Pergamon

0042-6989(95)00084-4

Vision Res. Vol. 36, No. 2, pp. 259–270, 1996
Copyright © 1996 Elsevier Science Ltd
Printed in Great Britain. All rights reserved
0042-6989/96 $9.50 + 0.00

Detection of First-order Structure in Optic Flow Fields

SUSAN F. TE PAS,*† ASTRID M. L. KAPPERS,* JAN J. KOENDERINK*

Received 27 April 1994; in revised form 29 September 1994; in final form 21 February 1995

We measured psychophysical thresholds for the detection of four different optic flow components in the presence of a translational velocity. We also measured thresholds for detection of rotation in the presence of expansion and for expansion in the presence of rotation. Our stimuli consisted of sparse random dot patterns. Detection thresholds are similar for all four optic flow components. Thus, our experiments indicate that our subjects use a factor that is similar in all first-order flow components, namely the relative orientation of the velocity vectors. First-order components can be extracted independently of both each other and translational velocity.

Optic flow Motion Looming Psychophysics Detection

INTRODUCTION

The *optical flow field* (Gibson, 1950) is a rich source of information about the geometrical structure of the world around us and our position relative to it. Much of this information is contained in the local differential structure of the flow field and not, for instance, in the average speed or direction of the flow. Koenderink and van Doorn (1975) expanded the optic flow field in a Taylor-like fashion to derive the zero- and first order optic flow components.‡ They have mathematically shown that any small deformation of a patch of the optic flow field can be decomposed (in first order, using terms that are linear in velocity) in four basic components. Such a decomposition is common in the kinematics of continuous media. The zero-order optic flow component, the translational velocity, does not deform at all. The four first-order optic flow components are differential invariants for Euclidean motion, which means that they are independent of the coordinate system that is chosen. These differential invariants are the curl, the divergence and the two components of the deformation (Koenderink & van Doorn, 1975, 1976; Koenderink, 1986; Longuet-Higgins & Pradzny, 1980). They are mathematically independent in the sense that together they form a complete orthogonal set in Euclidean space. For orientation and navigation in space, for instance relative to a floor and walls, this first-order differential structure of the flow captures most of the relevant information, although for a task like object recognition you might need to study higher orders, since surface curvature affects the second-order flow structure.

An important problem in visual research is whether our visual system actually makes use of this mathematical decomposition and whether these components are independently processed by the visual system. There are few systematic approaches to this problem. Physiological studies yield evidence that MST cells sensitive to optic flow components (divergence, rotation and in one case deformation) do exist (Orban, Lagae, Verri, Raiguel, Xiao, Maes & Torre, 1992; Tanaka & Saito, 1989; Tanaka, Fukada & Saito, 1989). However, these authors all report that such cells are unable to extract a single component from a more complex flow pattern. As yet, it is not clear what might be the functional role of these cells. Moreover, Graziano, Andersen and Snowden (1994) report that there are also cells that are preferentially tuned to spiral motion (a linear combination of divergence and curl), thus arguing that the hypothesis that there are separate channels for rotation and expansion/contraction only, appears to be incorrect. Tanaka *et al.* (1989) report that the orientation of the velocity vectors seems to be the most important factor in stimulating the MST cells that are sensitive to optic flow components. Psychophysical experiments by Freeman and Harris (1992) support this conclusion. Experiments by Warren, Blackwell, Kurtz, Hatsopoulos and Kalish (1991) where subjects had to indicate their direction of heading also lead to the conclusion that the

---

*Helmholtz Instituut, Utrecht University, Princetonplein 5, 3584 CC Utrecht, The Netherlands [*Tel.* 31 30532809; *Fax* 31 30 522664; *Email* S.F.tePas@fys.ruu.nl].

†To whom all correspondence should be addressed.

‡With *n*th-order optic flow field we mean the *n*th-order term in the Taylor expansion of the optic flow field. Thus, the zero-order flow field is constant in velocity, the first-order optic flow field is linear in the velocity. Similarly, the second-order flow field depends on the velocity in a quadratic manner, etc.

direction of the velocity is the most important element in the stimulus. Psychophysical experiments by Regan and Beverley (1985), and Freeman and Harris (1992) show that the visual system uses more than just local zero-order velocity information.

Combinations of different flow components have been used in psychophysical experiments by Freeman and Harris (1992). They investigated whether the visual system is selectively sensitive to different types of relative motion using a masking technique. Their results show that expansion and rotation are analysed independently and they support the existence of a Relative Motion System. On the other hand, experiments performed by Sekuler (1992) do not support the notion that this Relative Motion System is encoded in a distinct channel. Lappin, Norman and Mowafy (1991) also concluded that rotations and expansions were visually independent, although they did not find this independence for other components. De Bruyn and Orban (1993) added different components transparently, though they also performed some pilot experiments with superimposed flow fields. These experiments indicate that extraction of one component in the presence of another is not a difficult task if both components are superimposed. Kappers, van Doorn and Koenderink (1994) used a systematic approach, presenting rotations in combination with translational velocity. They found thresholds for discrimination of clockwise and counter-clockwise rotation to be independent of translational velocity over the entire range of speeds at which subjects were able to perform the task. Also, dot density and lifetime had little influence on their results. Using the same systematic approach for divergence in the presence of translational velocity Kappers, te Pas and Koenderink (1993) performed experiments in which different aspects of a flow field (e.g. direction of the velocity or the velocity gradient etc.) were presented separately. These experiments indicate that human performance depends on the direction of velocity rather than on the velocity gradient or the exact value of the first-order component (rotation, divergence or deformation). In this paper we address the question whether different mechanisms are used for the detection of different first-order optic flow components. For all four first-order optic flow components we use the same paradigm that was used by Kappers *et al.* (1994) for rotation.

In order to find out whether optic flow components like rotation, divergence and deformation can be extracted independently by the visual system we ran a large number of experiments in which one of these components had to be detected in the presence of a translational velocity. The speed of this translational velocity was varied over a large range (0.03–320 deg/sec). The advantage of such a large number of experiments is that we use similar stimuli for a complete family of flow fields. In this way results can be compared directly. In a second experiment we also added a divergence to a curl in order to find out whether first-order components could be detected independently of each other.

METHODS

Apparatus

Our stimuli are generated on an Atari MegaST4 computer and shown on an Atari SM125 high resolution 70 Hz white phosphor P4 monochrome monitor (luminance 71 cd/m^2). The monitor dimensions are 13.6 × 21.7 cm (400 × 640 pixels). Subjects rest their heads in a chin-rest 34 cm from the screen. Thus the screen area is 22.6 × 35.4 degrees of visual angle and pixel separation is 3.2 min of arc. Experiments are performed monocularly with a natural pupil; in our case all subjects used only their right eye. Subjects were asked to fixate in the middle of the screen, where a fixation cross was presented immediately before and after the presentation of each stimulus. Subjects were placed in a dark room so that they would not be distracted by elements outside the stimulus. The only light came from the computer screen.

Stimulus

Stimuli are similar (in some cases identical) to those used by Kappers *et al.* (1994). We use pseudo-random dot patterns, consisting of dark dots on a light background. The size of the dots is 9.6 min of arc. The dots are situated on a perturbated hexagonal point raster to prevent subjects from recognizing local features arising from (random) clustering. The number of visible dots per frame and the lifetime of the dots determine the grid spacing. The stimulus window is circular with a radius of 190 pixels (10 degrees of visual angle). Moving patterns are generated by presenting sequences of frames stroboscopically. In between two consecutive frames all dots "move with the flow". The number of frames determines the total presentation time. We keep this constant at 16 frames (228 msec). Previous experiments (Kappers *et al.*, 1994), show that the effect of the number of dots per frame and the lifetime of the dots have very little effect on performance. Therefore, for most experiments, we kept this constant at 16 dots per frame with a lifetime of 16 frames. We added some experiments with other numbers of dots (64 and 4 dots) and different lifetimes of the dots (2, 3 and 4 frames) to check whether results were consistent. The start frame of each dot is chosen at random to avoid density cues in the case of a diverging stimulus.

The stimulus movement can be a divergence, a rotation or a deformation, the latter in either of two directions (horizontal or oblique). In the rest of the paper we will refer to these movements as div, curl, def and def45, where def is a deformation with its axis of contraction either horizontal or vertical, and def45 is a deformation with its axis of contraction in one of the diagonal directions. One should bear in mind that when we refer to "the curl", we actually mean "the curl of the (instantaneous) velocity field". In our stimuli the value of the first-order optic flow component is always the same over the entire stimulus area. It is important to note that if the curl has a value of 1 rad/sec this value will not change when we add a translational velocity.

For a given value of the first-order component, we vary the *speed* of the translational velocity inside the stimulus window by adding a vertical translational component. The *direction* of movement is always downward. If we add a translational velocity to one of the first-order optic flow components, the location of the "centre" of the flow field (the singular point where the speed is zero) will change. If the speed of the translational velocity increases, the singular point moves outward. The upper panel of Fig. 1 shows the addition of a clockwise curl and a translational velocity in the downward direction. The lower panel shows the addition of the same curl and a downward translational velocity that is three times as fast. Notice that the value of the curl is the same in both the upper left and right panels and the lower left and right panels. In the central panels the value of the curl is zero. As we chose to keep the movement of the stimulus downward, the location of this singular point is an important cue. Because the movement is very slow close to this singular point, pixel noise might give rise to unwanted stimulus features. Therefore, the singular point is always located outside the stimulus window on either a vertical (div and def) or a horizontal (curl and def45) axis through the centre of the stimulus. We do this by always adding at least the minimum amount of translational velocity that is necessary to keep the singular point outside the stimulus window. The singular point is thus located at eccentricities between 10 deg of visual angle (necessary to keep it outside the stimulus window) and about 80 deg of visual angle (for the maximum translation at which our best subject could still perform the task). Examples of the four different flow fields that are used are presented in Fig. 2a–d. All four figures show 64 dots with a lifetime of 16 frames and a translational velocity of 10 deg/sec. Although most measurements were done with 16 dots per frame, we chose to show these static images with 64 dots per frame because it provides a much clearer picture. In Fig. 2 the consecutive frames are superimposed to create a static image, and the value of the flow components was always one (in their respective units). In the following section we will discuss the meaning of "a curl of 1 rad/sec". From the examples it is easy to see that a combination of translation and curl (in this case 10 deg/sec and 1 rad/sec respectively, see Fig. 2a) leads to the same speed range inside the stimulus window as for instance a similar combination of translation and divergence (10 deg/sec and 1 /sec respectively, see Fig. 2b). Figure 2e shows the same rotation as in Fig. 2a, but this time the average value of the translational velocity is 40 deg/sec. Figure 2f again shows the same rotation and translational velocity as in Fig. 2e, but here there are 4 dots per frame and the lifetime of the dots is 2 frames.

Procedure

In an experimental session we vary only the amount of translational speed and one of the flow components (div, curl, def or def45). The other components are kept at zero. Subjects have to indicate the sign of this flow component, in other words they have to indicate whether the rotation was clockwise or counter-clockwise or whether the divergence is an expansion or contraction. For the deformation subjects have to indicate if the point of zero speed is situated below or above the stimulus window (def) or to the left or to the right of the stimulus window (def45). Psychophysical thresholds are measured by adding a stochastic noise component to the stimulus. A two-dimensional Gaussian perturbation

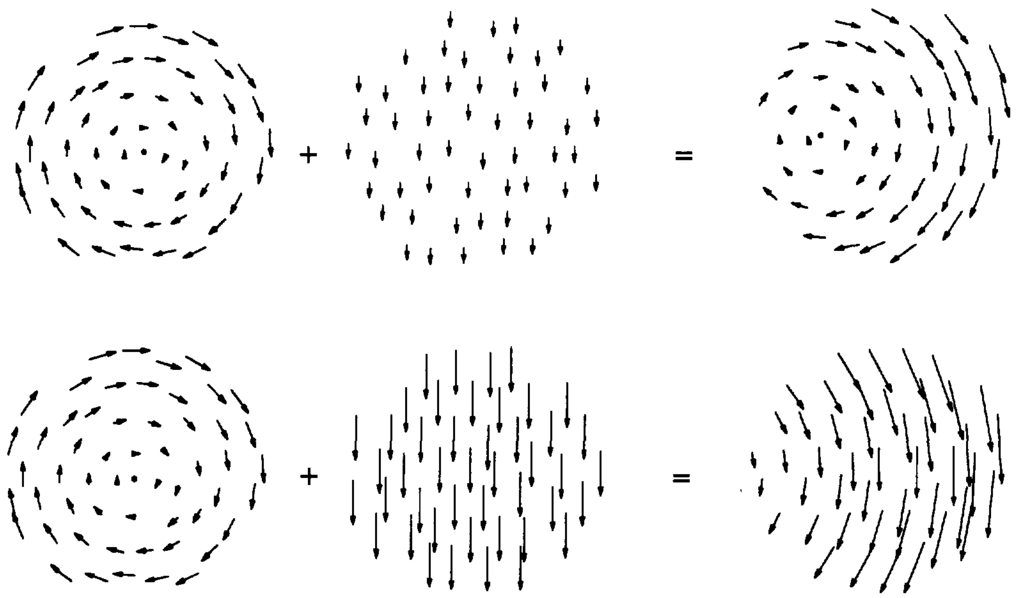

FIGURE 1. Upper panel: the addition of a clockwise rotation and a uniform downward translational velocity. Lower panel: the addition of the same clockwise rotation as above and a downward translational velocity that is three times as fast. Notice that the singular point moves to the left, and that the location of the singular point depends on the value of the translational velocity. In both the upper left and right panels and the lower left and right panels the curl has the same value. In the central panels the curl is zero.

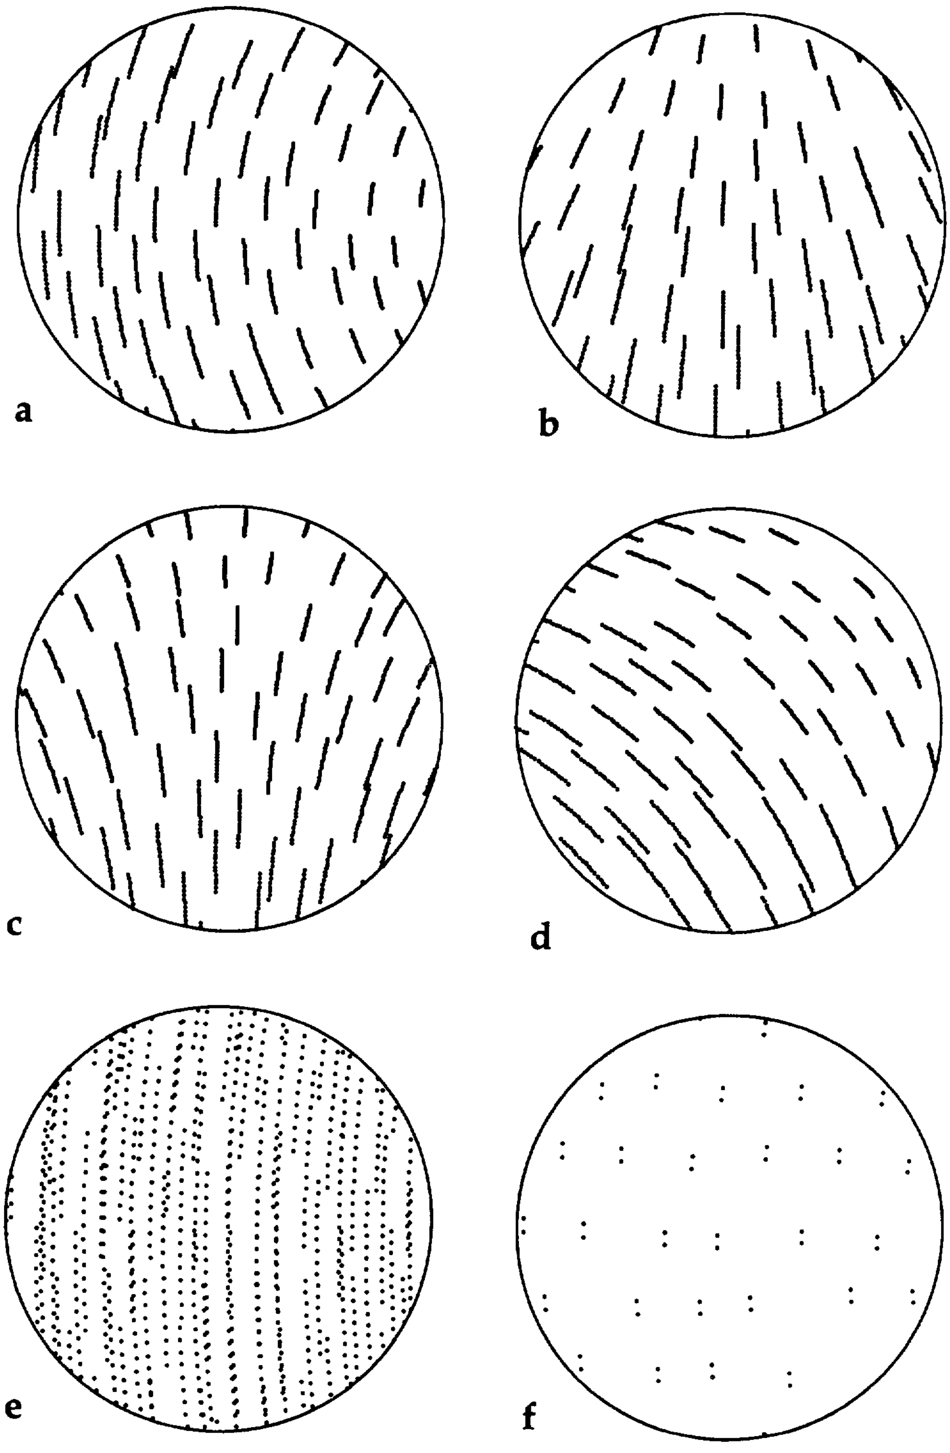

FIGURE 2. The 16 consecutive frames of the film were superimposed to provide static examples of our dynamic stimuli. For all examples the noise level is zero degrees. The stimuli differ in flowtype, translational velocity, number of dots and lifetime: (a) A rotating stimulus, with curl = 1 rad/sec, translational velocity = 10 deg/sec and 64 dots per frame with a lifetime of 16 frames. (b) An expanding stimulus, with div = 1/sec, translational velocity = 10 deg/sec and 64 dots per frame with a lifetime of 16 frames. (c) A deforming stimulus, with def = 1/sec, translational velocity = 10 deg/sec and 64 dots per frame with a lifetime of 16 frames. In this case the axis of contraction is horizontal. (d) The same as in (c) except that the axis of contraction is diagonal. (e) A rotating stimulus, with curl = 1 rad/sec, translational velocity = 40 deg/sec and 64 dots per frame a lifetime of 16 frames. (f) A rotating stimulus, with curl = 1 rad/sec and translational velocity = 40 deg/sec. In this case there are 4 dots per frame with a lifetime of 2 frames.

vector with variable amplitude is added to each new position of the dots. As the amplitude increases, the dots will start to drift. At a certain noise level the drift will be so large that the subject is unable to determine the sense of the flow. Figure 3a–d show stimuli with noise levels of 0, 0.05, 0.1, and 0.2 deg respectively (for a stimulus diameter of 20 deg of visual angle). As in Fig. 2 the 16 frames are superimposed. In all panels of Fig. 3, the value of the curl is 1 rad/sec, the translational velocity is 10 deg/sec and there are 16 dots per frame with a lifetime of 16 frames.

We use a 2AFC-paradigm to determine noise thresholds using an adaptive staircase procedure. We start at zero noise level, and the first step size is 0.1 of a degree (see Fig. 3c). The subjects do 16 trials, each of which can either be, for instance, a clockwise or a counter-clockwise rotation. After these 16 trials they receive information about both their score (percentage correct) and the noise level. If their score is more than 10% worse than it was at the previous noise level the step size is halved and the noise level decreases by one step. Otherwise, the step size stays the same and the noise level increases by one step. Then the subjects again do 16 trials and the whole procedure is repeated. The procedure stops when either the step size becomes smaller than 20% of the total noise level, or if subjects reach a score below 60% and their score was not more than 10% worse than it was at the previous noise level. A 75% correct threshold is calculated in the following manner. The data points are ordered according to increasing noise levels. If our subjects were ideal, the frequency of correct responses would decrease from 100% correct to 50%. However, as subjects are rarely ideal, data might be very messy, and we need a procedure that is insensitive to bad outliers (e.g. the median instead of the mean). If the frequency of correct response is higher than 75% we denote the data point with a plus, if it is lower than 75% we denote the data point with a minus. Data points equal to 75% do not contribute. We then find the noise level at which there are just as much plusses on the high noise side as there are minuses on the low noise side. This noise level is taken to be the 75% correct point. This turns out to be a very robust procedure. We repeated measurements for occasional data points to check for any learning effects and there were no indications that thresholds changed over a large period of time (sometimes up to 1 yr).

We sampled both the range of the translational speed and the range of the flow component by doubling (or halving) their respective values until the subject was unable to perform the task even at zero noise-level. For instance for a rotation of 1 rad/sec we obtained thresholds for a translational speed of 5, 10, 20, 40 etc. deg/sec until at a certain value, for instance 80 deg/sec,

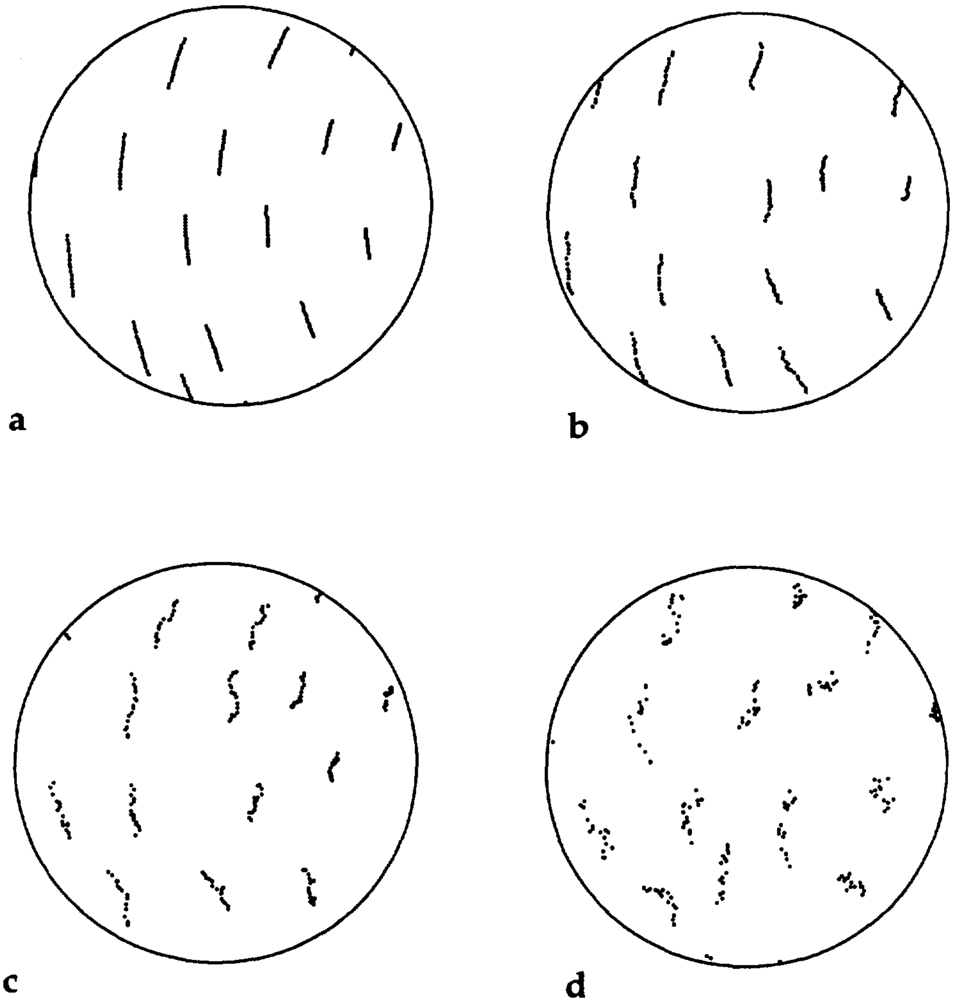

FIGURE 3. Stimuli with different noise levels. For all examples there were about 16 dots per frame and the life time of the dots was 16 frames. The curl is 1 rad/sec and the translational velocity is 10 deg/sec. (a) A noise level of 0 deg. (b) A noise level of 0.05 deg. (c) A noise level of 0.1 deg. (d) A noise level of 0.2 deg.

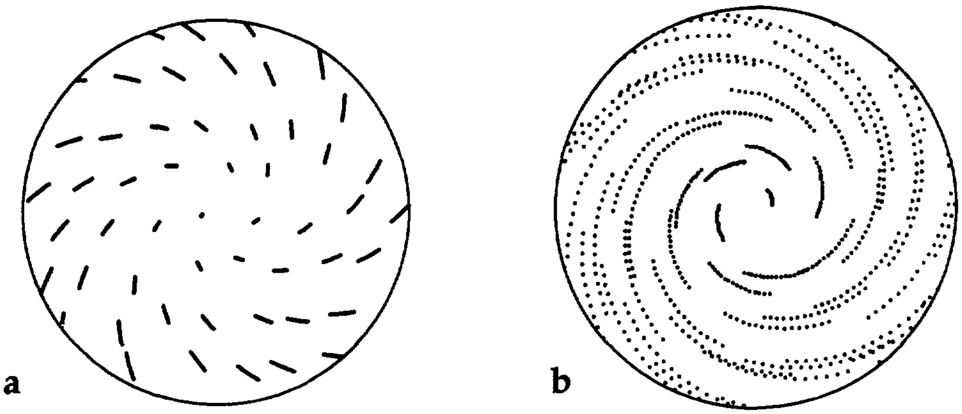

FIGURE 4. Combinations of rotation and expansion. Again the 16 consecutive frames are superimposed. The stimulus shows 64 dots per frame and the lifetime of the dots is 16 frames. The noise level is zero degrees. (a) Div = 2/sec, curl = 2 rad/sec. (b) Div = 2/sec, curl = 8 rad/sec.

the subject could not perform the task at zero noise-level. Then we doubled the value of rotation and started all over again. Thus our data points span the entire range of combinations of curl and translational speed that the subject could measure, provided the centre of rotation is kept outside the stimulus window. This procedure is employed for all four first-order optic flow components.

Because the literature tends to be inconsistent on this issue we explain here the standard meaning (in mathematics) of the phrases "the curl is 1 rad/sec" or "the def is 1/sec" etc. By *curl* we mean half the angular rotation per second. Thus a curl of 2 rad/sec means that an angle of one radian will be travelled in one second. By *divergence* we mean the relative area expansion (or contraction) per second. A divergence of 2/sec means that an area of size X will have expanded in one second by twice its own size to become $3X$. Perhaps the least intuitive quantity is the *deformation*. A deformation of 1/sec means a contraction of 1/sec in one direction and an expansion of 1/sec in the perpendicular direction. This results in a pure shape change, with no change in area and no rotation of the element. A direct advantage of this standard mathematical definition is that the four first-order components can easily be compared in terms of velocity distribution or deviation from parallel flow (as will be demonstrated in the Discussion section).

Second experiment

We performed a second experiment to check performance for different combinations of flow components. Instead of varying the translational speed inside the stimulus window we added a divergence to a curl. In this case the centre of rotation and divergence coincide with the centre of the stimulus; thus the average translational speed inside the stimulus window is always zero. Because the centre of rotation or divergence is always located in the centre of the stimulus, it will not provide a cue in this case. De Bruyn and Orban (1993) reported that in the case of a similar stimulus subjects found it easy to extract one component independently of the other. However, because this was a pilot experiment, these authors only looked at one combination of divergence and curl, whereas we measured the entire possible range. Once again the subject had to detect the sign of the flow by indicating whether the rotation was clockwise or counterclockwise. Thresholds were measured using the procedure described in the previous section. The sign of the divergence (either expansion or contraction) was chosen randomly. We will refer to this experiment as "curldiv".

We also ran the inverse experiment; in other words instead of asking the subject the sign of the rotation we presented the same stimulus, but asked for the sign of divergence (either an expansion or a contraction). This time, the sign of the curl was chosen randomly. We will refer to this experiment as "divcurl". These are essentially different tasks. In the first case the ability of subjects to extract the *curl* from a complex flow pattern (a spiral) is tested. In the second case we asked subjects to extract the *divergence* from a complex flow pattern (again a spiral). If subjects actually make use of a decomposition, two different mechanisms could be involved in these tasks.

Because all experiments with first order flow components in the presence of translational velocity indicate that dot density and correlation have very little effect on thresholds, we kept the number of dots at 16 dots per frame and the lifetime at 16 frames. Two examples of this stimulus are shown in Fig. 4, again the 16 frames are superimposed to create a static image. In Fig. 4a the divergence has a value of 2/sec and the curl has a value of 2 rad/sec. In Fig. 4b the divergence has a value of 2/sec and the curl has a value of 8 rad/sec. There are 64 dots per frame and the lifetime of the dots is 16 frames.

Subjects

Three subjects (AD and two of the authors, AK and SP) participated in a large number of the experiments; four others (AP, HV, IL and IV) participated only in some experiments in order to check whether the results were consistent for different subject. AK, SP, IL and IV are emmetropic. AD and HV are myopic corrected to normal. AP is myopic (−2 diopter). Table 1 shows the participation of subjects in different conditions. Numbers indicate how many different combinations of

TABLE 1. Participation of different subjects

	Subject						
Flow type	SP	AK	AD	AP	HV	IL	IV
Div	16	4		1	1	1	
Curl	4*	16	16				
Def	4	4					
Def45	4						
Curldiv	1						1*
Divcurl	1						1*

The numbers indicate how many different combinations of number of dots and liftime of dots where measured. In some cases we have not measured the entire range, these cases are marked by a '*'.

number of dots and lifetime of dots were measured. Thus the number four in the first row of subject AK means that she actually measured four sessions; an example of one such session is shown in Fig. 5. The stimulus was a divergence in the presence of translational velocity. She measured with four dots per frame and a lifetime of both two and sixteen frames, and with sixteen dots per frame and a lifetime of both two and sixteen frames, i.e. four different conditions. The experiments were conducted over a period of about 2 yr for AK and SP and over several months for AD. We repeated measurements for occasional data points, and the performance of our subjects did not change in any significant way, indicating that there are no learning effects. The results for the curl conditions for subjects AD and AK are taken from Kappers *et al.* (1994). Because the experiments are rather time consuming, we decided that in some cases measuring only a part of the range of div or curl would be enough to check whether results were consistent. These cases are marked with a "*" in Table 1.

RESULTS

Experiment 1

The results we will discuss in this section are qualitatively similar for all subjects over all conditions, although there were some quantitative differences in the height of the noise levels and the range that could be covered. However, we are not primarily interested in these detailed quantitative effects. Therefore we will not show every result we obtained. Instead we will present some representative graphs to illustrate our findings. The reader should bear in mind that the conclusions hold for all conditions and all subjects without exception unless noted explicitly.

Figure 5 shows data for subject AK for three different flow-types: a curl, a divergence and a deformation. In this case the stimulus consisted of 16 dots with a lifetime of 16 frames. For each value of the curl, def or div the 75% noise level is plotted as a function of the translational velocity inside the stimulus window. It should be noted that the points in each graph represent all conditions for which the subject was able to perform the task. The noise levels clearly depend on the curl, divergence and deformation in a similar way. Influence of translational velocity on noise levels appears to be very small over a large region (flat part of the curves), until at a certain critical value of the translational velocity performance drops steeply to zero. The point on the far left of each curve is determined by the experimental setup. It is the minimum translation necessary to locate the central of expansion, rotation or deformation outside the stimulus window. The point farthest to the right however is determined solely by the performance of the subject. It represents the maximum value of the translational velocity for which the subject was still able to perform the task. This may not seem obvious from the curve, but it should be realized that when the curve ends, the following point that we tested (which was at twice the speed) was really outside the range of the subject.

For all conditions we can characterize the performance of the subjects by two parameters. To represent the rather flat region we choose the maximum value of the curve. An alternative choice might be the average height

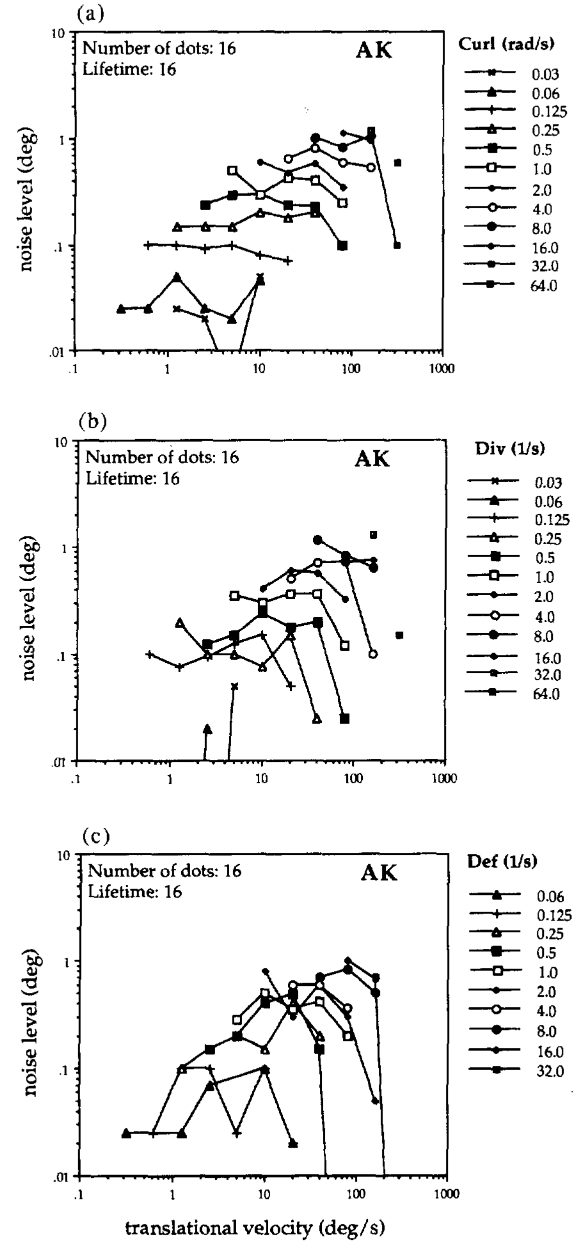

FIGURE 5. Data for subject AK for a stimulus with 16 dots per frame and a lifetime of 16 frames. (a) Results for curl. (b) Results for divergence. (c) Results for deformation (contraction axis is horizontal).

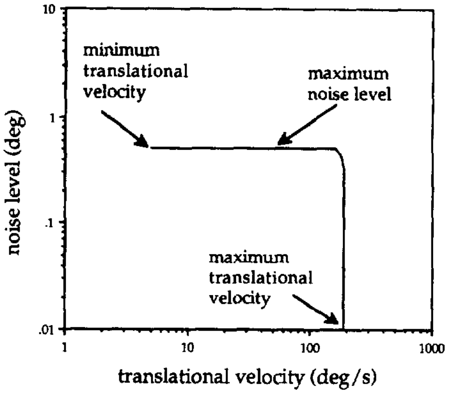

FIGURE 6. Representation of a typical curve for the 75% noise level as a function of the translational velocity. This curve can be characterized by a maximum noise level and a maximum translational velocity as indicated.

of the flat region, but it is difficult to define the exact end of the flat region. The other parameter is the maximum translational velocity, which we represented by the last data point in each curve. This representation is shown schematically in Fig. 6. Such a representation is of course a rough one, since the maximum noise level somewhat overestimates the height of the flat region of the curves. However, this does not create a problem because we are mainly interested in the overall trend, which remains unaffected. The advantage of this parametrization is the thorough condensation of the data, which permits a much clearer interpretation.

Figure 7a shows the maximum noise level (as defined in Fig. 6) as a function of divergence, curl and deformation for subject AK. Again, there were 16 dots per frame and the lifetime of the dots was 16 frames. Figures 7b and c show the maximum noise level as a function of divergence and deformation (in two directions) for subject SP. The number of dots per frame was 16 dots with a lifetime of 16 frames for Fig. 7b, and 4 dots with a lifetime of 2 frames for Fig. 7c. Figure 7a,b and c clearly show that the maximum noise level increases with curl, divergence and deformation. The most remarkable feature of Fig. 7 is that performance is independent of the flow-type that was actually used. Thus subjects appeared to be able to use some property of the flow field that was common in all four optic flow components, or alternatively subjects have detectors for different components that are tuned in the same way.

Figure 8 shows the maximum translational velocity as a function of divergence, curl and deformation. Conditions are the same as in Fig. 7; again results are given for two different subjects. Again we can observe a very similar performance for all four optic flow components, suggesting that a common factor is responsible for the results.

The only differences between subjects and conditions are of a quantitative nature. The main effect is that a different range can be covered. For instance, in Fig. 8a and b one can see that subject AK covers a range from 0.03/sec up to 64/sec for divergence, whereas subject SP covers a somewhat smaller range, from 0.125/sec up to 32/sec, for the same combination of lifetime and number of dots. When comparing Fig. 7b and c we can see the result of different lifetimes and dot densities. As the stimulus gets sparser there is a small effect on results: the subject is unable to cover such a wide range. The results of AK and SP are representative for the results of other subjects.

Experiment 2

Figure 9a and b show results for subject SP for the discrimination between expansion and contraction (in the presence of rotation) and the discrimination between

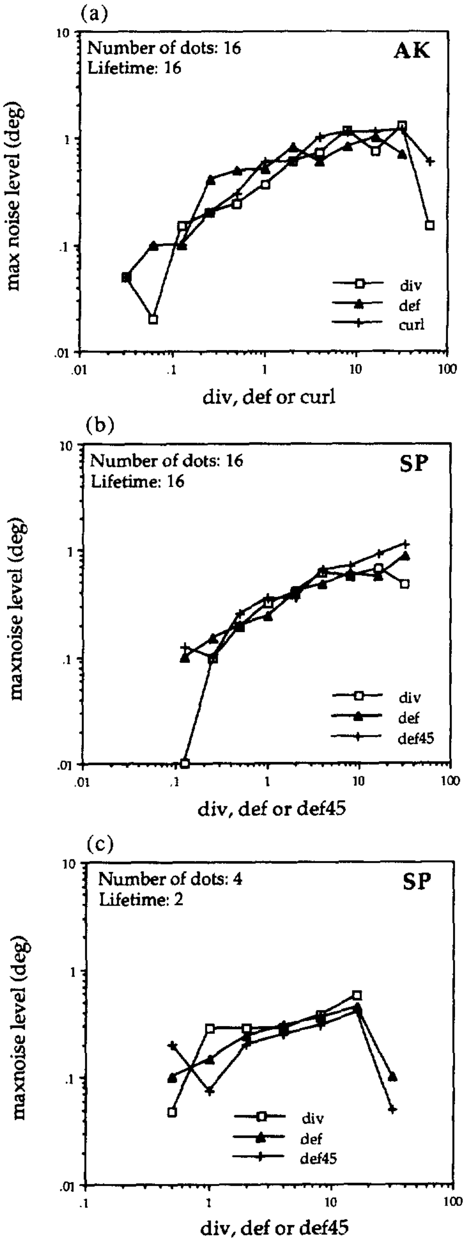

FIGURE 7. The maximum noise level as a function of the curl, divergence or deformation (in two directions). (a) Results for subject AK, 16 dots per frame and a lifetime of 16 frames. (b) Results for subject SP, 16 dots per frame and a lifetime of 16 frames. (c) Results for subject SP, 4 dots per frame and a lifetime of 2 frames.

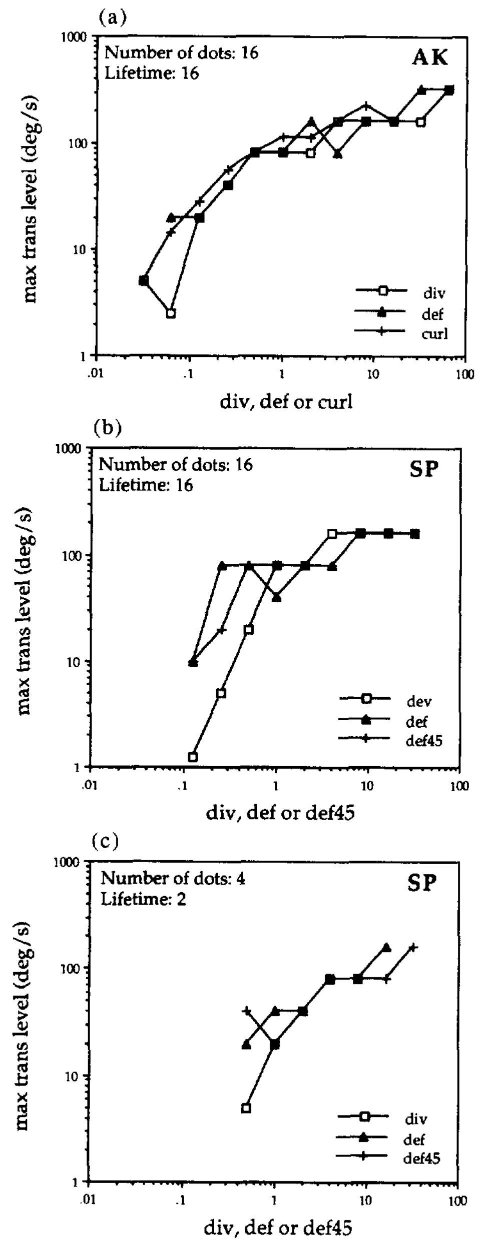

FIGURE 8. The maximum translational velocity as a function of the curl, divergence or deformation (in two directions) for the same conditions as in Fig. 7. (a) Subject AK, 16 dots per frame and a lifetime of 16 frames. (b) Subject SP, 16 dots per frame and a lifetime of 16 frames. (c) Subject SP, 4 dots per frame and a lifetime of 2 frames.

clockwise and counter-clockwise rotation (in the presence of divergence) respectively. Because each experiment takes so long (5–8 hr for a graph such as Fig. 9a) subject IV measured a more limited number of combinations. However, her results show the same trends as subject SP and we have no reason to believe this to be otherwise for the remaining conditions. Performance appears to be independent of the superimposed flow component over a wide range, then drops to zero when velocity gets too high. Again behaviour can be characterized by a maximum noise level and a maximum superimposed flow component (divergence or curl), as is schematically shown in Fig. 6. In this case the minimum value of the superimposed flow component is zero, since we do not need to locate the centre of rotation and divergence outside the stimulus window.

Figure 10a shows the maximum noise level as a function of curl (the divergence was varied). Performance for the discrimination between clockwise and counter-clockwise curl clearly increases when the value of the curl gets larger. Figure 10b shows the maximum noise level as a function of divergence, for the discrimination between expansion and contraction in the presence of curl as well as in the presence of translational velocity (as shown in Fig. 7). Both show remarkably similar curves. Detection of the sign of divergence obviously depends on both translational velocity and curl in a similar way. When we compare Fig. 10a and b, it is clear that detection of the sign curl and divergence yield comparable thresholds, which points to a similar detection mechanism. Figure 10c shows the maximum superimposed component as a function of divergence and curl. Again, results for the detection of the sign of div and curl show no differences. Here the limiting factor seems to be the overall speed rather than the type of the flow component. Results for our second experiment are very similar to those for our first experiment, suggesting that one optic flow component can be detected independently of other components, not just of the translational velocity.

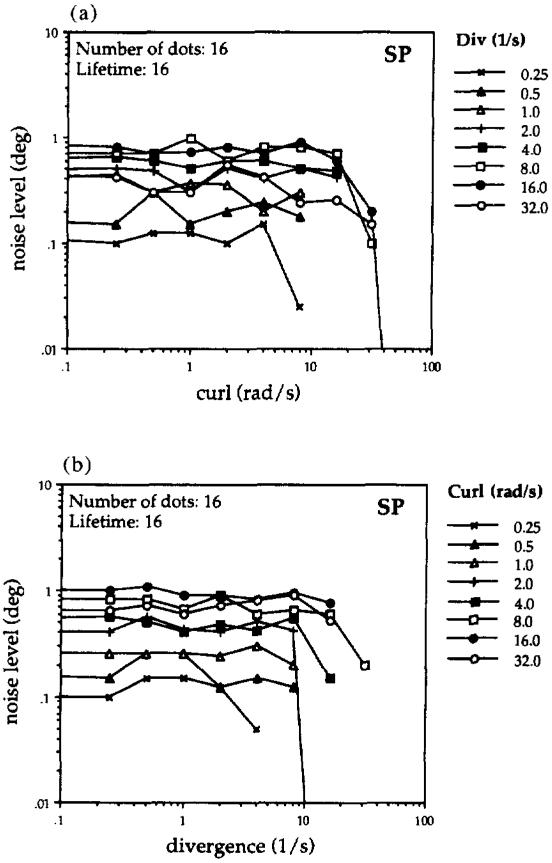

FIGURE 9. Data of subject SP for a stimulus with 16 dots per frame and a lifetime of 16 frames. (a) Results for detection of rotation in presence of divergence. (b) Results for detection of divergence in the presence of rotation.

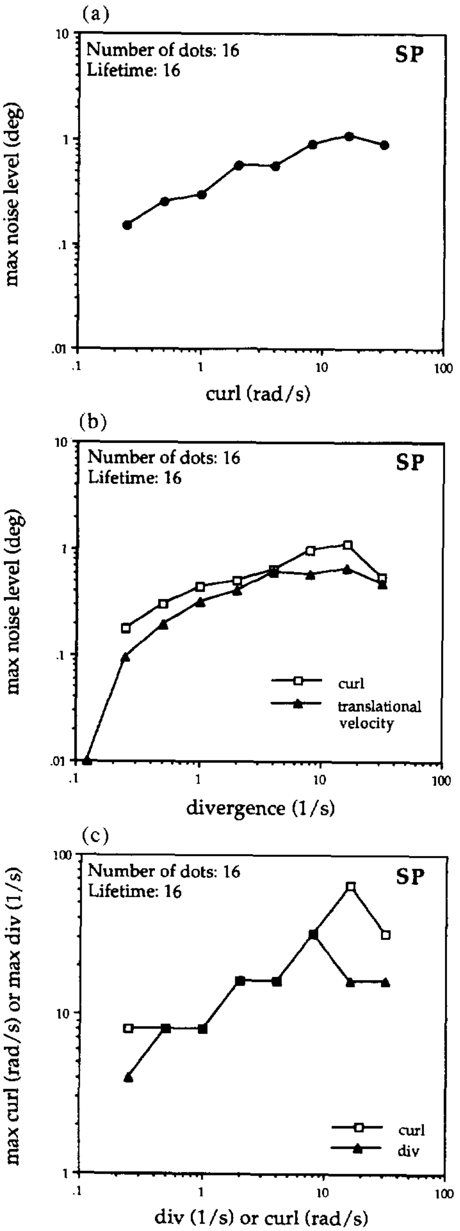

FIGURE 10. Results for subject SP, 16 dots per frame with a lifetime of 16 frames. (a) Maximum noise level as a function of curl. (b) Maximum noise level as a function of divergence. Results are plotted for both detection of divergence in the presence of rotation and for detection of divergence in the presence of translational velocity. (c) Maximum divergence as a function of curl and maximum curl as a function of divergence. The characterization is an analogue of what is shown in Fig. 6.

Control experiments

We ran various control experiments which we will mention briefly. First we changed the direction of the translational velocity to horizontal. When the direction of motion is horizontal instead of vertical results remain the same, indicating that there is no anisotropy. Because the direction of the translational velocity does not change during an experimental session, one could argue that subjects using that knowledge need only obtain the direction of one velocity vector (the movement of one dot between two frames) to perform the task. This could be an explanation for the fact that we found no difference in subjects' performance between lifetimes of 16 and 2 frames. When we choose a random direction of the translational velocity between 0 and 360 deg, however, results are not influenced at all. This indicates that subjects use at least two velocity vectors to obtain a result.

We also ran an experiment with static stimuli, i.e. we presented some of the div and curl stimuli as in our main experiment, but as in Fig. 2 the 16 consecutive frames were superimposed. The stimuli look much like Glass patterns (Glass & Pérez, 1973). Of course when we present just the static pattern, no conclusion can be drawn about the sign of the first-order flow field at all, the information is ambiguous. Therefore, we told the subjects beforehand that the direction of the movement in the equivalent dynamic task was always downward. Thus, we effectively reduced the task to a judgement of orientation. We introduced the same kind of noise to perturb the positions of the dots as we did in our main experiment and thresholds for this static task are similar to those of our dynamic experiment, although performance was slightly worse. Of course one cannot directly compare the results of the two experiments, because our dynamic stimulus presents no static cues (the individual frames consist of random dots, and we minimized persistance of vision by using dark dots on a light background). Moreover, in the static experiments the dot density was much larger than in the dynamic experiments, because we superimposed all 16 consecutive frames. Still, assuming that the static and the dynamic orientation judgement systems operate in a similar manner, the results lend some support to our findings that subjects use mainly the orientation of the velocity vectors to do the task.

In yet another experiment we changed the texture of the stimulus. This enabled us to investigate whether the nature of the carriers of the flow field had an influence on the results. To make comparisons easy we presented connected black and white triangles using the random dots from our main experiment as corners. We found only a small quantitative influence on results in that the maximum noise levels increased slightly. This was probably due to the amount of information available, which was much higher in the textured stimulus. Qualitatively, however, the results remained the same, indicating that the exact nature of the carriers of the flow field (e.g. random dots, triangles etc.) is of minor importance; we can therefore expand our previous conclusions to include textured stimuli.

DISCUSSION

In this paper we address the question of whether our visual system decomposes a first-order optical flow field into four independent components. While theoretical aspects of this decomposition have been looked at in detail (Koenderink & van Doorn, 1975, 1976; Koenderink, 1986; Longuet-Higgins & Pradzny, 1980) there have been very few systematical psychophysical studies.

We performed detection experiments for four different optic flow components (divergence, curl and two orthogonal components of deformation) in the presence of translational velocity. Similar performance is found for all four optic flow components over a maximally large range of experimental conditions. All subjects perform qualitatively the same, although there are some minor quantitative differences. Results strongly suggest that there is a common factor that the subjects detect, independent of the optic flow component presented. This conclusion is further supported by a second experiment where we added a divergence to a curl. Again we find no differences between curl and div detection, indicating that our findings hold for combinations of different optic flow components. More support is gained by a control experiment in which the texture was different. Again we find the same results as we found for the random dot stimuli. De Bruyn and Orban (1990) also found similar results for div, curl and def in an experiment where they disturbed the direction information of the individual motion vectors.

We analysed our stimulus with an ideal detector model. The detector was ideal in the sense that the correspondence problem was assumed to be solved. This model is based on a least squares estimation of the direction of the first-order flow component, using uncertainties in velocity estimation only. The ideal detector model shows that the maximum noise level should depend linearly on the first-order flow component if performance is optimal. For low velocities, human performance comes close to this prediction. For high velocities, however, human subjects are inferior to this ideal detector by a factor of about 10, or even higher. This means that in the high speed range subjects' performance is not limited by stimulus properties.

In earlier experiments we presented different aspects of a divergence flow field separately in order to find out which of these different aspects (e.g. direction of the velocity or velocity gradient etc.) subjects use for this kind of task (Kappers *et al.*, 1993). We found that performance depends on the direction of the velocity rather than on the velocity gradient or the exact value of the first-order component. This indicates strongly that the common factor our subjects detect is the *direction* of the velocity. Freeman and Harris (1992) also reported that expansion and rotation mechanisms are sensitive to directional relationships, while spatial speed gradients are not required. Physiological support for these findings is reported by Tanaka *et al.* (1989).

We found no evidence that different specific mechanisms are responsible for the detection of different flow components. Subjects respond to every flow type in a similar manner. The results suggest that subjects use a similar strategy for detection of first-order structure, regardless of the exact nature of the component. Subjects could easily detect one component in the presence of translational velocity, suggesting that divergence, curl and deformation are all independent of translational velocity. Moreover, our second experiment shows that divergence and curl are also independent of each other. This probably indicates that performance for all four first-order components is independent of all other first-order components.

A common factor in all four first-order flow fields is the deviation from parallel flow. This is a direct advantage of the standard mathematical definitions for the first-order optic flow components given in the section about the experimental procedure (last paragraph). Figure 11a and b show the displacement of one dot in between two frames for divergence and curl respectively. The deviation from parallel flow, denoted by θ, is the same (in first order approximation and disregarding the sign) in an expanding flow field with a divergence of 1/sec (Fig. 11a), as it is in a counter-clockwise rotating flow field with a curl of 1 rad/sec (Fig. 11b). For deformation we find the same angle θ. Detecting this direction difference is a task similar to discriminating the direction of the velocity vector. In earlier experiments we already found that subjects use mostly the direction of the velocity vector to perform the task (Kappers *et al.*, 1993). Thus, in our experiment subjects probably detected this common factor in the stimulus rather than a first-order optic flow component. While this might explain the similarities between results for different first-order flow components, it cannot explain why results are the same for different translational velocities. When one changes the translational velocity, the angle θ will of course change as well. This aspect is not explained by our common factor (the direction of the velocity).

Our experiments do not rule out the existence of different first-order optic flow component detectors. On the other hand they indicate that the underlying mechanisms of such detectors should at least be similar for different first-order optic flow components. There are a number of physiological studies that show that there are cells that are sensitive to rotation, expansion and even deformation in one case (Orban *et al.*, 1992; Tanaka & Saito, 1989; Tanaka *et al.*, 1989). However, up till now these authors have found only detectors that depend on the value of the other flow components. It is not yet clear what the functional role of these detectors is, but in our view the response from a real first-order optic flow component detector should be independent of the values of other components in the field. In our experiments the value of the curl, for instance, is constant over the entire stimulus area, and we feel that the response of a

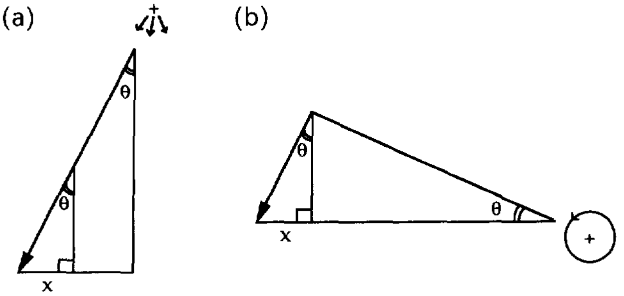

FIGURE 11. A similarity between the vector fields (movement between two frames) of (a) a divergence and (b) a curl. Notice that if the value of div and curl is the same, and both have the same translational velocity, then the angles θ are the same in both figures.

curl-detector should therefore be independent of the fixation point. However, one can easily see that if an optic-flow cell is simply integrating the response of local motion detectors, the same cell might respond differently for different fixation points, because the local velocity vectors are not the same. Another cell, tuned for a different location in the flow field, could take over, so that the total response is the same. Unfortunately, the neurophysiological studies can not make a distinction between one cell responding to one value of the curl, or a population of cells responding to one value of the curl, as they can test only single cells with different patterns, and not an entire population at the same time. Graziano *et al.* (1994) have found cells in MST that respond to spiral motion, indicating that if component detectors exist, they are not limited to curl, divergence and deformation only. The evidence from psychophysical studies is ambiguous: though there are psychophysical studies that suggest that, for instance, looming or rotation detectors do exist (Regan & Beverly, 1985), other experiments (Milne & Snowden, 1993) suggest that the visual field does not have different detectors for div, def and curl.

It is possible that specific mechanisms that are sensitive to different optic flow components all operate in a similar way. In that case we should expect similar thresholds for all optic flow components. This could be realised by a Relative Motion System, as discussed by Freeman and Harris (1992). However, experiments by Sekuler (1992) show that there is no separate channel for relative motion. Although the existence of specific component detectors which operate differently cannot be ruled out, subjects appear not to use such detectors over a large range of conditions.

REFERENCES

De Bruyn, B. & Orban, G. A. (1990). The role of direction information in the perception of geometric optic flow components. *Perception & Psychophysics*, *47*, 433–438.

De Bruyn, B. & Orban, G. A. (1993). Segregation of spatially superimposed optic flow components. *Journal of Experimental Psychology: Human Perception and Performance*, *19(5)*, 1014–1027.

Freeman, T. C. A. & Harris, M. G. (1992). Human sensitivity to expanding and rotating motion: effects of complementary masking and directional structure. *Vision Research*, *32(1)*, 81–87.

Gibson, J. J. (1950). *The perception of the visual world.* Boston: Houghton Mifflin.

Glass, L. & Pérez, R. (1973). Perception of random dot interference patterns. *Nature*, *246(5432)*, 360–362.

Graziano, M. S. A., Andersen, R. A. & Snowden, R. J. (1994). Tuning of MST neurons to spiral motions. *Journal of Neuroscience*, *14(1)*, 54–67.

Kappers, A. M. L., te Pas, S. F. & Koenderink, J. J. (1993). Detection of divergence in optic-flow fields. *Perception*, *22 (Suppl.)*, 83a.

Kappers, A. M. L., van Doorn, A. J. & Koenderink, J. J. (1994). Detection of vorticity in optical flow fields. *Journal of the Optical Society of America A*, *11(1)*, 48–54.

Koenderink, J. J. (1986). Optic flow. *Vision Research*, *26(1)*, 161–180.

Koenderink, J. J. & van Doorn, A. J. (1975). Invariant properties of the motion parallax field due to the movement of rigid bodies relative to an observer. *Optica Acta*, *22*, 773–791.

Koenderink, J. J. & van Doorn, A. J. (1976). Local structure of movement parallax of the plane. *Journal of the Optical Society of America*, *66(7)*, 717–723.

Lappin, J. S., Norman, J. F. & Mowafy, L. (1991). The detectability of geometric structure in rapidly changing optical patterns. *Perception*, *20*, 513–528.

Longuet-Higgins, H. C. & Pradzny K. (1980). The interpretation of a moving retinal image. *Proceedings of the Royal Society of London B*, *208*, 385–387.

Milne, A. B. & Snowden, R. J. (1993). Is there anything special about expansion and rotational flow fields? *Perception*, *22 (Suppl.)*, 95a.

Orban, G. A., Lagae, L., Verri, A., Raiguel, S., Xiao, D., Maes, H. & Torre, V. (1992). First-order analysis of optical flow in monkey brain. *Proceedings of the National Academy of Sciences U.S.A.*, *89*, 2595–2599.

Regan, D. & Beverly, K. I. (1985). Visual responses to vorticity and the neural analysis of optic flow. *Journal of the Optical Society of America A*, *2(2)*, 280–283.

Sekuler, A. B. (1992). Simple-pooling of unidirectional motion predicts speed discrimination for looming stimuli. *Vision Research*, *32 (12)*, 2277–2288.

Tanaka, K. & Saito, H.-A. (1989). Analysis of motion of the visual field by direction, expansion/contraction, and rotation cells clustered in the dorsal part of the medial superior temporal area of the macaque monkey. *Journal of Neurophysiology*, *62–3*, 626–641.

Tanaka, K., Fukada, Y. & Saito, H.-A. (1989). Underlying mechanisms of the response specificity of expansion/contraction, and rotation cells in the dorsal part of the medial superior temporal area of the macaque monkey. *Journal of Neurophysiology*, *62–3*, 642–656.

Warren, W. H., Blackwell, A. W., Kurtz, K. J., Hatsopoulos, N. G. & Kalish, M. L. (1991). On the sufficiency of the velocity field for perception of heading. *Biological Cybernetics*, *65*, 311–320.

Acknowledgements—This research is supported by the Insight II project of the ESPRIT basic research action of the European Commission and the Netherlands Organization for Scientific Research (NWO). We would like to thank all subjects for their valuable contribution.