PERGAMON

A model using MT-like motion-opponent operators explains an illusory transformation in the optic flow field

Constance S. Royden *, Daniel M. Conti

Department of Mathematics and Computer Science, College of the Holy Cross, P.O. Box 116A, Worcester, MA 01610, USA

Abstract

Previous studies have shown that a physiologically based model using motion-opponent operators to compute heading performs accurately for simulated observer translations. Here we show how this model can explain an illusory shift in the perceived focus of expansion of a radial flow field that occurs when a field of laterally moving dots is superimposed on a field of radially moving dots. Furthermore, we can use the model to predict the perceptual shift of the focus of expansion for novel visual stimuli. These results support the hypothesis that this illusion results from motion subtraction during the processing of optic flow fields.
© 2003 Elsevier Ltd. All rights reserved.

Keywords: Optic flow; Heading; Motion opponent; Computational model; Middle temporal area (MT)

1. Introduction

In order to navigate through the environment, people must accurately judge their direction of motion, or "heading". Psychophysical studies show that people judge their heading well under a variety of conditions when approaching a stationary scene (Crowell & Banks, 1993; Rieger & Toet, 1985; van den Berg, 1992; Warren & Hannon, 1988, 1990). Much recent research has examined how cells in the visual cortex might compute these heading parameters given the known responses of these cells to motion stimuli. Several models have been developed that show how these cells may process motion information to compute the direction of observer translation for observers moving through a stationary scene (Beintema & van den Berg, 1998; Cutting, Springer, Braren, & Johnson, 1992; Hatsopoulos & Warren, 1991; Lappe & Rauschecker, 1993; Perrone, 1992; Perrone & Stone, 1994; Royden, 1997). All these models perform as well as people when tested with simulations of observer motion in a straight line, and most can compute heading in the presence of rotations generated by eye movements. Therefore, one cannot

determine which models best describe the human mechanisms based on these tests alone.

One can further test these models by determining whether they respond similarly to humans under conditions for which they were not specifically developed. In addition one can generate predictions of human performance based on the model's responses to novel stimuli, and then test these predictions with visual psychophysics. One revealing stimulus comes from an illusory transformation of the optic flow field that occurs when a plane of laterally moving dots is superimposed on a plane of dots moving in a radial pattern as diagrammed in Fig. 1. Ordinarily, for a single plane of dots moving in a radial pattern (Fig. 1a), people accurately perceive the center of the radial pattern, known as the focus of expansion (FOE), which coincides with their perceived direction of motion (Gibson, 1950). When shown the stimulus with overlapping lateral and radial fields (Fig. 1c), people see an illusory shift of the focus of expansion (FOE) in the direction of the lateral motion (Duffy & Wurtz, 1993). Here we examine the response to this stimulus of a physiologically based model of heading detection (Royden, 1997) that uses "motion-opponent" operators, i.e. operators with adjacent excitatory and inhibitory regions within their receptive fields, similar to cells found in the middle temporal visual area (MT). We show that this model also shows this shift in the computed heading direction when presented with

---

*Corresponding author. Tel.: +1-508-793-2472; fax: +1-508-793-3530.

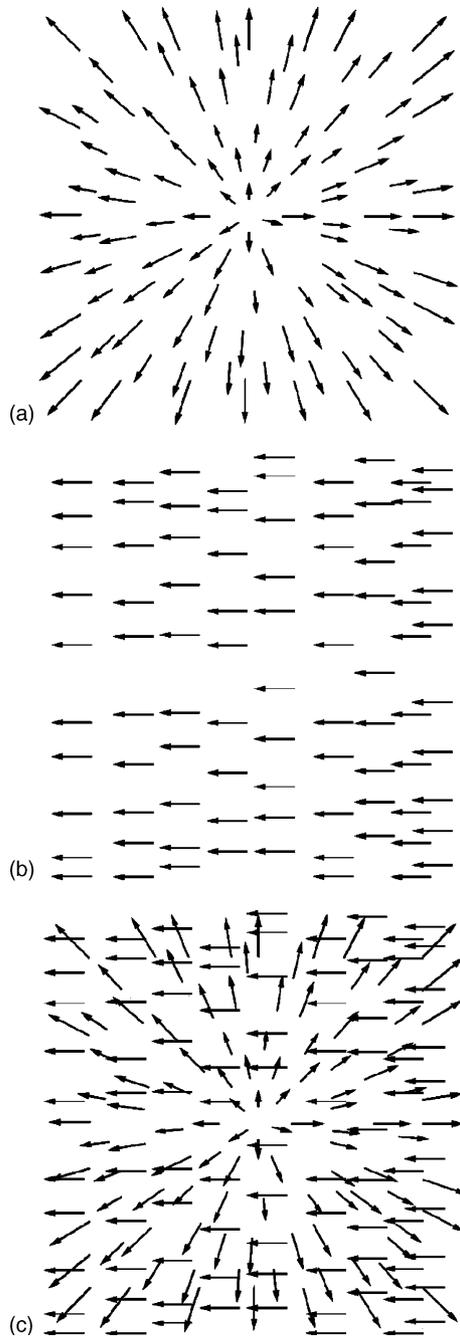*E-mail address:* croyden@mathcs.holycross.edu (C.S. Royden).

Fig. 1. Diagram of illusory stimulus. (a) Radial motion field. The arrows indicate the velocity and direction of individual points in the image. (b) Lateral motion field. All points move uniformly to the left. (c) Illusory stimulus with radial and lateral fields overlapping.

overlapping lateral and radial fields. We further show that the model predicts the magnitude and direction of other illusory shifts when presented with novel stimuli consisting of two superimposed radial flow fields. We present psychophysical results showing that people also experience the predicted illusory shifts when presented with these stimuli. These results support the hypothesis that the human visual system performs a motion sub-

traction in the process of computing heading from optic flow.

2. The model

The model tested here uses motion-opponent operators to compute heading for an observer undergoing translation and rotation. It has been described in detail elsewhere (Royden, 1997, 2002), so only the essential details will be given here. Based on the analysis of Longuet-Higgins and Prazdny (1980), this model uses subtraction of image motions in adjacent regions of the visual field to eliminate the rotational components of the image motion, leaving only the translational components. The remaining translational components can be used to compute the observer's heading.

Any observer motion can be described as a combination of translational and rotational motion along or around the three coordinate axes. Consider a point $P = (X, Y, Z)$ in the scene that is projected onto an image plane located one unit in front of the observer. The velocity of the projected point $p = (x, y)$ on the image plane is given by the following equations (Longuet-Higgins & Prazdny, 1980; Royden, 1997):

$$
\begin{aligned}
v_x &= \frac{xT_z - T_x}{Z} + xyR_x - (1 + x^2)R_y + yR_z \\
v_y &= \frac{yT_z - T_y}{Z} + (1 + y^2)R_x - xyR_y - xR_z
\end{aligned}
\tag{1}
$$

where T_x, T_y, and T_z are the three components of the observer's translational velocity and R_x, R_y, and R_z are the three components of the observer's rotational velocity. x and y are the coordinates of the point $P = (X, Y, Z)$ projected onto the image plane where $x = X/Z$ and $y = Y/Z$.

The velocity of the image point can be separated into two terms. The first term depends on the observer's translation, but not rotation. This term also depends on the distance, Z, of the point P from the observer. The second term depends only on observer rotation and is independent of the distance, Z. Longuet-Higgins and Prazdny (1980) suggested that if one can measure the image velocities for two points along a line of sight, for example at the border between two objects at different distances, then one can eliminate the rotation component by subtracting one of the image velocities from the other. The remaining difference vector depends only on observer translation and points directly toward or away from the observer's direction of translation. The difference vectors are given as

$$
\begin{aligned}
v_{xd} &= (-T_x + xT_z)\left(\frac{1}{Z_1} - \frac{1}{Z_2}\right) \\
v_{yd} &= (-T_y + yT_z)\left(\frac{1}{Z_1} - \frac{1}{Z_2}\right)
\end{aligned}
\tag{2}
$$

where v_{xd} is the horizontal component and v_{yd} is the vertical component of the difference vector and Z_1 and Z_2 are the distances from the two different surfaces. Rieger and Lawton (1985) showed that this approach works well even when the vector subtraction occurs for points that are spatially separated by a small amount. Hildreth (1992) extended the model to accommodate moving objects.

We used the idea of motion subtraction to develop a physiological model to compute observer translation direction in the presence of rotations (Royden, 1997). The motion subtraction is carried out by motion-opponent operators, shown in Fig. 2, that are based on

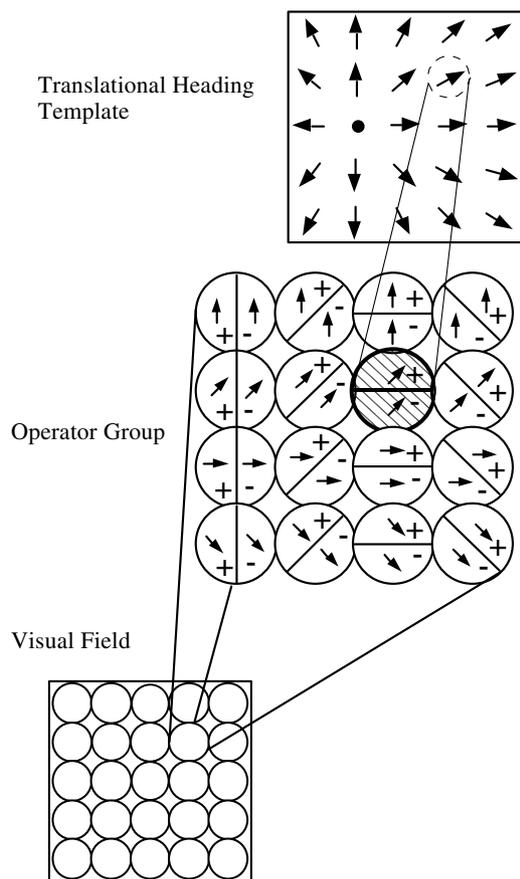Translational Heading Template

Operator Group

Visual Field

Fig. 2. Diagram of the model to compute heading. The bottom square shows the visual field divided into individual regions representing the receptive fields of the motion-opponent operators. In the model simulations, these receptive fields overlapped one another. The middle layer illustrates the group of motion-opponent operators that process the image motion from a single region of the visual field. These operators vary in their preferred direction of motion and in the angle of the line dividing the excitatory and inhibitory regions of the operator. Only a subset of the total number of operators is shown. An example of a hypothetical cell with maximal response is indicated with the diagonal hatch lines. The top layer illustrates a single template cell. The maximally responding operators in the middle layer project to a layer of cells whose receptive fields are templates tuned to radial patterns of input. The position of the center of this radial pattern, indicated by the black circle, varies for different template cells.

the receptive fields of cells in the primate visual area MT (Allman, Miezin, & McGuiness, 1985; Maunsell & van Essen, 1983a; Raiguel, Van Hulle, Xiao, Marcar, & Orban, 1995; Xiao, Raiguel, Marcar, Koenderink, & Orban, 1995). As shown in Fig. 2, each region of the visual field is processed by a group of these operators that differ in their preferred direction and the angle of the line dividing the excitatory and inhibitory regions. The precise spatial layout of the operators is not crucial to the results of the model. Previous experiments showed that using motion-opponent operators that have a center-surround spatial layout generates similar results to those generated by the version of the model shown here (Royden, 1997, 2002). In the current implementation, the direction tuning of each operator is given by a cosine function (i.e. the response to motion in the receptive field decreases with the cosine of the angle between the preferred motion direction and the direction of image motion within the receptive field). However, Royden (1997) showed that the response of the model does not depend critically on the tuning width.

Within a group of operators processing a given region of the visual field, the operator that responds most strongly to a given optic flow stimulus has a preferred direction of motion that points approximately toward or away from the point on the image plane that coincides with the intersection of the observer's direction of translation with the image plane. These maximally responding operators project to a second layer of cells that are templates for radial patterns of input from the motion-opponent operators. These template cells, which have some properties similar to cells in the medial superior temporal visual area (MST), vary in the location of the center of their preferred radial pattern (Duffy & Wurtz, 1991a, 1995; Graziano, Andersen, & Snowden, 1994; Saito et al., 1986; Tanaka & Saito, 1989). In other words, each template cell receives and sums the input only from those motion-opponent operators whose preferred directions of motion are consistent with the preferred radial pattern of that template cell. The template cell that responds most strongly will have a center that coincides with the observer's direction of translation. This model computes translational heading well in the presence of rotations (Royden, 1997) and shows heading biases similar to those exhibited by humans in the presence of moving objects (Royden, 2002).

Because the model is based on subtraction of image velocities in neighboring regions of the visual field, one can predict the response of the model to the stimulus that generates the illusory transformation by computing the difference vectors that result from subtracting the radial image vectors from the lateral image vectors and finding the point of intersection of these difference vectors. One can model the radial flow as the image motion for approach toward a frontoparallel plane at a distance Z from the observer. In this case, since there is no

rotation, the image velocity for a point on the plane is given by equations (1) with Rx, Ry and Rz all set to zero:

$$v_x = \frac{xT_z - T_x}{Z}$$
$$v_y = \frac{yT_z - T_y}{Z} \tag{3}$$

The image velocity for the laterally moving dots is a constant, v_{lat}, in the horizontal direction and zero in the vertical direction. If each radial velocity vector is matched with a lateral velocity vector, the difference vector generated by subtracting one from the other is given by

$$v_{xd} = \frac{(xT_z - T_x)}{Z} - v_{\text{lat}}$$
$$v_{yd} = \frac{(yT_z - T_y)}{Z} \tag{4}$$

If the radial flow field has a focus of expansion (FOE) in the center of the visual field, $T_x = 0$ and $T_y = 0$, so Eq. (4) reduces to

$$v_{xd} = \frac{xT_z}{Z} - v_{\text{lat}}$$
$$v_{yd} = \frac{yT_z}{Z} \tag{5}$$

Eq. (5) describes a radial field of difference vectors, with the focus of expansion located at

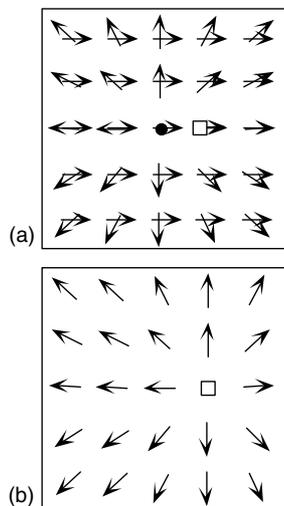$$x = \frac{v_{\text{lat}}Z}{T_z}$$
$$y = 0 \tag{6}$$

Fig. 3. Illustration of difference vectors for the illusory stimulus. (a) Overlapping radial and lateral image vectors in the illusory stimulus. Each arrow indicates an image velocity for a point in the image. (b) Difference vectors for the stimulus in (a). Each arrow indicates the difference vector generated by subtracting a lateral motion vector from a radial motion vector. The filled circle shows the center of the radial pattern of velocities in radial field. The open square shows the position of the center of the radial pattern of difference vectors.

Thus, the location of the center of the difference vector field is displaced horizontally away from the center of the radial field in the direction of lateral dot motion. This is diagrammed in Fig. 3.

3. Simulation 1: Response to the illusory stimulus

3.1. Methods

Eq. (6) implies that the motion-opponent operator model will also show a shift in the computed observer heading in the direction of the lateral dot motion when presented with a plane of laterally moving dots superimposed on a plane of radially moving dots similar to the illusion described by Duffy and Wurtz (1993), shown in Fig. 1. To test this prediction, we ran a full simulation of the model given this input stimulus.

In the following experiments we ran computer simulations of the Royden (1997) model on conditions similar to those used by Duffy and Wurtz (1993). The model parameters were the same as used in previous simulations (Royden, 1997, 2002). Each motion-opponent operator had a receptive field radius of 2 deg. Receptive field positions were spaced every 2 deg, so that they overlapped. Each region of the visual field was analyzed by 192 operators, representing 24 preferred directions of motion, evenly spaced between 0 and 360 deg, and eight angles of the axis between excitatory and inhibitory regions, evenly spaced between 0 and 180 deg. In order to increase the speed of computation, we used operators that could give a negative response to motion in the anti-preferred direction within the excitatory region of the receptive field and positive responses in the inhibitory region for anti-preferred motion, as described in Royden (1997). This allowed us to use a single operator to represent two neurons, e.g. one with an excitatory region on the right and inhibitory on the left and another with excitatory on the left and inhibitory on the right. Negative neural responses were ignored in the computation of heading. This is computationally equivalent to doubling the number of operators and allowing only positive responses in the excitatory region and negative responses in the inhibitory region.

The motion-opponent operators were distributed to cover a $40° \times 40°$ field of view. The heading templates had receptive field sizes that covered the entire viewing window. The input strength of each motion-opponent operator is weighted by a Gaussian function of the distance between the center of the motion-opponent operator's receptive field and the center of the preferred radial pattern of the template cell. The response of the template cell was computed as the sum of these weighted inputs from these motion-opponent operators. Preferred headings of the template cells were spaced every 2 deg both horizontally and vertically.

The conditions of this simulation were calculated to replicate the conditions used in the first experiment of Duffy and Wurtz (1993) as closely as possible. The visual field was made up of 300 dots whose initial positions were randomized within the field of view. Half of the dots in the field moved radially, expanding from the center of the viewing window. The other half of the dots moved laterally. In one set of simulations, the radial motion was calculated as the image velocities generated by the motion of an observer moving with a speed of 89.97 cm/s towards a stationary, transparent plane of dots located 50 cm from the observer. This resulted in a radial speed of 40 deg/s at a location 25 deg from the center of the field. The lateral motion was calculated as the motion of a transparent plane of dots moving horizontally across the observer's field of view with no change in depth relative to the observer. The speed of the lateral field was varied between 0, 9, 17 or 24 deg/s, both leftward and rightward. In a second set of simulations, the lateral speed was held constant at 17 deg/s and the observer speed was varied between 57.01, 72.32, 89.97, 111.04, or 137.24 cm/s to generate radial image speeds of 28, 34, 40, 46, or 52 deg/s measured at 25 deg from the FOE. Each condition was run 30 times.

3.2. Results

The results of the model's heading calculation are shown in Fig. 4. Each data point represents the average response of the model over a total of 30 trials. The model responded similarly to the human observers in the Duffy and Wurtz (1993) experiments. Leftward planar motion caused a shift in the computed FOE to the left and rightward caused a shift to the right. As shown in Fig. 4a, the magnitude of the shift increased with the speed of the lateral motion. The data in the Duffy and Wurtz (1993) paper showed an average linear regression line with a slope of $+0.74$ ($r = 0.86$). Linear regression of the model results, shown in Fig. 4, generates a line with a slope of $+0.56$ ($r = 1.00$). Fig. 4a also shows the amount of shift of the FOE predicted from calculating the focus of expansion of the difference vectors from Eq. (6). The model's results are very close to those calculated from the equations (slope $= +0.57$, $r = 1.00$).

Fig. 4b shows that the magnitude of the lateral shift decreases as the speed of the radial field increases. Duffy and Wurtz (1993) reported that this was true for human observers as well, although they did not report an average slope for this decrease.

4. Experiment 1. Creating a more robust illusion

The results presented by Duffy and Wurtz (1993) showed a great deal of variability among observers. When graphing the magnitude of the apparent shift due to overlapping planar motion, the slopes of regression lines for individual observers ranged between 0.3 and 1.2. We hypothesized that some of this variation could be due to spatial imbalances between the number of radially moving dots and the number of laterally moving dots in a given region of the visual field. This imbalance might lead to patches of pure lateral motion and patches of pure radial motion. It is possible that the visual system might interpolate between patches of lateral or radial motion to generate the perception of two
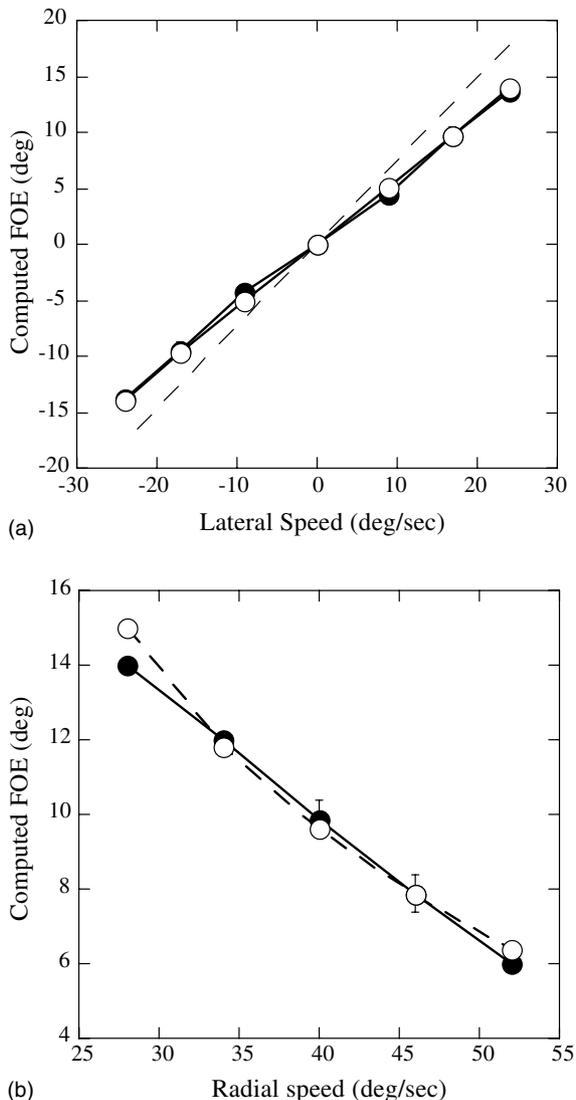

Fig. 4. Results of simulation 1: Graphs showing the computed FOE from model simulations. (a) The radial speed is held constant as the lateral speed is varied. (b) The lateral speed is constant as the radial speed is varied. Negative values indicate a position or motion to the left of center, while positive values indicate a position or motion to the right. Filled symbols show the response of the model, averaged over 30 trials. Open symbols show the predicted shift of the FOE based on the calculated difference vectors given in Eq. (6). The dotted line in (a) shows the average slope of reported shifts in Duffy and Wurtz (1993). Error bars for the model results indicate ±1 standard deviation. (Error bars not shown are smaller than the plot symbols.)

transparent planes, allowing trained observers to perceive the true FOE of the radial plane with higher accuracy. This idea is related to a similar explanation by Qian, Andersen, and Adelson (1994) of transparent motion perception for leftward and rightward moving planes. They showed that the perception of transparency could be eliminated by pairing each leftward moving dot with a rightward moving dot within a limited spatial region. We reasoned that we could eliminate transparency in our illusion in the same way. By pairing each radially moving dot with a laterally moving dot within a small spatial window, there would be no imbalance of motion signals. Therefore, if our hypothesis of local motion subtraction is correct, the resulting percept would be the result of the subtraction of one radial motion image velocity from one lateral motion image velocity at each location in the visual field. We hypothesized that this stimulus would generate a stronger illusion and human observers would show less variability between subjects.

4.1. Methods

This experiment used a computer controlled display of random dots to simulate observer motion through a scene. The scene consisted of two transparent planes composed of dots within a 25×25 deg viewing window. The radial motion was generated by simulating observer motion with a speed of 42 cm/s (note that this is a slower speed than used in simulation 1) toward the center of a stationary plane, located at 50 cm from the observer. The second plane consisted of dots that were assigned a uniform horizontal speed to generate the lateral motion. The dots were equally distributed between the two planes. The dot density for each trial was 0.64 dots/deg^2.

The stimuli were generated by a Power Mac G4 and presented on an Apple 21 inch CRT monitor. The display was set at 800×600 pixels with a refresh rate of 85 Hz. The dots were 2×2 pixel white squares subtending a 0.09 square degree area presented on a black background. The dots remained the same size over the entire sequence. Each dot had a lifetime of 240 ms. When a dot's lifetime expired it disappeared and was replaced by a new dot in a random place in the viewing window. At the beginning of each trial, each dot was randomly assigned an initial lifetime at random between 0 and 240 ms so the dots would expire at different times during a trial. After a dot was replaced the first time, its lifetime became 240 ms for all subsequent lifetimes in the trial. If a dot's motion brought it outside the viewing window, it was recreated at a random position inside the viewing window. Each sequence lasted 0.8 s and consisted of 24 frames.

Each observer viewed the screen with both eyes 50 cm from the screen, with their heads positioned using a chin and forehead rest. The observers were allowed free eye movements. The room was completely dark except for the light of the display. At the start of each trial, the first frame of the motion sequence was displayed, showing a static field of random dots. Observers initiated each trial by tapping the space bar, which set the dots in motion. At the conclusion of each trial, a cursor appeared. The observers were instructed to move the cursor to the position of the perceived FOE and click on the mouse button. The location of the cursor was recorded for each trial.

The scenes were constructed using two conditions: "matched point" and "non-matched point". The non-matched point condition randomly positioned every dot in both fields across the entire visual field, similar to the experiments of Duffy and Wurtz (1993). The matched point condition randomly distributed the dots of the radial field across the viewing window. Each dot from the lateral field was then placed in the same position as a dot from the radial field, creating a field with two dots at each position. The two dots would separate due to their respective motions during the trial. The lateral dots were assigned to move horizontally left or right at speeds of 0, 2, 6, or 10 deg/s. All seven speeds were tested using both matched point and non-matched point conditions for comparison between the methods. The conditions were presented in random order, with each different lateral speed being presented 10 times, using both matched point and non-matched point, for a total of 140 trials. The simulated speed of the observer, and therefore of the radial dot motion, was kept constant for each trial.

Eight observers with normal vision participated in Experiment 1. Two of the observers were aware of the experimental hypotheses and had previous experience as psychophysical observers. Two were unaware of the experimental hypotheses but had experience as psychophysical observers. The remaining four observers were all naive and had no prior experience. All observers volunteered to take part in the experiment and were not compensated for their participation. The naive observers participated in a practice session to familiarize them with the experiment. The practice session consisted of 14 trials similar to those in the actual experiment and was done with an investigator in the room to answer questions.

4.2. Results

The results of this experiment are shown in the Fig. 5, which shows the results for both the non-matched point (Fig. 5a) and the matched point conditions (Fig. 5b). For comparison, the response of the model and the calculated results from Eq. (6) are also included in each graph. For the matched point condition, the response of the model was calculated for the largest point separa-
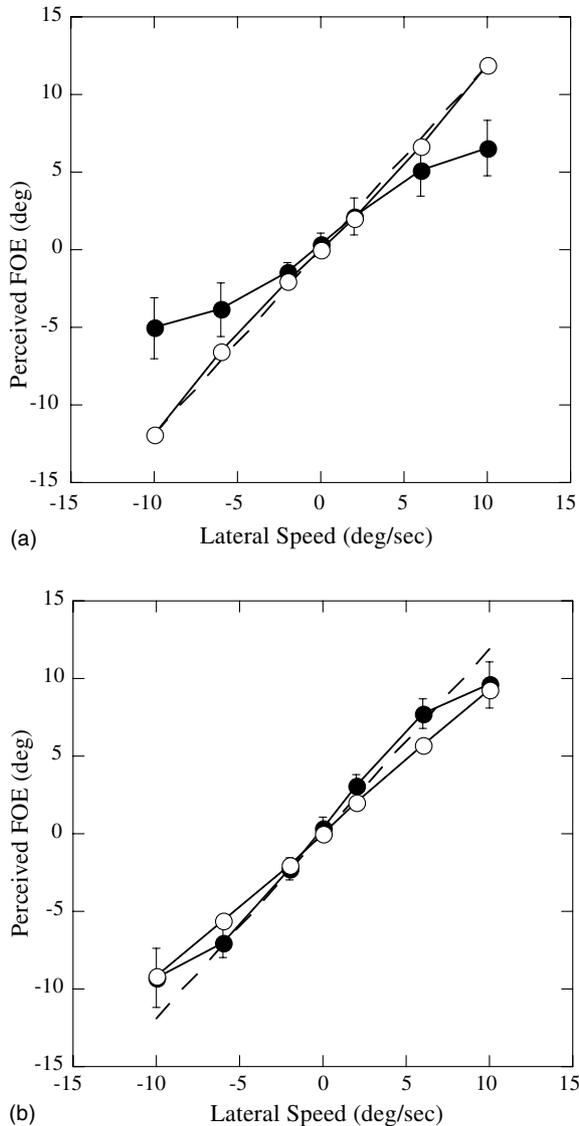
Fig. 5. Results of Experiment 1. (a) Perceived FOE position for different lateral speeds in the non-matched point condition. (b) Perceived FOE position for the matched point condition. Filled symbols show average response from human observers. Error bars are ±1 standard deviation. Open symbols show the response of the model averaged over 30 trials. The dashed line indicates the shift calculated based on difference vectors in Eq. (6).

tion, which occurs at $t = 240$ ms, because the model only performs subtractions between spatially separated points, due to the spatial extent of the receptive fields used. Fig. 5a shows the results of the non-matched point field responses. All of the subjects saw an illusory shift in the direction of lateral motion, however the amount of the perceived shift was considerably less than that shown by the model and the calculated shift from Eq. (6). Fig. 5b shows the results for the matched point condition. In this case subjects saw a much larger illusory shift than in the non-matched point condition, and much closer to the model results and the calculated shift.

The trend line for the matched point results has a slope of +1.03 ($r = 0.99$), nearly the same as the slope of the line showing the model results (slope +0.93, $r = 1.00$) and the calculated results (slope of +1.19, $r = 1.00$). The differences between the model and the human responses ranged from 0.1 to 2.0 deg, with an average difference of 0.8 deg. In every case the difference between the model response and the average human response was smaller than the largest differences seen between individual human subjects. Thus the model is responding as closely to the average human result as one might expect for a given human observer.

The magnitude of illusory shift varied considerably between subjects under the non-matched point scenario, particularly at the 6 and 10 deg/s lateral speeds with standard deviations of 1.74 and 2.00 deg respectively (for the left and right data combined). The results of the matched point scenario had less variation between subjects at these speeds with standard deviations of 1.01 and 1.67 deg for the 6 and 10 deg/s lateral speeds. Observers often commented that the matched point scheme created a stronger and less ambiguous illusion than the non-matched scheme. Some experienced observers became adept at seeing two transparent planes in the non-matched case, a perception that tended to nullify the perceptual shift. In contrast, these same observers said that they could not separate out two planes of motion in the matched point scheme. This is consistent with the idea of motion transparency diminishing the overall illusory shift of the FOE.

Because we allowed free eye movements, the retinal stimulus presented here may not exactly match the stimulus on the screen, since the observers may be making eye movements and tracking points in the stimulus. This should make no difference to the results of the model, since the motion subtraction eliminates the component of image motion resulting from an eye rotation. To verify that the model is unaffected by rotations due to eye movements, we repeated simulation 1 using the same stimuli with an added rotational component of 5 deg/s either about the X-axis, the Y-axis, the Z-axis or both the X and Y axes. The results were essentially the same as the results without the added rotations, with each data point differing from the corresponding data point in the other conditions by less than the standard deviation of the data (data not shown). This confirms our hypothesis that eye rotations do not affect the output of the model.

5. Experiment 2. Testing predictions for a novel stimulus

Experiment 1 and Simulation 1 show that a model using motion-opponent operators can account for an illusory shift in the perceived FOE when radial and lateral planes of dots overlap. One can further test the

idea that motion subtraction underlies this illusion by creating novel stimuli and using the model to predict the perceptual shift seen by humans. We reasoned that if motion subtraction accounts for the perceptual shift in the FOE seen when planes of radial and lateral motion overlap, then the same motion subtraction should lead to predictable perceptual shifts when two planes of radial motion with different FOE positions are superimposed, as diagrammed in Fig. 6. For this stimulus, one can calculate the predicted perceptual shift by calculating the radial center of the difference vectors generated when the velocity vectors from one of the radial planes are subtracted from the velocity vectors of the second plane, as described below.

One interesting aspect of this novel stimulus is that it can be used to test a competing theory of how the illusory shift is generated. Several researchers have proposed that the horizontal motion of the laterally moving dots stimulate a visual mechanism for detecting smooth pursuit eye movements (Duffy & Wurtz, 1993; Lappe & Rauschecker, 1995; Pack & Mingolla, 1998). The theory states that the system then compensates for these eye movements by shifting the perceived location of the center of expansion in the direction opposite the detected eye movement (which is in the same direction as the laterally moving dots). The radial patterns used here are unlikely to stimulate an eye movement system very strongly, since eye movements would tend to generate lateral flow, and thus one would expect that this stimulus would eliminate or greatly reduce the illusory effect if an eye movement compensation mechanism is the explanation. In contrast, the motion subtraction model should still lead to robust and predictable illusory shifts of the center of expansion.

In the following experiment we test both the model and human responses to two overlapping flow fields. These flow fields were generated by simulating observer motion toward two overlapping planes, as diagrammed in Fig. 7. The first flow field, referred to as the "radial field", was created by simulating observer motion toward a stationary plane, generating a radial pattern of image velocities with an FOE in the center of the viewing window. The second field is created by simulating observer motion toward a plane moving laterally across the observer's field of view at a given depth, parallel to the stationary plane. This flow field is referred to as the "lateral field". This simulated motion toward a laterally moving plane generates a radial pattern of image velocities whose FOE is shifted in the direction opposite the direction of the plane's lateral motion. Therefore the scene has two FOEs, one in the center corresponding to the stationary, radial flow field, and another shifted FOE corresponding to the lateral flow field.

Motion toward the stationary plane corresponds to translation only in the "Z" direction, letting T_x and T_y
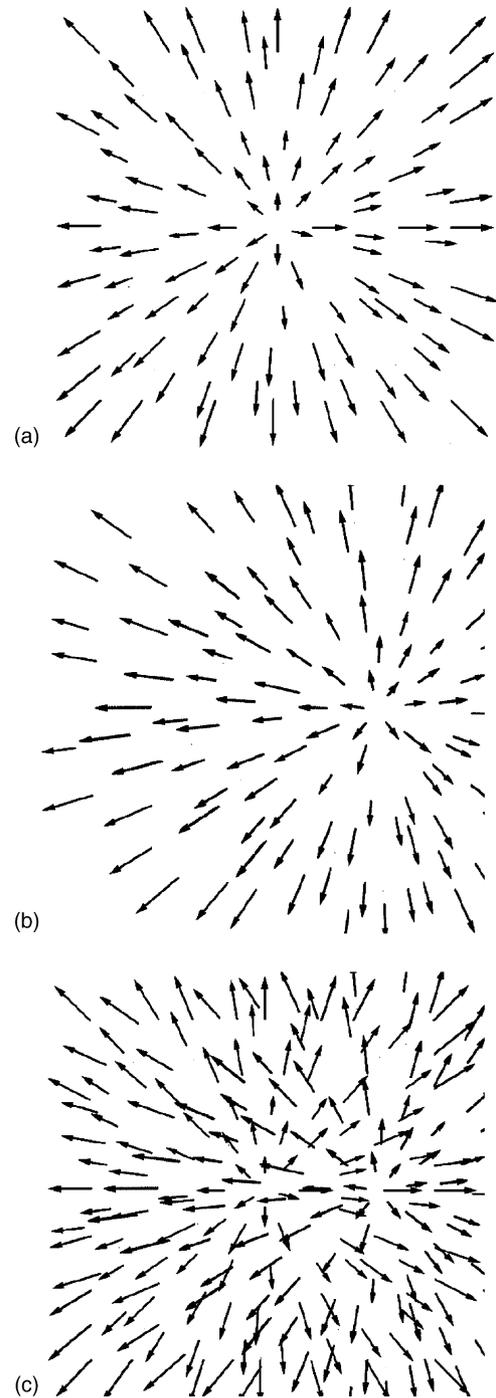

Fig. 6. Diagram of novel illusory stimulus used in Experiment 2. (a) Radial motion of dots in the "radial field" with FOE in the center of the field. (b) Radial motion of dots in the "lateral field" with FOE shifted to the right. (c) Combined stimulus with planes of motion from (a) and (b) overlapping.

equal zero. Let the distance to the radial field be denoted as Z_{rad}. From Eq. (3) we get

$$v_{x,\mathrm{rad}} = \frac{xT_z}{Z_{\mathrm{rad}}}$$
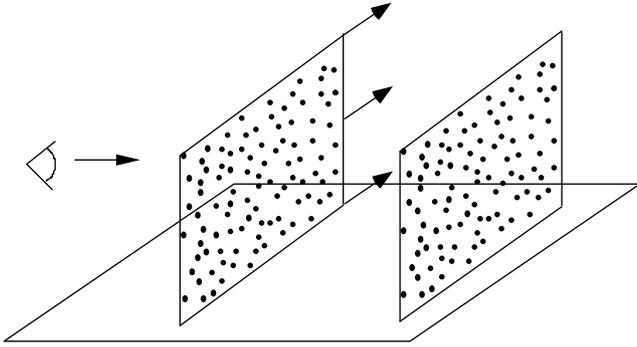$$v_{y,\mathrm{rad}} = \frac{yT_z}{Z_{\mathrm{rad}}}$$

$$(7)$$

Fig. 7. Diagram of simulated observer and planar motion that generates stimulus for Experiment 2. The observer moves straight toward the center of the stationary far plane generating image motion with FOE in the center of the field. The near plane moves laterally to the left or right as the observer approaches it, generating image velocities with a shifted FOE.

For the second plane, note that motion toward a plane moving in one direction gives the same flow field as would translational motion by the observer in the opposite direction. Therefore the flow field created by observer motion toward a laterally moving plane would be the same flow field associated with an observer having both forward and horizontal (T_z and T_x) translational motion. Let T_{lat} be the speed of the lateral moving plane and Z_{lat} be the initial distance between the lateral field and the observer. The flow field generated with these parameters would be the same as that for an observer moving forward with speed T_z and horizontally with speed $T_x = -T_{lat}$. The flow field for the lateral field is then

$$v_{x,lat} = \frac{xT_z - (-T_{lat})}{Z_{lat}}$$
$$v_{y,lat} = \frac{yT_z}{Z_{lat}} \tag{8}$$

The difference vector field is computed by taking the difference between the two vector fields defined by Eqs. (7) and (8)

$$v_{xd} = v_{x,lat} - v_{x,rad} = \frac{xT_z + T_{lat}}{Z_{lat}} - \frac{xT_z}{Z_{rad}}$$
$$v_{yd} = v_{y,lat} - v_{y,rad} = \frac{yT_z}{Z_{lat}} - \frac{yT_z}{Z_{rad}} \tag{9}$$

The center of the radial pattern of difference vectors defined in Eq. (9) is the location (x, y) where v_{xd} and v_{yd} are both zero. v_{yd} is clearly zero when $y = 0$. Setting v_{xd} to zero and solving Eq. (9) for x gives

$$x = \frac{-T_{lat}}{T_z} \left(\frac{Z_{rad}}{Z_{rad} - Z_{lat}} \right) \tag{10}$$

The positions of the FOEs of the difference vector field and the lateral field can also be given in terms of visual angles θ and ϕ respectively, where

$$\tan \theta = x$$
$$\tan \phi = \frac{-T_{lat}}{T_z} \tag{11}$$

Substituting the Eq. (11) into (10) gives the following equation:

$$\tan \theta = \left(\frac{Z_{rad}}{Z_{rad} - Z_{lat}} \right) \tan \phi \tag{12}$$

The angle of the FOE of the combined field is proportional to the angle of the FOE of the lateral field. This proportionality depends on the relative depths of the two fields. Consider a constant of proportionality, β such that

$$Z_{lat} = \beta Z_{rad} \tag{13}$$

Substituting Eq. (13) into (12) gives a general equation for calculating the position of the center of the radial pattern for the difference vectors given the ratio β and the position of the FOE of the lateral field, ϕ

$$\tan \theta = \left(\frac{1}{1 - \beta} \right) \tan \phi \tag{14}$$

Eq. (14) shows that two factors affect the position of the FOE of the difference vector field. The proportionality between the depths of the two fields, β, and the visual angle, ϕ, of the lateral field's FOE, which corresponds to the speed of the lateral field. One should therefore be able to manipulate the perceived location of the FOE by varying these two quantities. In the following experiment we tested the response of the model and of human observers to this new stimulus to determine whether the model can predict the perceptual shift of the FOE experienced by humans.

5.1. Methods

Nine observers with normal vision participated in Experiment 2. Six of the nine observers had participated in Experiment 1 and had familiarity with the workings of the experiment. The remaining three observers were all naive and had no prior experience. All observers volunteered to take part in the experiment and were not compensated for their participation. The naive observers participated in a practice session to familiarize them with the experiment. The practice session consisted of 14 movies similar to those in the actual experiment and was done with an investigator in the room to answer questions.

Experiment 2 was divided into two sets. Set 1 tested the effect on the perceived focus of expansion of varying the speed of the laterally moving field. The simulated observer motion toward the scene was held constant at a speed of 42 cm/s. The distances to the lateral and radial field were kept constant at 50 and 100 cm, respectively. The speeds of the lateral field were chosen so the

Table 1
Distances to two planes simulated in set 2 of Experiment 2 and corresponding β values

Experiment 2, Set 2 Field Depths								
β	0.1	0.3	0.5	1.5	2.0	3.0	2.22	1.67	0.23
Z_{rad}	500	166.67	100	100	100	100	67.5	90	645
Z_{lat}	50	50	50	150	200	300	150	150	150

calculated FOEs would remain inside the viewing window. Lateral speeds that corresponded to an FOE at visual angles ϕ of ±0, 2.5, 5, and 7.5 deg were used. Under the above conditions, the lateral speed of the field was 0, 2.1, 4.2, and 5.3 deg/s, respectively. For each condition, we tested the response to stimuli using the matched point and non-matched point conditions as described in Experiment 1.

Set 2 examined the effects of different depth relationships between the two fields on the perceived FOE position. The FOE of the lateral field was held constant at $\phi = 5°$, and β was varied by varying the distances to the two fields. The distances to the fields were calculated in terms of β as described in Eq. (13). The values used are listed in Table 1. Only the matched point stimuli were used in set 2.

The trials of sets 1 and 2 were randomly interleaved in a single run of the experiment. Each trial was shown fives times for a total of 115 trials (14 conditions for set 1 and 9 conditions for set 2). The conditions were shown in random order.

The model simulations were run using the same model parameters described in simulation 1. For these simulations we used the velocity vectors associated with the stimuli shown to human observers as input to the model. Each condition was run 30 times with different randomized dot positions within the two planes.

5.2. Results

Fig. 8 shows the results from set 1. Viewing both non-matched (Fig. 8a) and matched points (Fig. 8b), the human subjects and the computer model showed a shift in the direction opposite that of the lateral motion component of the lateral field. While this shift is predicted by the model, it is in the opposite direction from the shifts seen in Experiment 1. Fig. 8a shows the results for the non-matched points and Fig. 8b shows the results for the matched point condition. The magnitude of the illusory shift indicated by the human observers was greater for the matched point fields than for the non-matched points. The trendline for the non-matched points had a slope of -1.57 ($r = 1.00$) while the slope for the matched point condition was -1.99 ($r = 1.00$). Overall, the average human response for the matched point condition was close to the model's response (slope $= -2.28$, $r = 1.00$) as well as the calculated response (slope $= -2.68$, $r = 1.00$). As with the previous
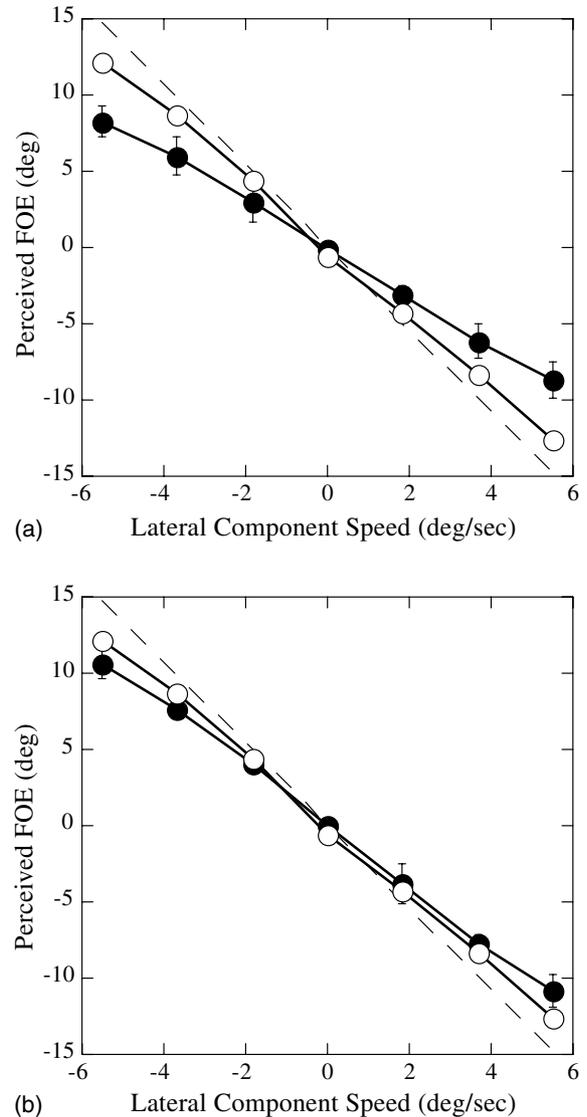

Fig. 8. Results of the first set of conditions in Experiment 2. (a) Position of the perceived FOE as the speed of the lateral motion component of the "lateral field" is varied for the non-matched point condition. (b) Position of the perceived FOE for the matched point condition. Negative values indicate a direction to the left of the center of the stimulus. Filled symbols indicate the average response of human observers. Error bars indicated ±1 standard deviation. Open symbols indicate the model response averaged over 30 trials. The dashed line indicates the calculated position of center of the difference vector field calculated from Eq. (6).

experiment, the illusory shift was smaller and more variable for the non-matched point condition than for the matched point condition. The differences between

the model and human results for the matched point condition range between 0.4 and 1.8 deg, with an average difference of 0.9 deg. In every case the difference between the model response and the average human response was smaller than the largest differences seen between individual human subjects. Thus, the model performance is similar to what might be expected from an individual human observer.

The results from set 2 are shown in Fig. 9. Again, the human results are very close to the results of the model, which both closely follow the calculated results. The differences between the average human results and the model responses range from 0.02 to 1.1 deg. As with the previous experiments, this is less than the difference seen between individual human observers. Therefore the model is again responding within the range expected of individual human responses. The calculated results in this set are asymptotic around $\beta = 1.0$. As β approaches 1.0 the results begin to go towards infinity. This same tendency was seen in the human results. On both sides of the asymptote, shown by a vertical dotted line in Fig. 9, the magnitude of the responses from the model and the humans increase.

One interesting result from this set of conditions is that the direction of the perceived shift can be either left or right depending on the value of β, even though the FOE of the lateral flow field was held constant at 5 deg. Thus the lateral component of motion was always to the left, and yet for values of β less than 1.0 the shift is to the right. The fact that the human results closely match those predicted by the model gives strong support to the hypothesis that the visual system uses motion subtraction to process flow fields.

6. Discussion

In these experiments we have tested the hypothesis that an illusory shift of the perceived FOE seen when a radial flow field and a lateral flow field are superimposed is caused by a motion subtraction mechanism. Others have presented similar hypotheses, suggesting that the effect could be explained by induced motion, which in effect subtracts some of the planar motion from the radial motion (Meese, Smith, & Harris, 1994) or by models that use local differential motion (Lappe & Rauschecker, 1995). Pack and Mingolla (1998) suggested that this illusory transformation could be partly accounted for using a center-surround type motion-opponent operator, similar to the operators used here. However this is the first time that the hypothesis has been tested by running simulations with a full implementation of such a model that uses spatially extended receptive fields to compute the motion subtraction. We have shown that the model shows an illusory shift of the computed FOE in the same direction as that perceived by people. In addition, in previous psychophysical work, the data have suggested that the visual system only partially compensates for the lateral motion. That is, the perceived shift was smaller than that predicted by subtracting the lateral motion from the radial motion vectors. We have developed a stronger version of the illusion by pairing the lateral and radial points. With this new stimulus, the shift in the FOE perceived by people closely matches the shift shown by the model. The fact that the model results fit so well with human results is consistent with the hypothesis that the illusion results from motion subtraction within the visual system. That the model predicts the perceived shift of the FOE in novel stimuli strengthens the support for this idea that the visual system uses some kind of motion subtraction mechanism to process optic flow fields.

6.1. Other possible architectures

While the model tested here produces results very similar to those of humans, it is not the only possible neural architecture that could produce these results. The key is in the motion subtraction, rather than in the exact organization of the neural network. For example, the subtraction could be the result of cross inhibition between cells of similar preferred directions of motion, similar to the architecture suggested by Qian et al. (1994) in their studies of motion transparency. The main
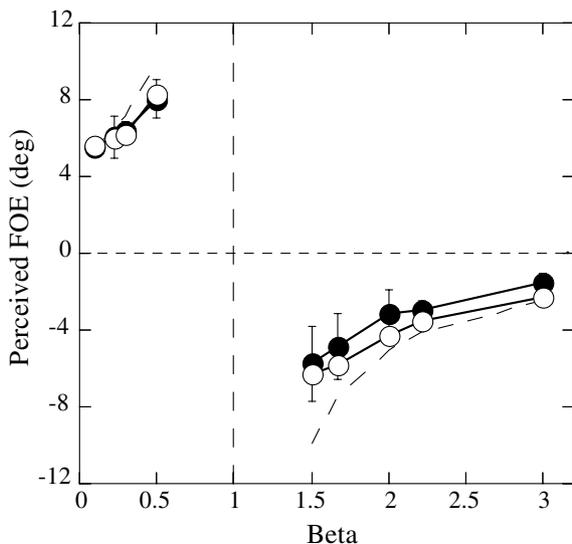

Fig. 9. Results of the second set of conditions in Experiment 2. The graph shows the location of the perceived FOE as the value of β, defined in Eq. (13), is varied. The location of the FOE of the lateral field is held constant at 5 deg to the right of the center of the visual stimulus. Filled symbols show the average response for human observers. Error bars indicate ±1 standard deviation. Open symbols indicate the response of the model averaged over 30 trials. The dashed line indicates the calculated center of the radial pattern of difference vectors, from Eq. (14).

difference between the architecture proposed here and that of Qian et al. (1994) is that in their architecture cells with opposite preferred directions of motion inhibited one another, while in the Royden model the inhibition occurs between adjacent regions with the same preferred direction of motion. The fact that our matched point experimental design gives the most powerful version of the illusion, while the non-matched point stimulus often leads to reported perceptions of transparent planes may be an indication of the relationship between the illusory stimulus and transparent motion perception. Qian et al. (1994) also found that pairing points in their stimulus eliminated the perception of transparency.

Alternatively, the subtraction could occur at the MST layer of processing, through pre-synaptic inhibition among the incoming nerve fibers from MT or through cross inhibition among the MST cells themselves. The model proposed by Lappe and Rauschecker (1993) contains this kind of inhibition by virtue of the weights of the connections between the first and second layers of their model. However, their model relies on a more global mechanism than ours, with the inputs to the second layer of cells coming from a widely distributed area of the visual field. Pack and Mingolla (1998) also argue for at least partial contribution of a global mechanism, integrating motion information from across the visual field to gain evidence for a rotation which is then subtracted from the flow field. The fact that our matched point stimulus generated a stronger and more reliable illusion than the unmatched point stimulus suggests that the subtraction occurs locally rather than globally. In other words, the subtraction could be accomplished by motion-opponent operators, as demonstrated here, or by local cross inhibition as described above.

If the MT cells are the primary units carrying out the motion subtraction, it is possible that these feed into the MST layer in some architecture different from the template pattern described above. There is some physiological evidence that suggests this is the case. Duffy and Wurtz (1991b) have found that responses of many MST cells cannot be explained as a simple mosaic of directional inputs from the previous layer. Furthermore, Lappe and Duffy (1999) have compared the responses of MST cells to the illusory stimulus described here with their responses to radial patterns of dot motion. They found that although the majority of MST cells exhibited a shift in their preferred center of expansion when presented with the illusory stimulus, in most neurons the shift was not as large as the perceptual shift shown in the psychophysical experiments. They showed that the distribution of responses is consistent with the behavior of cells in the Lappe and Rauschecker model (1993). One possible explanation for the varying amount of shift is that the random distribution of points used in the illusory stimulus for this experiment leads to unbalanced motion subtraction by the MT cells. This can lead to partial shifts in the computed focus of expansion. It would be interesting to test whether the shifts in the preferred center of expansion of MST cells were larger for the matched point condition we have developed here. We have not examined the responses of individual template cells in the model presented here, however one might expect them to show large shifts in their preferred center of expansion when presented with the illusory stimulus, consistent with only a minority of the cells found in MST. The exact behavior of the template cells in response to different stimuli is currently under investigation in our laboratory, and we hope the results will lead to modifications of the model to give a second layer of cells more similar to those found in MST.

We do not know the spatial extent over which the motion subtraction occurs. We were unable to find a visual stimulus in which the motions of the paired points were close enough spatially to eliminate the illusory effect. In the Royden model, if the points are very close together, so that their motions fall into the same half of the receptive field of an operator, then no subtraction occurs for the instantaneous image motion. This results in disappearance of the illusion. In fact, since the lateral and radial motions are averaged together in this case, the shift actually reverses direction. However, because the dots separate during the course of their lifetime, it may be that this separation is enough to allow the subtraction to occur in all real motion stimuli. Certainly our model, when run on a stimulus with the matched points at their maximum separation, still exhibits a shift in the computed FOE similar to that seen by humans. Whether the motion subtraction occurs between spatially separate, but adjacent, regions of the visual field as modeled here or whether it occurs between neurons with the same spatial locations of their receptive fields remains a question for future experimental investigation. However, the fact that the motion-opponent neurons in area MT are capable of carrying out this subtraction makes them good candidates for carrying out this process.

The operators used in this model are highly simplified versions of the neurons found in MT, intended to illustrate that a motion-opponent mechanism can account for the illusory effects described by Duffy and Wurtz (1993). We did not attempt to model the neurons in detail. However, we expect that the results will hold even for more detailed models of these neurons. We have shown in previous simulations with this model (Royden, 1997, 2002) that the precise spatial organization of the receptive fields does not affect the output of the model very much. For example, cells in area MT have a variety of spatial arrangements, including an asymmetric arrangement shown here as well as center-surround structures (Raiguel et al., 1995; Xiao et al., 1995). Previous work showed that, under a variety of conditions,

the center-surround structure shows behavior very similar to the behavior of the asymmetric operators shown here (Royden, 1997, 2002). The robustness of the output of the model to changes in receptive field structure suggests that this model will fare well when implemented with more realistic versions of MT neurons. In fact, a more detailed model of these neurons could explain a phenomenon that the current model does not. Pack and Mingolla (1998) reported that the illusion continued to strengthen as the lateral field was extended beyond the borders of the radial field. This could be explained by the large size of the inhibitory surrounds seen in MT (Allman et al., 1985; Raiguel et al., 1995; Xiao et al., 1995). The current model uses much more spatially limited surrounds, and so would not likely show this same effect. One future test of this model would be to simulate the more extended surrounds seen in MT neurons and test to see if they would generate this effect.

6.2. Other possible mechanisms

Ours is not the only model that can explain the illusion as described by Duffy and Wurtz (1993). Some other models (Beintema & van den Berg, 1998; Lappe & Rauschecker, 1993, 1995) also exhibit a shift in the computed translation direction when presented with these illusory stimuli. Interestingly, these other models also perform subtractions. The Lappe and Rauschecker model performs a subtraction at a higher level of processing than the MT cells, at the level of the input weights to the second layer of neurons. The Beintema and van den Berg model makes use of the difference between two templates tuned for opposite directions of rotation when computing heading direction. It seems likely that this subtraction leads to the shift seen with both these models. It would be an important test of these models to determine whether they respond in the same way that humans do to the novel stimuli developed here, using two overlapping planes of dots moving radially. If it is the case that motion subtraction is the main feature that allows all of these models to exhibit the shift of the center of expansion for the illusory stimulus, it may be difficult to develop experiments to distinguish among them. One possibility would be to address the local versus the global characteristics of the motion subtraction as described above. Our matched point conditions suggests that local subtraction mechanisms may be important to the generation of this illusion. This is the subject of ongoing research in our lab.

Several researchers have suggested that the illusory shift is due to a mechanism to compensate for an observer's eye movements (Duffy & Wurtz, 1993; Lappe & Rauschecker, 1995; Pack & Mingolla, 1998; Zemel & Sejnowski, 1998). The original argument proposed that the lateral motion of the dots stimulates a visual mechanism that detects rotations due to eye movements.

This rotation is then somehow subtracted from the flow field resulting in the perceptual shift seen. This idea relies on subtraction of motions, as does our model, which is why the two ideas both account for the illusion. However the results presented here suggest that the emphasis on eye movement compensation may be misplaced. The mechanism tested here will work independent of eye movements and is therefore more general. While our model can compensate for eye movements by motion subtraction, it can also compensate for rotations generated in other ways, such as an observer's motion on a curved path. In addition, the results obtained in Experiment 2, using the stimulus that consists of two radial flow fields, are incompatible with the eye movement compensation model. First, this stimulus has only radial flow fields, one of which has a shifted center of motion. A radial field would be unlikely to stimulate an eye movement system as strongly as laterally moving dots would, and therefore one would expect a weaker or non-existent illusion in this case. However, this stimulus generates a strong, reproducible illusory shift. It therefore seems likely that this shift is generated by a mechanism that does not depend on eye movement information. It is possible that the eye movement system picks up on the lateral component of flow in the radial field with the shifted center of expansion. This field can be decomposed into the sum of a radial field with a central focus of expansion and a lateral field. One might then argue that the eye movement system compensates for this lateral component of motion. However, this explanation cannot account for the conditions in which the illusory shift is in the opposite direction of this lateral component (when $\beta < 1.0$). The eye movement compensation model cannot account for the shifts in these conditions. Finally, the information about possible eye movements generated by the lateral motion of dots in the display is inconsistent with the actual eye movements being made by the observer. In other words, the lateral dot motion in the display is always in addition to the lateral motion generated by the observer's actual eye movements. If the system has access to extra-retinal information about the speed and direction of the actual eye movements, this extra-retinal information would be in conflict with the added visual information from the lateral motion field. It would therefore be unnecessary and inaccurate for the system to interpret the added lateral motion as being due to eye movements. Therefore, for all of the above reasons, it seems more likely that this illusory shift is the result of a more general motion subtraction that occurs throughout the optic flow field independent of the presence of absence of eye movements.

Finally, we have considered here only physiologically based models of heading computation, as we are primarily concerned with the physiological mechanisms underlying heading perception. These models are

necessarily constrained by the physiological properties of neurons. It is possible that other computational approaches would be more efficient or accurate in computing heading for mobile robots, which do not have all of the same constraints as do the biological systems.

6.3. Influence of depth ordering

Recently Grigo and Lappe (1998) examined the effect of stereoscopically presented stimuli on the strength of the illusion described here. They showed that the strength of the illusion decreased when the two patterns of motion (radial and lateral) are separated in depth using a stereoscopic depth cue. This reduction was less pronounced when the lateral motion was shown as behind the radial motion as opposed to the condition when the lateral motion was in front. They interpreted this finding as consistent with the idea that the illusion results from a mechanism that compensates for the eye movements. In the context of the model presented here, the decrease in the strength of the illusion when the planes are presented at different stereo disparities could be a result of the tuning for stereo disparity of individual MT cells (Maunsell & van Essen, 1983b). Although the current model does not take into account tuning for disparity, one can predict that dots at different disparities would be likely to stimulate different populations of MT cells, resulting in a perception of two transparent planes sliding across one another, similar to that seen by Qian and colleagues (Qian et al., 1994). The increased perception of two transparent planes would likely lead to a weaker illusory effect as discussed above. The asymmetry of the effect seen by Grigo and Lappe is intriguing, however speculation on whether it might arise from the disparity tuning of MT cells must await further electrophysiological data regarding the relationship of the disparity tuning for the center and surround of the motion-opponent subpopulation of these cells.

In our Experiment 1 there are no explicit cues to depth. The only monocular cue is the average dot speed within the two overlapping planes. A faster speed would be consistent with a closer plane. Thus, as the lateral dot speed increases, this would be consistent with the laterally moving plane being closer to the observer than the radially moving plane. If depth ordering were important in this case, the illusion should be weaker as lateral speed increases, bringing the lateral plane in front. Clearly this is not the case, so speed cues to depth do not have the same effect as stereo cues. In Experiment 2 we have suggested one physical situation that could lead to the stimulus described, however it is not unique and the relative distances of the planes could differ in other situations that could lead to the same image motions. However, if one examines the situation shown in Fig. 7, one can ask whether there is an asymmetry in the strength of the illusion that depends on the relative depths of the two planes, as described in the parameter, β. Although the direction of the shift changes depending on whether the laterally moving plane is in front ($\beta < 1$) or behind ($\beta > 1$) the radial plane, the magnitude of the shift is comparable in both situations and predicted by the motion subtraction mechanism described above. The shift is slightly larger when the laterally moving plane is in front of the radial plane than when it is behind it (compare results for $\beta = 0.5$ to those for $\beta = 2.0$). This is opposite the effect of stereoscopic disparity found by Grigo and Lappe. Thus the simplest explanation for the magnitudes of the shifts seen in this set of experiments is probably not related to the depth of the two planes per se, but rather to the relative dot speeds in the two planes which lead to the motion differences that predict these shifts.

6.4. Relationship to moving objects

We recently showed that the Royden (1997) model can account for biases seen in judgments of heading for scenes containing moving objects. A moving object in an otherwise stationary scene can cause biases in heading perception depending on the motion of the object (Royden & Hildreth, 1996; Warren & Saunders, 1995). An object moving laterally relative to the observer causes a bias in the direction of object motion, while certain looming objects cause a bias in the direction opposite the lateral component of motion, i.e. in the direction of the looming object's FOE. It has been previously suggested that the biases seen with moving objects are related to the illusory shift examined here (Pack & Mingolla, 1998; Warren & Saunders, 1995). The data presented here and for the moving object experiments (Royden, 2002) clarify this relationship. Both phenomena can be explained by a motion subtraction process. The bias seen with the lateral moving object is akin to the illusory shift seen in the original Duffy and Wurtz (1993) illusion. The bias seen with the looming objects is similar to that seen with the two radial planes of motion used here, when $\beta < 1.0$. The fact that the motion-opponent model for heading perception can so neatly account for observer biases in multiple conditions for which it was not originally developed provides strong support for this model as the mechanism for human optic flow processing.

7. Conclusion

We have demonstrated that a model using motion-opponent operators to compute heading can account for an illusory shift in the perceived FOE when a radial and a lateral field of dots are superimposed. We have shown that the illusion can be enhanced in humans by pairing each radially moving dot with a laterally moving

dot, suggesting that local interactions are important for the illusion. The human perception with this enhanced illusion matches the model results very closely. In addition, we have shown that we can use this model to predict the illusory shifts perceived for novel stimuli consisting of two overlapping radial fields of dot motion. This result supports the idea that optic flow is processed by a motion subtraction mechanism. While this motion subtraction could be carried out in a number of different ways, the fact that cells in visual area MT exist that have the motion-opponent properties required for this subtraction makes them excellent candidates for processing the optic flow field. Future work will examine models that use more detailed versions of these cells' receptive fields to gain further insight into the mechanisms of human heading perception.

Acknowledgements

References

Allman, J., Miezin, F., & McGuiness, E. (1985). Direction- and velocity-specific responses from beyond the classical receptive field in the middle temporal visual area (MT). *Perception, 14*, 105–126.

Beintema, J. A., & van den Berg, A. V. (1998). Heading detection using motion templates and eye velocity gain fields. *Vision Research, 38*, 2155–2179.

Crowell, J. A., & Banks, M. S. (1993). Perceiving heading with different retinal regions and types of optic flow. *Perception and Psychophysics, 53*, 325–337.

Cutting, J. E., Springer, K., Braren, P. A., & Johnson, S. H. (1992). Wayfinding on foot from information in retinal, not optical, flow. *Journal of Experimental Psychology: General, 121*, 41–72.

Duffy, C. J., & Wurtz, R. H. (1991a). Sensitivity of MST neurons to optic flow stimuli. I. A continuum of response selectivity to large field stimuli. *Journal of Neurophysiology, 65*, 1329–1345.

Duffy, C. J., & Wurtz, R. H. (1991b). Sensitivity of MST neurons to optic flow stimuli. II. Mechanisms of response selectivity revealed by small-field stimuli. *Journal of Neurophysiology, 65*, 1346–1359.

Duffy, C. J., & Wurtz, R. H. (1993). An illusory transformation of optic flow fields. *Vision Research, 33*, 1481–1490.

Duffy, C. J., & Wurtz, R. H. (1995). Response of Monkey MST neurons to optic flow stimuli with shifted centers of motion. *Journal of Neuroscience, 15*, 5192–5208.

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

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

Grigo, A., & Lappe, M. (1998). Interaction of stereo vision and optic flow processing revealed by an illusory stimulus. *Vision Research, 38*, 281–290.

Hatsopoulos, N. G., & Warren, W. H. (1991). Visual navigation with a neural network. *Neural Networks, 4*, 303–317.

Hildreth, E. C. (1992). Recovering heading for visually-guided navigation. *Vision Research, 32*, 1177–1192.

Lappe, M., & Duffy, C. J. (1999). Optic flow illusion and single neuron behaviour reconciled by a population model. *European Journal of Neuroscience, 11*, 2323–2331.

Lappe, M., & Rauschecker, J. P. (1993). A neural network for the processing of optic flow from ego-motion in man and higher mammals. *Neural Computation, 5*, 374–391.

Lappe, M., & Rauschecker, J. P. (1995). An illusory transformation in a model of optic flow processing. *Vision Research, 35*, 1619–1631.

Longuet-Higgins, H. C., & Prazdny, K. (1980). The interpretation of a moving retinal image. *Proceedings of the Royal Society of London B, 208*, 385–397.

Maunsell, J. H. R., & van Essen, D. C. (1983a). Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed and orientation. *Journal of Neurophysiology, 49*, 1127–1147.

Maunsell, J. H. R., & van Essen, D. C. (1983b). Functional properties of neurons in middle temporal visual area of the macaque monkey. II. Binocular interactions and sensitivity to binocular disparity. *Journal of Neurophysiology, 49*, 1148–1167.

Meese, T. S., Smith, V., & Harris, M. G. (1994). Induced motion may account for the illusory transformation of optic flow fields found by Duffy and Wurtz. *Vision Research, 35*, 981–984.

Pack, C., & Mingolla, E. (1998). Global induced motion and visual stability in an optic flow illusion. *Vision Research, 38*, 3083–3093.

Perrone, J. A. (1992). Model for the computation of self-motion in biological systems. *Journal of the Optical Society of America A, 9*, 177–194.

Perrone, J. A., & Stone, L. S. (1994). A model of self-motion estimation within primate extrastriate visual cortex. *Vision Research, 34*, 2917–2938.

Qian, N., Andersen, R. A., & Adelson, E. H. (1994). Transparent motion perception as detection of unbalanced motion signals. I. Psychophysics. *Journal of Neuroscience, 14*, 7357–7366.

Raiguel, S., Van Hulle, M. M., Xiao, D. K., Marcar, V. L., & Orban, G. A. (1995). Shape and spatial distribution of receptive fields and antagonistic motion surrounds in the middle temporal area (V5) of the macaque. *European Journal of Neuroscience, 7*, 2064–2082.

Rieger, J. H., & Lawton, D. T. (1985). Processing differential image motion. *Journal of the Optical Society of America A, 2*, 354–360.

Rieger, J. H., & Toet, L. (1985). Human visual navigation in the presence of 3D rotations. *Biological Cybernetics, 52*, 377–381.

Royden, C. S. (1997). Mathematical analysis of motion-opponent mechanisms used in the determination of heading and depth. *Journal of the Optical Society of America A, 14*, 2128–2143.

Royden, C. S. (2002). Computing heading in the presence of moving objects: A model that uses motion-opponent operators. *Vision Research, 42*, 3043–3058.

Royden, C. S., & Hildreth, E. C. (1996). Human heading judgments in the presence of moving objects. *Perception and Psychophysics, 58*, 836–856.

Saito, H., Yukie, M., Tanaka, K., Hikosaka, K., Fukada, Y., & Iwai, E. (1986). Integration of direction signals of image motion in the superior temporal sulcus of the macaque monkey. *Journal of Neuroscience, 6*, 145–157.

Tanaka, K., & Saito, H. (1989). Analysis of motion in 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*, 626–641.

van den Berg, A. V. (1992). Robustness of perception of heading from optic flow. *Vision Research, 32*, 1285–1296.

Warren, W. H., & Hannon, D. J. (1988). Direction of self-motion is perceived from optical flow. *Nature, 336*, 162–163.

Warren, W. H., & Hannon, D. J. (1990). Eye movements and optical flow. *Journal of the Optical Society of America A, 7*, 160–169.

Warren, W. H., & Saunders, J. A. (1995). Perceiving heading in the presence of moving objects. *Perception, 24*, 315–331.

Xiao, D. K., Raiguel, S., Marcar, V., Koenderink, J., & Orban, G. A. (1995). Spatial heterogeneity of inhibitory surrounds in the middle temporal visual area. *Proceedings of the National Academy of Science USA, 92*, 11303–11306.

Zemel, R. S., & Sejnowski, T. J. (1998). A model for encoding multiple object motions and self-motion in area MST of primate visual cortex. *Journal of Neuroscience, 18*, 531–547.