to the constraint imposed by $v^\perp(s)$.




Figure 9. Rotating Polygon. (a) The true velocity field for a polygon rotating rigidly in the image about the point O . (b) The initial perpendicular velocity vectors. (c) The computed velocity field of least variation.

4.1 Ideal Smooth Curves

We begin with some simple curves, undergoing rigid motion. For this first set of examples, the curves and their perpendicular components of velocity were generated analytically from a known velocity field, and therefore represent ideal input for the algorithm. Many of the examples were chosen because perceptual studies indicate that human observers see motions that differ from the true motion of the curves.

1. Rigid Translation in the Image Plane

In Figures 8a and 8c, a sampling of the true velocity field is shown for two curves, a polygon and an ellipse, translating rigidly in the image plane. Figures 8b and 8d illustrate the perpendicular velocity vectors $v^\perp(s)\mathbf{u}^\perp(s)$, which form the input to the algorithm. In this case, the velocity fields of least variation (not shown here) are identical to the true velocity fields, as expected.

2. Polygon Rotating in the Image Plane

Figure 9 illustrates the true velocity field, initial perpendicular velocity vectors, and smoothest




Figure 10. Rotating Ellipse. (a) The true velocity field for an ellipse rotating rigidly in the image about the point O . (b) The initial perpendicular velocity vectors. (c) The computed velocity field of least variation.

velocity field for a polygon rotating rigidly in the image plane, about the point O . From the theoretical results of Section 3.3 and Appendix II, we expect to obtain the correct velocity field in this case. Figure 9c shows that the computed smoothest velocity field is in fact the correct one.

3. Rotating Ellipse

Figure 10 illustrates the true velocity field, initial perpendicular velocity vectors, and smoothest velocity field for an ellipse rotating rigidly in the image plane, about the point O . In this case, the smoothest velocity field is quite different from the true velocity field. There is a reduced rotational component of velocity, and added radial component in the computed velocity field. The difference in total variation is significant.

At first glance, one might not consider the smoothest velocity field in this case to be a plausible solution. In some earlier perceptual experiments by Wallach, Weisz and Adams (1956), however, they noted that a rigid ellipse does not appear rigid under rotation; it appears to deform continuously. In their experiments, simple geometric figures were placed on a rotating turntable, and observers described the perceived motion of the figures while fixating the center of the turntable (conditions of free and tracking eye movements were also used). Ellipses of various




Figure 11. Rotating Ellipses of Different Aspect Ratios. (a) The true velocity field for a rotating ellipse with aspect ratio 23:23.5. (b) The computed velocity field of least variation for the ellipse in (a). (c) The true velocity field for a rotating ellipse with aspect ratio 5:1. (d) The computed velocity field of least variation for the ellipse in (c).

aspect ratios were observed. Wallach, Weisz and Adams (1956) found that when an ellipse whose axes measured 25 and 23.5 cm was rotated about its center, it appeared to stand still while its contour pulsated. The largest effects were observed for an ellipse whose aspect ratio was 3:2; the entire figure appeared fluid, undergoing a strong deformation, as well as a rotation. For some observers, the deformation was more restricted, and did not occur in the immediate vicinity of the poles of the major axis of the ellipse. As the aspect ratio of the ellipse was increased, it appeared more rigid. The perceived deformation of the ellipse was the same, regardless of whether it rotated about its center, or was placed eccentrically on the turntable.

For the ellipse of Figure 10, the aspect ratio is 2:1. The computed velocity field of least variation clearly implies a significant distortion of the contour, in addition to a rotation. In the immediate vicinity of the poles of the major axes, the smoothest velocity field is very similar to




Figure 12. Rotating Spiral. (a) The true velocity field for a logarithmic spiral rotating in the image about the point O . (b) The initial perpendicular velocity vectors. (c) The computed velocity field of least variation.

the true velocity field. In Figure 11, the true and smoothest velocity fields are shown for rotating ellipses whose aspect ratios are 25:23.5 and 5:1. When the ellipse is nearly circular, the smoothest velocity field indicates a strong inward and outward pulsation of the contour (Figure 11b). The smoothest velocity field for the narrower ellipse, shown in Figure 11d, is closer to the true velocity field than in the case of the ellipse with aspect ratio 2:1, implying less distortion of the contour. Finally, the smoothest velocity field for an eccentrically rotated ellipse differs from that of the centrally rotated ellipse only by the addition of a uniform translation along the contour. Thus, the same deformations of the ellipses is implied by the smoothest velocity field obtained for eccentric rotations. We conclude that the perception of the movement of rotating ellipses is at least qualitatively consistent with the computation of the velocity field of least variation.

4. Rotating Spiral

It is well known that a spiral appears to expand or contract, when it undergoes pure rotation about its center (Holland 1965). The perceived velocity field thus contains a large radial component, while the true velocity field contains only a rotational component of velocity. Figure 12 illustrates the true velocity field, initial perpendicular velocity vectors, and computed smoothest velocity field, for a single arm of a rotating logarithmic spiral. The smoothest velocity field exhibits a large

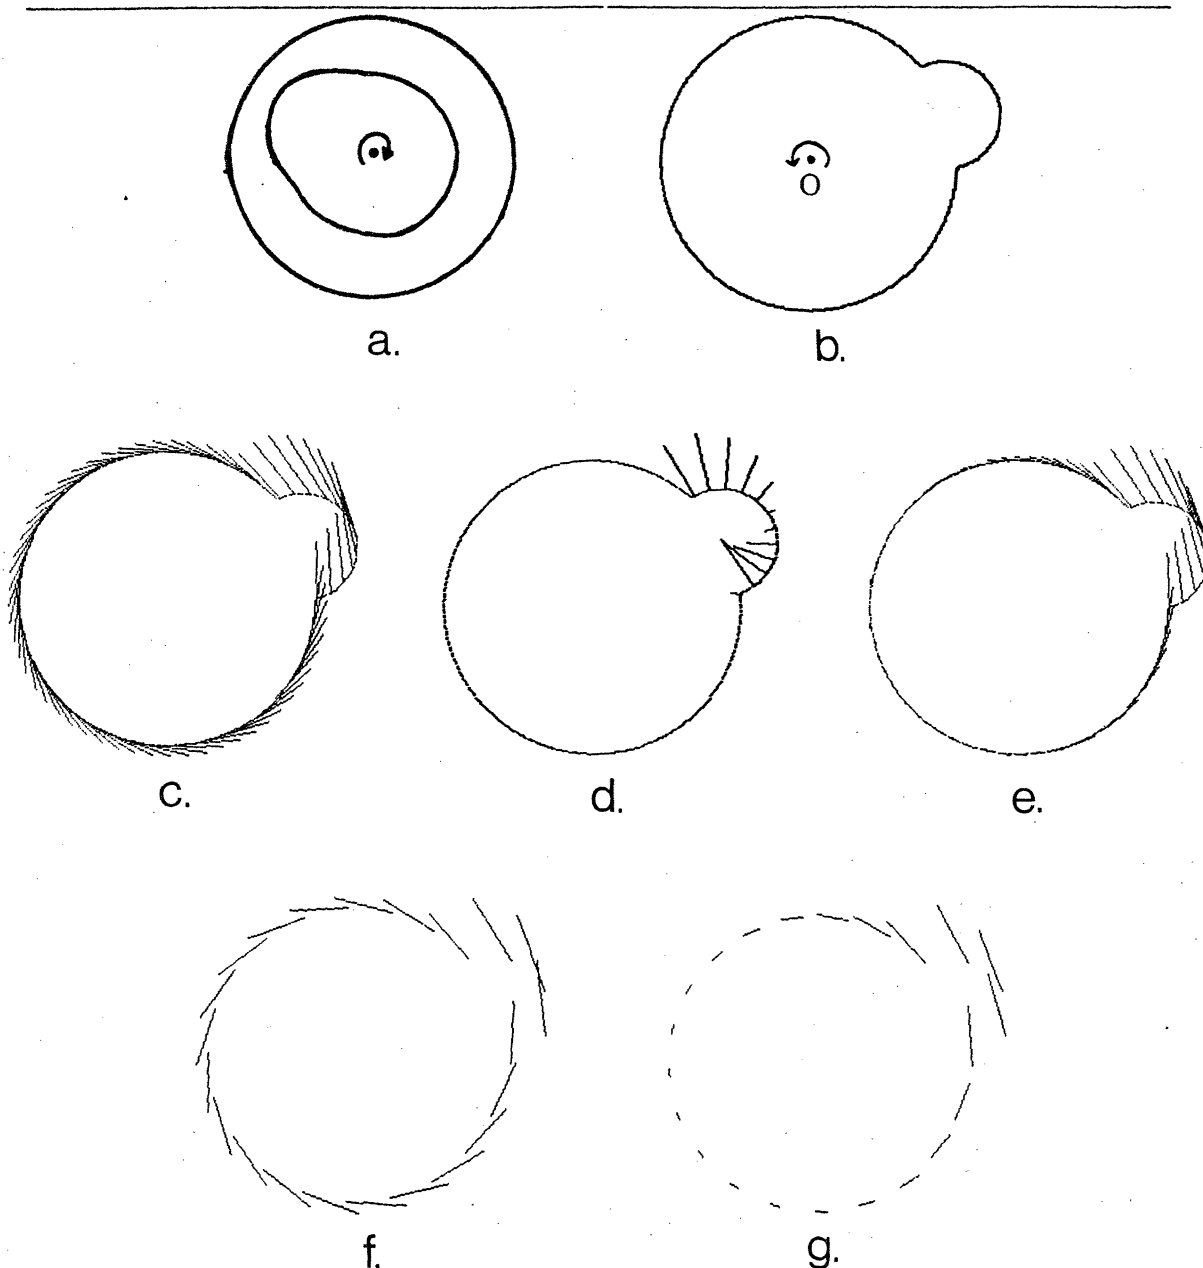


Figure 13. The Deformed Circle. (a) The rotating deformed circle used by Wallach, Weisz and Adams. (b) A deformed circle undergoes pure rotation in the image, about the point O . (c) The true velocity field. (d) The perpendicular velocity vectors. (e) The computed velocity field of least variation. (f) and (g) A coarse sampling of the true and computed velocity fields, without the underlying curve.

radial component of motion at the center of the spiral, which decreases toward the periphery. The perception of the movement of the spiral is qualitatively more consistent with the smoothest velocity field, than with either the true or initial velocity fields.




Figure 14. The Kinetic Depth Effect. (a) The three-dimensional wire-frame stimulus used by Wallach and O'Connell, shown from three different viewpoints. (b) The true projected velocity field that results from rotation of the figure about the central vertical axis. (c) The perpendicular components of velocity. (d) The computed velocity field of least variation.

5. Deformed Circle in Rotation

Wallach, Weisz and Adams (1956) placed a deformed circle at the center of a turntable, and observed its motion as the turntable rotated (Figure 13a). Again, the figure in rotation did not appear rigid. Rather, a continuous deformation took place; the body of the circle appeared to remain still while the deformation traveled about its perimeter. The rotating deformed circle of Figure 13b is based on this demonstration. The true velocity field and perpendicular velocity vectors are shown in Figures 13c and 13d. The perpendicular component of velocity is zero around the circle, due to its symmetry. The computed smoothest velocity field is shown in Figure 13e. To better illustrate the true and computed velocity fields, they are shown in Figures 13f and 13g, without the underlying contour (the velocity field is also more coarsely sampled here). The directions of velocity are very similar for the two velocity fields, but the magnitudes are quite different. In the smoothest solution, the magnitude of velocity decreases significantly as distance from the area of the bump increases, consistent with the perceptual studies.

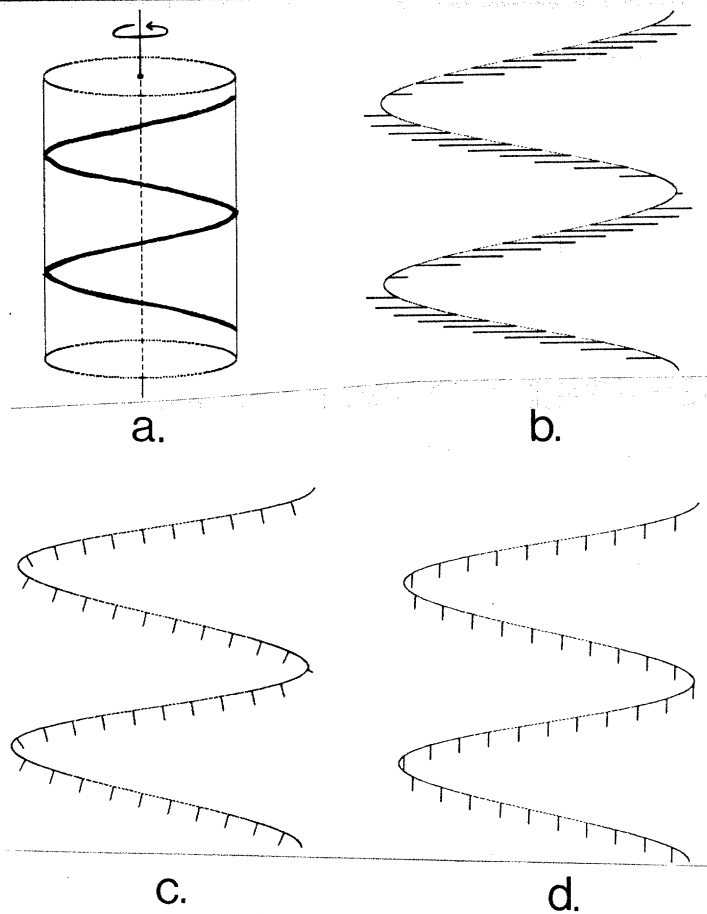


Figure 15. The Barberpole Illusion. (a) A circular helix on an imaginary cylinder, rotating about the vertical axis of the cylinder. (b) The two-dimensional projection of the helix and its velocity field. (c) The initial perpendicular velocity vectors. (d) The computed velocity field of least variation.

6. The Kinetic Depth Effect

The ability of the human visual system to interpret the three-dimensional structure of objects from their projected two-dimensional motion was first explored systematically by Wallach and O'Connell (1953) in a study of what they termed the *kinetic depth effect*. In their experiments, an unfamiliar wireframe object was rotated behind a translucent screen, and the shadow of its parallel projection was viewed from the other side of the screen. Given the projection of the wireframe figure shown in Figure 14a, for example, observers were able to give a correct description of the hidden object's three-dimensional structure. This particular figure is an example of an object for which it can be proven that the true and smoothest velocity fields are equivalent; this equivalence is demonstrated in Figures 14b and 14d. The perception of the correct three-dimensional structure of this figure suggests that the human visual system also derives its correct two-dimensional motion.

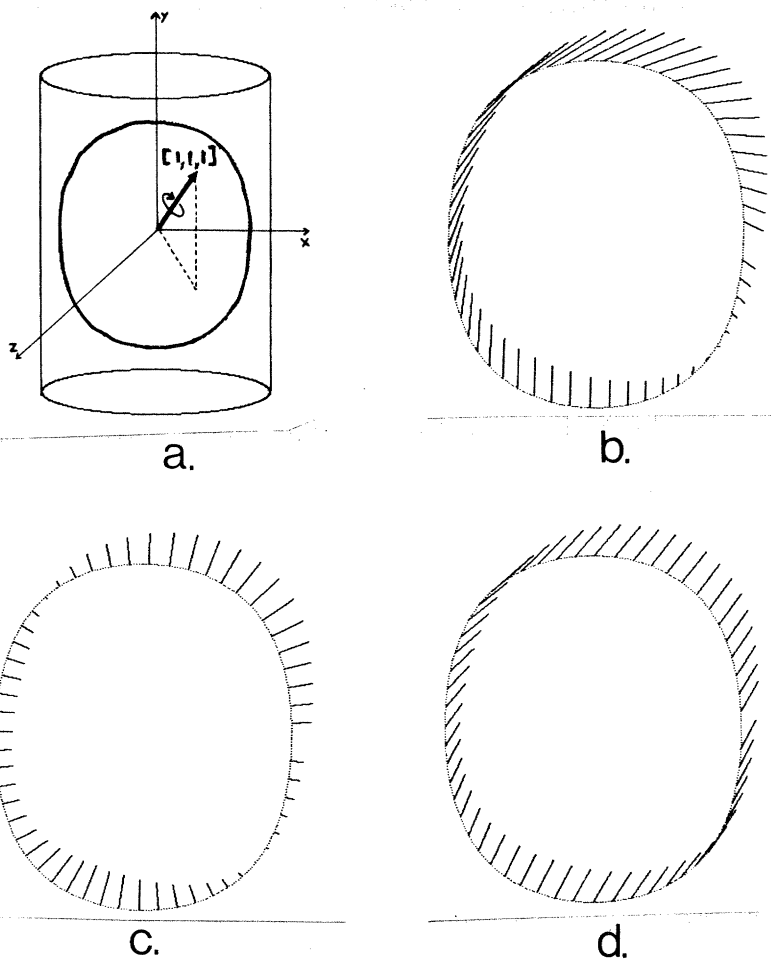


Figure 16. Smooth Curve in Rotation. (a) A circle, wrapped around a cylinder, rotating around the axis $\mathbf{n} = [1, 1, 1]^T$. (b) The two-dimensional projection of the curve and its true velocity field. (c) The initial perpendicular velocity vectors. (d) The computed velocity field of least variation.

7. The Barberpole Illusion

In Figure 15a, a circular helix on an imaginary cylinder is shown, rotating about its central vertical axis. The projection of the curve onto the image plane, together with the resulting two-dimensional velocity vectors, are shown in Figure 15b. Figures 15c and 15d show the initial perpendicular velocity vectors, and computed velocity field. The true velocity vectors are strictly horizontal, while those for the smoothest velocity field, which actually correspond to pure translation of the curve, are strictly vertical. The motion of this curve is analogous to the moving stripes of a barberpole, which appear to move downward as the pole rotates.

8. Smooth Curve in Rotation

In Figure 16, the initial contour is a circle that has been wrapped around the surface of an




Figure 17. Smooth Curve in Rotation. The curve of Figure 16 is rotated 90° from its initial position, and then undergoes an instantaneous rotation about the axis $\mathbf{n} = [1, 1, 1]^T$. (a) The two-dimensional projection of the curve and its true velocity field. (b) The initial perpendicular velocity vectors. (c) The computed velocity field of least variation.

imaginary cylinder. This contour is rotated rigidly about the axis $\mathbf{n} = [1, 1, 1]^T$, as shown in Figure 16a. The projected two-dimensional contour and its true velocity field are shown in Figure 16b. The same contour, at different positions in its trajectory, is shown in Figures 17 and 18. (These two contours are rotated 90° and 180° from the initial position shown in Figure 16.) The initial perpendicular velocity vectors are shown in Figures 16c, 17b and 18b. The computed velocity fields of least variation are shown in Figures 16d, 17c and 18c. At the initial position of the curve, shown in Figure 16, the smoothest velocity field is significantly different from the true velocity field, exhibiting a large translational component toward the upper right corner of the image. At other positions of the curve, however, the true and smoothest velocity fields are qualitatively very similar. The analysis of the predicted velocity fields for smooth curves such as this may be useful for the design of perceptual experiments that investigate whether a constraint such as smoothness underlies motion measurement in the human visual system.

4.2 A Natural Motion Sequence

In this section, we present the results of the algorithm for the sequence of aerial photographs shown in Figure 19. When viewing the two images, alternated back and forth in rapid succession, we observe first, that displacements increase in magnitude from the top to the bottom of the image, and second, that there is a rotation between the two frames, with displacements at the top of the image having a larger horizontal component, and those toward the bottom having a large vertical component. To compute the velocity field, the images were first convolved with a $\nabla^2 G$ operator, and zero-crossings detected. At each zero-crossing, $v^\perp(s)$ was computed, using the spatial and temporal gradient of the filtered image. The conjugate gradient algorithm was then used to compute the velocity field along the zero-crossing contours. Because of the large error




Figure 18. Smooth Curve in Rotation. The curve of Figure 16 is rotated 180° from its initial position, and then undergoes an instantaneous rotation about the axis $\mathbf{n} = [1, 1, 1]^T$. (a) The two-dimensional projection of the curve and its true velocity field. (b) The initial perpendicular velocity vectors. (c) The computed velocity field of least variation.




Figure 19. A Natural Image Sequence. (a) and (b) Two natural images, taken in sequence from an airplane. The images contain 256×256 picture elements.

in the input that is typical of natural images, the velocity field solution was only required to approximately satisfy the constraints imposed by $v^\perp(s)$, rather than exactly satisfy these constraints. (Details of this analysis are given in Hildreth (1983).) The results of the velocity field algorithm are shown in Figure 20. Velocity vectors at evenly spaced points are shown in black, superimposed on the original image, shown with reduced contrast. The results compare well with the actual displacements between the two frames.



Figure 20. The Computed Velocity Field. At evenly spaced points in the image, the computed velocity vectors are shown in black, superimposed on the original image of Figure 19a, reduced in contrast.

5. SUMMARY AND CONCLUSIONS

We have presented a computational study of the measurement of visual motion. The problem was formulated as the computation of an instantaneous two-dimensional velocity field from the changing retinal image. This formulation rests on the assumption that motion analysis may be divided into two stages: the first is the measurement of projected two-dimensional motion, and the second is the derivation of three-dimensional structure from motion. It was suggested that this formulation may be appropriate for the short range motion process in human vision.

A computational theory for the derivation of the velocity field was proposed, with three main components. First, initial measurements of motion in the image take place at the location of significant intensity changes, which give rise to zero-crossings in the output of the convolution of the image with a $\nabla^2 G$ operator, as suggested by Marr and Ullman (1981). These initial

measurements provide only the component of velocity in the direction perpendicular to the local orientation of the contours. Second, these initial measurements are integrated along contours to compute the two-dimensional velocity field. A fundamental problem for this integration stage is that the velocity field is not determined uniquely from information available in the changing image. The third component of the theory is then the formulation of the additional constraint of smoothness of the velocity field, based on the physical assumption that surfaces are generally smooth, which allows the computation of a unique velocity field. A theoretical analysis of the conditions under which this computation yields the correct velocity field suggests that the solution is physically plausible. Empirical studies show that in situations for which the true and smoothest velocity fields differ, the smoothest velocity field may be more consistent with human motion perception. It also appears that this formulation of the motion measurement problem may be biologically feasible, in that it leads naturally to algorithms that involve simple, local, parallel operations.

The various aspects of the computational theory are somewhat independent. For example, the suggestion that initial motion measurements take place in the vicinity of significant intensity changes does not rely on the use of zero-crossings; there may be other mechanisms by which intensity changes are located. The need for a second stage of analysis that combines the initial motion measurements is a consequence of the *local* nature of these measurements. Motion detection schemes that are based on more primitive measurements of changing intensity (such as those suggested by Hassenstein and Reichardt (1956; Reichardt 1961), and Barlow and Levick (1965)), as well as those based on more symbolic tokens (Ullman 1979a), also require this combination stage. Finally, the need for additional constraint is fundamental to the analysis of visual motion. The particular constraint of smoothness can be extended to other formulations of the motion measurement problem.

There are many computational aspects of motion measurement that remain to be addressed. For example, if two adjacent surfaces undergo different motions, they generally give rise to a discontinuity in the velocity field along their boundary. Across such a boundary, the physical assumption of surface smoothness is not valid, and may yield error in the computed velocity field. The explicit detection of motion discontinuities is important for the type of computation proposed in this paper. A second area for investigation is the use of multiple channels, with differing spatial and temporal frequency characteristics, in order to increase the range and resolution of velocities that may be analyzed reliably. The present work only addresses the type of analysis that may take place within a single channel. The experiments of Adelson and Movshon (1982) suggest that multiple channels may not be strongly coupled, for the measurement of motion. A third important

aspect of motion measurement is the role of eye movements, both saccadic and smooth pursuit. The analysis of the perceptual demonstrations in Section 4 assumed a fixed eye position.

A final aspect of motion measurement concerns the role of the short range and long range processes in an integrated system. For some visual tasks, such as the detection of a sudden movement, or object segmentation, rough and instantaneous motion measurements may be sufficient (Marr and Ullman 1981). Computational studies of the recovery of three-dimensional structure suggest that a more complete and accurate measurement of motion is required for this task (Ullman 1980, 1983a,b; Prazdny 1980; Clocksin 1980; Longuet-Higgins and Prazdny 1981; Longuet-Higgins 1981). The work presented in this paper suggests that a computation based on short range motion measurements can, in principle, yield a very precise local velocity field. The reliable derivation of three-dimensional structure, however, requires that these measurements be integrated over a large time frame (Ullman 1983b). The long range process may play a role in this integration process. Studies by Petersik (1980) suggest that the long range process may be crucial to the recovery of structure from motion in the human visual system. A second possible role for the long range process is the maintenance of continuity of objects moving through a complex environment, in which objects are sometimes occluded from view (Marr 1982; Braddick 1983), for example, when an animal runs through a forest. This task could also be guided by information supplied by the short range motion process. The general picture that emerges is one expressed by Ullman (1981); that the short range and long range motion processes may interact in complementary ways to produce a flexible and reliable system for motion measurement.

ACKNOWLEDGEMENT: I thank Shimon Ullman for valuable supervision and useful comments on this paper.

REFERENCES

- Adelson, E. H. & Movshon, J. A. 1982 Phenomenal coherence of moving visual patterns. *Nature, Lond.* **300**, 523-525.
- Anstis, S. M. 1970 Phi movement as a subtraction process. *Vision Res.* **10**, 1411-1430.
- Anstis, S. M. 1980 The perception of apparent motion. *Phil. Trans. R. Soc. Lond. B.* **290**, 153-168.
- Anstis, S. M. & Rogers, B. J. 1975 Illusory reversal of visual depth and movement during changes of contrast. *Vision Res.* **15**, 957-961.
- Attneave, F. & Block, G. 1973 Apparent motion in tridimensional space. *Percept. Psychophys.* **13**, 301-307.
- Barlow, H. B. & Levick, R. W. 1965 The mechanism of directional selectivity in the rabbit's retina. *J. Physiol., Lond.* **173**, 477-504.
- Bell, H. H. & Lappin, J. S. 1973 Sufficient conditions for the discrimination of motion. *Percept. Psychophys.* **14**, 45-50.
- Braddick, O. J. 1974 A short-range process in apparent motion. *Vision Res.* **14**, 519-527.
- Braddick, O. J. 1980 Low-level and high-level processes in apparent motion. *Phil. Trans. R. Soc. Lond. B.* **290**, 137-151.
- Braddick, O. J. 1983 Characteristics of the short range motion process in human vision. Paper presented at the Marr Conference, Cold Spring Harbor, NY, April.
- Bruss, A. & Horn, B. K. P. 1983 Passive navigation. *Comp. Vision Graph. Image Processing* **21**, 3-20.
- Clocksink, W. F. 1980 Perception of surface slant and edge labels from optical flow: a computational approach. *Perception* **9**, 253-269.
- Corbin, H. H. 1942 The perception of grouping and apparent motion in visual space. *Arch. Psychol.* **273**.
- Davis, L., Wu, Z. & Sun, H. 1982 Contour-based motion estimation. In *Proceedings: Image Understanding Workshop (Palo Alto, California)*, pp. 124-131. Arlington, Virginia: Science Applications.
- Fennema, C. I. & Thompson, W. B. 1979 Velocity determination in scenes containing several moving objects. *Comp. Graph. Image Processing* **9**, 301-315.
- Graham, C. H. 1965 Perception of movement. In *Vision and visual perception* (ed. C. H. Graham), New York: John Wiley & Sons.
- Grimson, W. E. L. 1981 *From images to surfaces. A computational study of the human early visual system*. Cambridge, Massachusetts: MIT Press.
- Hassenstein, B. & Reichardt, W. 1956 Systemtheoretische analyse der zeit-, reihenfolgen- und vorzeichenbewertung bei der bewegungs-perzeption der russelkafers. *Chlorophanus. Z. Naturf. IIb*, 513-524.
- Hildreth, E. C. 1983 *The measurement of visual motion*. PhD Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology.
- Holland, H. C. 1965 *The spiral after-effect*. Oxford: Pergamon Press.
- Holvand, C. I. H. 1935 Apparent movement. *Psychol. Bull.* **32**, 755-778.
- Horn, B. K. P. & Schunck, B. G. 1981 Determining optical flow. *Artif. Intell.* **17**, 185-203.
- Julesz, B. 1971 *Foundations of cyclopean perception*. Chicago: The University of Chicago Press.

- Kahneman, D. & Wolman, R. 1970 Stroboscopic motion: effects of duration and interval. *Percept. Psychophys.*, **8**, 161-164.
- Kolers, P. A. 1972 *Aspects of motion perception*. New York: Pergamon Press.
- Korte, A. 1915 Kinematoskopische untersuchungen. *Zeitschrift für Psychol.* **72**, 193-296.
- Lappin, J. S. & Bell, H. H. 1976 The detection of coherence in moving random dot patterns. *Vision Res.* **16**, 161-168.
- Lawton, D. T. 1983 Processing translational motion sequences. *Comp. Vision Graph. Image Processing* **22**, 116-144.
- Leese, J. A., Novak, C. S. & Taylor, V. R. 1970 The determination of cloud pattern motion from geosynchronous satellite image data. *Pattern Recognition* **2**, 279-292.
- Lillestrand, R. L. 1972 Techniques for change detection. *IEEE Trans. Computers* **c-21**, 654-659.
- Longuet-Higgins, H. C. 1981 A computer algorithm for reconstructing a scene from two projections. *Nature* **293**, 133-135.
- Longuet-Higgins, H. C. & Prazdny, K. 1981 The interpretation of moving retinal images. *Proc. R. Soc. Lond. B.* **208**, 385-397.
- Luenberger, D. G. 1973 *Introduction to linear and nonlinear programming*. Reading, Massachusetts: Addison-Wesley.
- Marr, D. 1982 *VISION*. San Francisco: W. H. Freeman.
- Marr, D. & Hildreth, E. C. 1980 Theory of edge detection. *Proc. R. Soc. Lond. B.* **207**, 187-217.
- Marr, D. & Poggio, T. 1979 A computational theory of human stereo vision. *Proc. R. Soc. Lond. B.* **204**, 301-328.
- Marr, D. & Ullman, S. 1981 Directional selectivity and its use in early visual processing. *Proc. R. Soc. Lond. B.* **211**, 151-180.
- Nagel, H.-H. 1982 On change detection and displacement vector estimation in image sequences. *Pattern Recognition Letters* **1**, 55-59.
- Neff, W. 1936 A critical investigation of the visual apprehension of movement. *Am. J. Psychol.* **XLVIII**, 1-42.
- Neuhaus, W. 1930 Experimentelle untersuchung der scheinbewegung. *Archiv für die Gesamte Psychologie* **75**, 315-348.
- Pantle, A. J. & Picciano, L. 1976 A multistable display: evidence for two separate motion systems in human vision. *Science* **193**, 500-502.
- Petersik, J. T. 1980 The effect of spatial and temporal factors on the perception of stroboscopic rotation stimulations. *Perception* **9**, 271-283.
- Petersik, J. T., Hicks, K. I. & Pantle, A. J. 1978 Apparent movement of successively generated subjective patterns. *Perception* **7**, 371-383.
- Poggio, T. 1983 Visual Algorithms. In *Physical and biological processing of images* (ed. O. J. Braddick & A. C. Sleight), Berlin: Springer-Verlag.
- Potter, J. L. 1975 Velocity as a cue to segmentation. *IEEE Trans. Systems, Man, Cybernetics* **SMC-5**, 390-394.
- Potter, J. L. 1977 Scene segmentation using motion information. *Comp. Graph. Image Processing* **6**, 558-581.
- Prazdny, K. 1980 Egomotion and relative depth map from optical flow. *Biol. Cybernetics* **36**, 87-102.

- Regan, D. & Spekrijse, H. 1970 Electrophysiological correlate of binocular depth perception in man. *Nature, Lond.* **225**, 92-94.
- Reichardt, W. 1961 Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In *Sensory communication* (ed. W. A. Rosenblith), Cambridge, Massachusetts: MIT Press.
- Rudin, W. 1973 *Functional analysis*. New York: McGraw-Hill.
- Smith, E. A. & Phillips, D. R. 1972 Automated cloud tracking using precisely aligned digital ATS pictures. *IEEE Trans. Computers* **c-21**, 715-729.
- Ternus, J. 1926 Experimentelle untersuchung uber phanomenale identitat. *Psychol. Forsch.* **7** 81-136. Translated in *A source book of gestalt psychology* (ed. W. D. Ellis), 1967, New York: Humanities Press.
- Thompson, W. B. & Barnard, S. T. 1981 Lower-level estimation and interpretation of visual motion. *IEEE Computer*, August, 20-28.
- Ullman, S. 1979a *The interpretation of visual motion*. Cambridge, Massachusetts: MIT Press.
- Ullman, S. 1979b Relaxation and constrained optimization by local processes. *Comp. Graph. Image Processing* **9**, 115-125.
- Ullman, S. 1980 The interpretation of three-dimensional structure from motion. *Proc. R. Soc. Lond. B* **203** 405-426.
- Ullman, S. 1981 Analysis of visual motion by biological and computer systems. *IEEE Computer*, August, 57-69.
- Ullman, S. 1983a Computational studies in the interpretation of structure and motion: summary and extension. *MIT Artif. Intell. Lab. Memo* **706**.
- Ullman, S. 1983b Maximizing rigidity: the incremental recovery of 3-D structure from rigid and rubbery motion. *MIT Artif. Intell. Lab. Memo* **721**.
- Ullman, S. & Hildreth, E. C. 1983 The measurement of visual motion. In *Physical and biological processing of images* (ed. O. J. Braddick & A. C. Sleight), Berlin: Springer-Verlag.
- Wallach, H. 1976 On perceived identity: 1. the direction of motion of straight lines. In *On Perception* (ed. H. Wallach), New York: Quadrangle.
- Wallach, H. & O'Connell, D. N. 1953 The kinetic depth effect. *J. Exp. Psychol.* **45**, 205-217.
- Wallach, H., Weisz, A. & Adams, P. A. 1956 Circles and derived figures in rotation. *Am. J. Psychol.* **69**, 48-59.
- Wertheimer, M. 1912 Experimentelle studien uber das sehen von bewegung. *Zeitschrift fur Psychol.* **61**, 161-265.
- Wolferts, K. 1974 Special problems in interactive image processing for traffic analysis. *Proc. Second International Joint Conf. Pattern Recognition*, 1-2.
- Yuille, A. L. 1983 Mathematics for motion. *MIT Artif. Intell. Lab. Memo*, in preparation.
- Zeeman, W. P. C. & Roelofs, C. O. 1953 Some aspects of apparent motion. *Acta. Psychol.* **9**, 159-181.

APPENDIX I

In this appendix, we use the following theorem to prove the uniqueness of the velocity field solution, as formulated in Section 3:

Theorem: *Suppose there exists a complete semi-norm Θ on a space of functions H , and that Θ satisfies the parallelogram law. Then, every nonempty closed convex set $E \subset H$ contains a unique element v of minimal norm, up to an element of the null space. Thus, the family of minimal functions is*

$$\{v + s \mid s \in S\}$$

where

$$S = \{w \mid v + w \in E\} \cap \mathcal{N}$$

and \mathcal{N} is the null space of the functional

$$\mathcal{N} = \{u \mid \Theta(u) = 0\}.$$

Let V be the space of all real two-dimensional vector-valued functions along a curve that are continuous, and have continuous first partial derivatives. We first show that the functional $\Theta(V) = \left\{ \int \left| \frac{\partial \mathbf{V}}{\partial s} \right|^2 ds \right\}^{\frac{1}{2}}$ is a complete semi-norm on this space that satisfies the parallelogram law. We then show that the subset $W \subset V$ of vector functions that satisfy the image constraints is convex. It then follows that there exists a unique element $w \in W$ of minimal norm, up to possibly an element of the null space.

To show that this functional is a semi-norm that satisfies the parallelogram law, we take advantage of the following: If μ is a semi-inner product, defined on the space V , and $v \in V$, then $\Theta(v) = \mu(v, v)^{\frac{1}{2}}$ is a semi-norm on the space. A semi-inner product is a function $\mu: V \times V \mapsto \mathfrak{R}$ written $\mu(v, w)$ satisfying the following ($w, x, y, z \in V$):

1. $\mu(y, x) = \mu(x, y)$
2. $\mu(x + y, z) = \mu(x, z) + \mu(y, z)$
3. $\mu(ax, y) = a\mu(x, y)$, $a \in \mathfrak{R}$
4. $\mu(x, x) \geq 0$

The functional $\mu(v, w) = \int \left(\frac{\partial v}{\partial s} \cdot \frac{\partial w}{\partial s} \right) ds$ is a semi-inner product:

1. $\mu(y, x) = \int \left(\frac{\partial y}{\partial s} \cdot \frac{\partial x}{\partial s} \right) ds = \int \left(\frac{\partial x}{\partial s} \cdot \frac{\partial y}{\partial s} \right) ds = \mu(x, y)$

$$2. \mu(x + y, z) = \int \left(\frac{\partial(x+y)}{\partial s} \cdot \frac{\partial z}{\partial s} \right) ds = \int \left(\frac{\partial x}{\partial s} + \frac{\partial y}{\partial s} \right) \cdot \frac{\partial z}{\partial s} ds = \int \left\{ \left(\frac{\partial x}{\partial s} \cdot \frac{\partial z}{\partial s} \right) + \left(\frac{\partial y}{\partial s} \cdot \frac{\partial z}{\partial s} \right) \right\} ds$$

$$= \mu(x, z) + \mu(y, z).$$

$$3. \mu(ax, y) = \int \left(\frac{\partial(ax)}{\partial s} \cdot \frac{\partial y}{\partial s} \right) ds = a \int \left(\frac{\partial x}{\partial s} \cdot \frac{\partial y}{\partial s} \right) ds = a\mu(x, y).$$

$$4. \mu(x, x) = \int \left(\frac{\partial x}{\partial s} \cdot \frac{\partial x}{\partial s} \right) ds \geq 0.$$

If $\mu(v, w)$ is a semi-inner product, then $\Theta(v) = \mu(v, v)^{\frac{1}{2}} = \left\{ \int \left(\frac{\partial v}{\partial s} \cdot \frac{\partial v}{\partial s} \right) ds \right\}^{\frac{1}{2}} = \left\{ \int \left| \frac{\partial v}{\partial s} \right|^2 ds \right\}^{\frac{1}{2}}$ is a semi-norm. $\Theta(v)$ also satisfies the parallelogram law:

$$\begin{aligned} [\Theta(v + w)]^2 + [\Theta(v - w)]^2 &= \\ &= \int \left(\frac{\partial(v+w)}{\partial s} \cdot \frac{\partial(v+w)}{\partial s} \right) ds + \int \left(\frac{\partial(v-w)}{\partial s} \cdot \frac{\partial(v-w)}{\partial s} \right) ds \\ &= 2 \int \left(\frac{\partial v}{\partial s} \cdot \frac{\partial v}{\partial s} \right) ds + 2 \int \left(\frac{\partial w}{\partial s} \cdot \frac{\partial w}{\partial s} \right) ds \\ &= 2[\Theta(v)]^2 + 2[\Theta(w)]^2 \end{aligned}$$

It now remains to show that the semi-norm is complete. To do this, it must be shown that every Cauchy sequence $\{V_n\}$, where $V_n \in V$, converges under this semi-norm. We begin with the functional $\Theta(V) = \left\{ \int \left| \frac{\partial V}{\partial s} \right|^2 ds \right\}^{\frac{1}{2}}$, where $\Theta: V \mapsto \mathfrak{R}$. Let $f(V) = \frac{\partial V}{\partial s}$, and define a new functional $\Phi(f) = \left\{ \int |f|^2 ds \right\}^{\frac{1}{2}}$, with $\Phi: \partial V \mapsto \mathfrak{R}$, where ∂V denotes the space of all partial derivatives of elements of V . $\Phi(f)$ is an L^2 norm, and is therefore complete on ∂V . Let $\{V_n\}$ be a Cauchy sequence of two-dimensional vector functions in V , under the norm Θ ; then the sequence $\left\{ \frac{\partial V_n}{\partial s} \right\}$ is a Cauchy sequence in ∂V , under the norm Φ . Since Φ is complete in ∂V , then any Cauchy sequence in ∂V must converge under Φ ; that is, there exists an element $f \in \partial V$ such that:

$$\lim_{n \rightarrow \infty} \Phi \left(\frac{\partial V_n}{\partial s} - f \right) = 0$$

Now, let $V = \int f ds$. V contains only those vector functions along a curve that are continuous, with continuous first partial derivatives; it follows that f and V are continuous, so that $V \in V$. (V may not be unique, but this is to be expected, because the functional Θ is only a semi-norm.) V satisfies the following:

$$\lim_{n \rightarrow \infty} \Theta(V_n - V) = 0$$

Therefore the Cauchy sequence $\{V_n\}$ converges, and our semi-norm is complete.

The second statement that we want to show is that the subset of velocity fields satisfying the image constraints is convex. Let $W \subset V$ be the subset of V that satisfies the constraints derived from the image. These constraints are as follows:

1. $v^\perp(s)$ is known everywhere along the curve.
2. There may exist a set of points $\{p_i\}$ for which $v^\perp(s)$ and $v^\top(s)$ are known.
3. There may exist a set of points $\{p_i\}$ for which the direction of velocity, and hence $\frac{v^\perp(s)}{v^\top(s)}$ is known.

If W is convex, it must be the case that:

$$tw_1 + (1-t)w_2 \in W, \text{ for all } w_1, w_2 \in W, \text{ and } 0 \leq t \leq 1$$

For convenience of notation, we write v^\perp and v^\top for $v^\perp(s)$ and $v^\top(s)$.

1. If $w_1, w_2 \in W$, then $w_1^\perp|_{p_i} = w_2^\perp|_{p_i} = v^\perp|_{p_i}$, hence:

$$[tw_1 + (1-t)w_2]^\perp|_{p_i} = tw_1^\perp|_{p_i} + (1-t)w_2^\perp|_{p_i} = v^\perp|_{p_i}$$

2. Similarly, if $w_1, w_2 \in W$, then $w_1^\top|_{p_i} = w_2^\top|_{p_i} = v^\top|_{p_i}$, hence:

$$[tw_1 + (1-t)w_2]^\top|_{p_i} = tw_1^\top|_{p_i} + (1-t)w_2^\top|_{p_i} = v^\top|_{p_i}$$

3. If the direction of velocity at point p_i is the same for w_1 and w_2 , then any linear combination of w_1 and w_2 yields a velocity vector at p_i that is parallel to this direction.

Finally, if w_1 and w_2 are continuous, with continuous first partial derivatives, then this will also be true of $[tw_1 + (1-t)w_2]$. Therefore, the subset $W \subset V$, which satisfies the image constraints, is convex. This completes the proof.

APPENDIX II

In this appendix, we show the following:

Proposition. *Suppose that a rigid three-dimensional object, consisting of straight lines intersecting in space, projects onto the image plane, using orthographic projection, in such a way that line intersections are preserved (that is, two lines intersect in the image if and only if their generators intersect in space). Further, suppose that this object rotates in space with angular velocity ω about a single axis $\mathbf{n} = [n_x, n_y, n_z]^T$. Then the two-dimensional velocity field that satisfies $v^\perp(s)$ measured along lines in the image and minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$, is the correct two-dimensional projected velocity field.*

We divide our proof into two parts: the first shows that the constraint imposed by two lines meeting at a vertex is sufficient to compute the correct two-dimensional velocity of the vertex; the second shows that if a rigid line segment is moving in space, and the projected two-dimensional velocity vectors at its endpoints are known, then the velocity field satisfying $v^\perp(s)$ along the projection of the line onto the image, which minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$, is the correct projected velocity field. These results are extensions of some previous observations made by Alan Yuille (personal communication). Throughout this appendix, the argument, s , to functions that vary along a curve, has been omitted.

Proposition (a). *Suppose that two lines l_1 and l_2 in space, intersect at the point (x_1, y_1, z_1) , and the configuration rotates rigidly with angular velocity ω about a single axis in space, given by $\mathbf{n} = [n_x, n_y, n_z]^T$. Assume that l_1 and l_2 project onto lines l_1^p and l_2^p in the image plane, using orthographic projection, which have different orientations. Then the constraint provided by v^\perp on l_1^p and l_2^p is sufficient to compute the correct projected two-dimensional velocity at (x_1, y_1, z_1) .*

Proof. First, let α and β denote the orientations of l_1^p and l_2^p in the image plane. The true velocity field, \mathbf{V} , is given by:

$$\mathbf{V} = \omega z \begin{bmatrix} n_y \\ -n_x \end{bmatrix} + \omega n_z \begin{bmatrix} -y \\ x \end{bmatrix}$$

We therefore have $\mathbf{V} = (\omega z n_y - \omega y n_z, -\omega z n_x + \omega x n_z)$. Let \mathbf{u}^\top and \mathbf{u}^\perp denote the unit tangent and normal vectors, and v^\top and v^\perp denote the tangential and normal components of velocity.

We then have:

$$\mathbf{V} = v^\top \mathbf{u}^\top + v^\perp \mathbf{u}^\perp$$

For l_1^p , we have the following:

$$\mathbf{u}^\top = (\cos \alpha, \sin \alpha)$$

$$\mathbf{u}^\perp = (\sin \alpha, -\cos \alpha)$$

$$\mathbf{v}^\top = \mathbf{V} \cdot \mathbf{u}^\top = (\omega z_1 n_y - \omega y_1 n_z) \cos \alpha + (\omega x_1 n_z - \omega z_1 n_x) \sin \alpha$$

$$\mathbf{v}^\perp = \mathbf{V} \cdot \mathbf{u}^\perp = (\omega z_1 n_y - \omega y_1 n_z) \sin \alpha - (\omega x_1 n_z - \omega z_1 n_x) \cos \alpha$$

For l_2^p , we have the following:

$$\mathbf{u}^\top = (\cos \beta, \sin \beta)$$

$$\mathbf{u}^\perp = (\sin \beta, -\cos \beta)$$

$$\mathbf{v}^\top = \mathbf{V} \cdot \mathbf{u}^\top = (\omega z_1 n_y - \omega y_1 n_z) \cos \beta + (\omega x_1 n_z - \omega z_1 n_x) \sin \beta$$

$$\mathbf{v}^\perp = \mathbf{V} \cdot \mathbf{u}^\perp = (\omega z_1 n_y - \omega y_1 n_z) \sin \beta - (\omega x_1 n_z - \omega z_1 n_x) \cos \beta$$

It is assumed that \mathbf{v}^\perp is known, and that the tangential component \mathbf{v}^\top is unknown. The velocity of the point (x_1, y_1, z_1) must satisfy constraints imposed by both lines l_1^p and l_2^p . Let c_1 and c_2 denote the unknown \mathbf{v}^\top along the lines l_1^p and l_2^p at the point (x_1, y_1, z_1) . We then have two equations for the velocity at this point, in terms of the two unknowns c_1 and c_2 . From l_1^p we have:

$$\mathbf{V} = [(\omega z_1 n_y - \omega y_1 n_z) \sin \alpha - (\omega x_1 n_z - \omega z_1 n_x) \cos \alpha] \begin{bmatrix} \sin \alpha \\ -\cos \alpha \end{bmatrix} + c_1 \begin{bmatrix} \cos \alpha \\ \sin \alpha \end{bmatrix}$$

From l_2^p :

$$\mathbf{V} = [(\omega z_1 n_y - \omega y_1 n_z) \sin \beta - (\omega x_1 n_z - \omega z_1 n_x) \cos \beta] \begin{bmatrix} \sin \beta \\ -\cos \beta \end{bmatrix} + c_2 \begin{bmatrix} \cos \beta \\ \sin \beta \end{bmatrix}$$

These two equations are satisfied for:

$$c_1 = (\omega z_1 n_y - \omega y_1 n_z) \cos \alpha + (\omega x_1 n_z - \omega z_1 n_x) \sin \alpha$$

$$c_2 = (\omega z_1 n_y - \omega y_1 n_z) \cos \beta + (\omega x_1 n_z - \omega z_1 n_x) \sin \beta$$

c_1 and c_2 are equivalent to the true \mathbf{v}^\top stated previously; therefore, we have obtained the correct projected two-dimensional velocity at the point (x_1, y_1, z_1) .

Proposition (b). *Suppose a line segment l in space rotates rigidly with angular velocity ω about a single axis, given by $\mathbf{n} = [n_x, n_y, n_z]^T$. Let l^p denote the orthographic projection of l onto the image plane. Further, suppose that the projected two-dimensional velocity vectors at the endpoints, denoted by \mathbf{V}_0 and \mathbf{V}_1 , are known. Then the two-dimensional velocity field that satisfies the constraint given by \mathbf{V}_0 , \mathbf{V}_1 and \mathbf{v}^\perp along l^p , and minimizes $\int |\frac{\partial \mathbf{V}}{\partial s}|^2 ds$, is the correct projected two-dimensional velocity field.*

Proof. First, since \mathbf{u}^\perp and \mathbf{u}^\top are constant along the line segment, we have:

$$\frac{\partial \mathbf{V}}{\partial s} = \frac{\partial v^\perp}{\partial s} \mathbf{u}^\perp + \frac{\partial v^\top}{\partial s} \mathbf{u}^\top$$

Total variation of velocity along the segment is then given by:

$$\int_{s_0}^{s_1} \left| \frac{\partial \mathbf{V}}{\partial s} \right|^2 ds = \int_{s_0}^{s_1} \left[\left(\frac{\partial v^\perp}{\partial s} \right)^2 + \left(\frac{\partial v^\top}{\partial s} \right)^2 \right] ds$$

where s_0 and s_1 denote arclength at the two endpoints. Since $\frac{\partial v^\perp}{\partial s}$ is known, the above expression is minimized when the second term is minimized:

$$\int_{s_0}^{s_1} \left(\frac{\partial v^\top}{\partial s} \right)^2 ds$$

Let v_0^\top and v_1^\top denote the known tangential components of velocity at the endpoints of l . We want to minimize the above functional, subject to these known velocities. We can formulate the problem as an *Isoperimetric* problem, and use the Euler-Lagrange equations to solve it. The general problem can be stated as follows:

$$\text{minimize } I = \int_{s_0}^{s_1} F \left[v^\top, \frac{\partial v^\top}{\partial s}, s \right] ds$$

$$\text{subject to } \int_{s_0}^{s_1} \Phi \left[v^\top, \frac{\partial v^\top}{\partial s}, s \right] ds = c_k$$

In this case, we seek to minimize:

$$\int_{s_0}^{s_1} \left(\frac{\partial v^\top}{\partial s} \right)^2 ds$$

subject to the following constraint:

$$\int_{s_0}^{s_1} \left(\frac{\partial v^\top}{\partial s} \right) ds = v_1^\top - v_0^\top$$

The general solution is given by the Euler-Lagrange equation:

$$\frac{d}{ds} \left(\frac{\partial \Psi}{\partial \left(\frac{\partial v^\top}{\partial s} \right)} \right) - \frac{\partial \Psi}{\partial s} = 0$$

where $\Psi = F + \mu \Phi$. μ denotes the Lagrange multiplier. In this case:

$$F = \left(\frac{\partial v^\top}{\partial s} \right)^2 \quad \Phi = \frac{\partial v^\top}{\partial s}$$

We therefore have the following differential equation:

$$\frac{d}{ds} \left[2 \left(\frac{\partial v^\top}{\partial s} \right) + \mu \right] = 0$$

which leads to the equation:

$$\frac{\partial^2 v^\top}{\partial s^2} = 0$$

The solution is therefore: $v^\top = as + b$. From the known tangential components of velocity at the endpoints, we obtain the following solution:

$$v^\top = \frac{v_1^\top - v_0^\top}{s_1 - s_0} (s - s_0) + v_0^\top$$

Since the true v^\top and computed tangential components are both linear functions that satisfy the same endpoint velocities, they must be the same function. Therefore the smoothest velocity field is the correct projected two-dimensional velocity field.

The original statement of the problem now follows straightforwardly from these two results. In addition, we have the following corollary:

Corollary. Suppose a simple polygon is rotating rigidly in the image plane. Then the two-dimensional velocity field that satisfies v^\perp measured along line segments in the image, and minimizes $\int \left| \frac{\partial \mathbf{V}}{\partial s} \right|^2 ds$, is the correct two-dimensional velocity field.

Finally, the above results also hold for an additional translation of the object in space.