Available online at www.sciencedirect.com
SCIENCE DIRECT
Vision Research 46 (2006) 1823–1837
Vision Research
www.elsevier.com/locate/visres

Effects of local and global factors in the Pinna illusion

Rick Gurnsey *, Geneviève Pagé

Department of Psychology, Concordia University, Montréal, Que., Canada H4B 1R6

Received 21 June 2005; received in revised form 9 September 2005

Abstract

The Pinna illusion (Pinna & Brelstaff, 2000) consists of two concentric rings of micropatterns that appear to counter-rotate when the observer moves towards the stimulus. There have been several reports that the illusion is stronger when the retinal expansion is produced by observer self-motion than when produced on a computer screen without observer self-motion. In fact, we found that the illusion is as strong (or stronger) when the retinal expansion is produced on a computer screen without observer self-motion. In a second series of experiments the strength of the Pinna illusion was inferred from the amount of physical counter-rotation required to null it. The strength of the illusion is relatively unaffected by changes to the global structure of the display but minor changes to the micropatterns comprised in the display can effectively eliminate the illusion. We provide a simple model of optical flow that is in very good agreement with many of the results reported.

Keywords: Motion; Illusions; Optical flow; Computational models; Psychophysical methodologies

1. Introduction

There are several striking illusions of motion arising from static images [e.g., the Ouchi illusion (Ashida, 2002; Ashida, Kitaoka, & Sakurai, 2005; Hine, Cook, & Rogers, 1997; Khang & Essock, 1997a, 1997b; Ouchi, 1977; Spillmann, Heitger, & Schüller, 1986; Spillmann, Tulunay-Keesey, & Olson, 1993), Kitaoka's "rotating snakes" illusion (Conway, Kitaoka, Yazdanbakhsh, Pack, & Livingstone, 2005) and variants of the Fraser and Wilcox (1979) "escalator illusion" that arises from asymmetric luminance-gradients (Faubert & Herbert, 1999; Naor-Raz & Sekuler, 2000)]. Fig. 1A shows another remarkable motion illusion that was discovered by Baingio Pinna and reported by Pinna and Brelstaff (2000). When the centre of Fig. 1A is fixated and moved towards the eyes, a compelling illusory rotation of the two rings of squares (micropatterns) composing the display is experienced. The illusion-like the other motion illusions just cited—demonstrates a clear failure of the visual system to solve the aperture problem (Bayerl & Neumann, 2002; Gurnsey, Sally, Potechin, & Mancini, 2002; Morgan, 2002; Pinna & Brelstaff, 2000). As the stimulus expands and contracts on the retina each micropattern travels along a straight line that connects it to the centre of the display, and yet the micropatterns do not appear to follow a straight path. Several recent papers have suggested that motion illusions such as the Pinna and Ouchi illusions can be explained by an "orthogonal bias" (Mather, 2000). On this account the visual system produces an interpretation of image flow that is biased towards the strongest normal velocities (i.e., velocities perpendicular to 1D contours) in the image. This kind of bias can be shown to occur when the process of optical flow estimation is contaminated by spatiotemporal noise (e.g., Fermüller & Malm, 2004; Fermüller, Pless, & Aloimonos, 2000; Weiss, Simoncelli, & Adelson, 2002; Weiss & Fleet, 2002).

This role of orthogonal bias in the original Pinna illusion (Fig. 1A) can be understood by considering the low frequency luminance gradients present in the micropatterns as shown in Fig. 1B. When Fig. 1A expands on the retina the two micropatterns directly to the right of centre (shown in Fig. 1B, top) translate to the right. The low frequency

* Corresponding author. Tel.: +1 514 848 2424 ext. 2243; fax: +1 514 8484545.
E-mail address: Rick.Gurnsey@concordia.ca (R. Gurnsey).

0042-6989/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2005.09.011

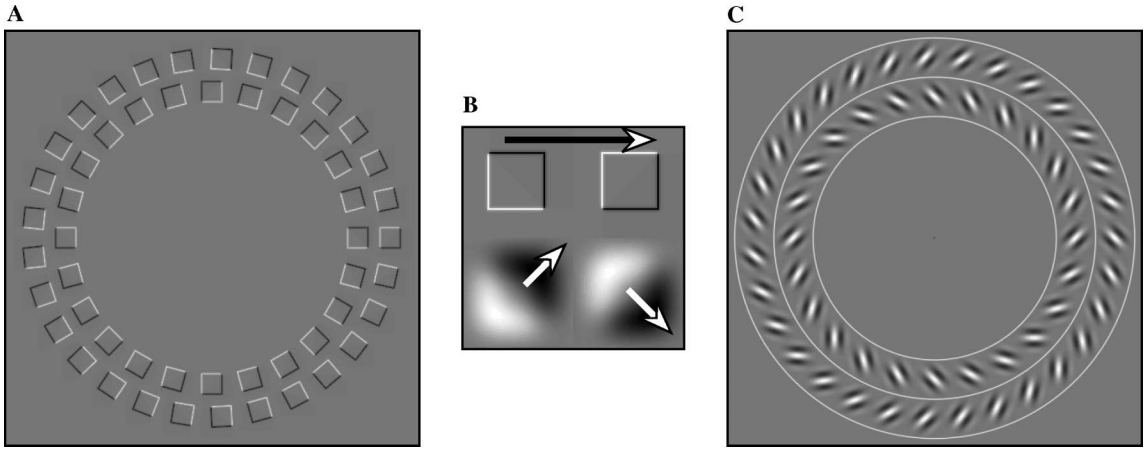

Fig. 1. (A) The original version of the Pinna illusion. When fixating on the central black dot and moving toward the paper, one should experience the circles counter-rotating. (B) The two micropatterns at the top have strongly oriented low-spatial frequency components. The two bottom micropatterns depict a blurred version of the two above, i.e., high frequencies have been removed from the image. When the micropatterns translate to the right the stimuli will most strongly stimulate neurons selective for the directions indicated by the white arrows (normal velocities). (C) A variant of the Pinna illusion comprising Gabor Patches.

luminance gradients present in the display (Fig. 1B, bottom) are oriented ±45° to the true 2D flow and their normal velocities (speed of motion in a direction perpendicular to the luminance gradient) would excite direction selective mechanisms tuned to up-to-the-right and down-to-the-right as indicated by the white arrows. This pattern of activation is replicated along all lines of expansion resulting in a clockwise flow component in the outer ring and a counterclockwise flow component in the inner ring. The importance of low-frequency luminance gradients is demonstrated when the box-like micropatterns of Fig. 1A are replaced with Gabor patches (Fig. 1C). Morgan (2002), Gurnsey et al. (2002), and Bayerl and Neumann (2002) showed that these displays produce a very compelling illusion. Gurnsey et al. demonstrated that the strength of the illusion depends on the number of Gabor patches in the display, their wavelengths, and the orientation difference between adjacent patterns in the inner and outer rings.

It might seem that the Pinna illusion can be understood completely in terms of the so-called orthogonal bias. However, there have been anecdotal reports that the illusion is strongest when the observer moves towards the paper on which the illusion is printed than when the paper is moved towards the eyes (Gurnsey et al., 2002; Morgan, 2002; Spillmann, Pinna, Stürzel, & Werner, 2003). If it were true that the illusion's strength depends on the manner in which identical retinal motions are created (self-motion vs. screen-motion) then extra-retinal contributions to the strength of the illusion would be clearly implicated. In fact, there are perceptual differences when retinal motions are produced in these two ways (Morgan, 2002). Under conditions of self-motion there is size constancy; the two rings of micropatterns appear to counter-rotate within a static ring of constant size. It is reasonable to suppose that vestibular inputs or efference-copy contribute to size constancy. Under conditions of screen-motion observers are clearly aware of the retinal expansion and contractions. It may be that the strength of the illusion is affected by differences in extra-retinal inputs under conditions of self-motion and screen-motion. In fact, Morgan (2002) made exactly this case, as did Gurnsey et al. (2002). However, a difference in salience between screen-motion and self-motion has not been demonstrated empirically so it is not clear whether an appeal to extra-retinal input is actually required to explain the illusion.

To assess the role of extra-retinal contributions to the Pinna illusion we conducted a two-part experiment; each part required participants to make *relative salience judgments*. In the first part, we measured the orientation dependence of the illusion under conditions of self-motion and screen-motion using the relative-salience method of Gurnsey et al. (2002). In this paradigm participants are presented with pairs of stimuli that differ only in the orientation difference between the micropatterns in the inner and outer rings (see Fig. 2) and asked to judge which of the two produced the stronger illusion. Using a set of stimuli comprising several orientation differences, the relative salience of each orientation difference could be established. The second part compared directly the strength of the illusion arising from self-motion and screen-motion.

A second series of studies was conducted to measure the *absolute strength* of the Pinna illusion under a variety of conditions. Rather than assessing the relative salience of particular orientation differences, the absolute strength of several orientation differences was measured using a nulling procedure (Bayerl & Neumann, 2002). In this paradigm, the strength of the illusion is inferred from the amount of physical rotation opposite to the illusory rotation that is needed to null the illusory percept. The nulling procedure was used to determine which stimulus factors affect the absolute strength of the illusion. These factors included the number or rings in the stimulus, the speed of stimulus motion and the characteristics of the constituent micropatterns, among several others.

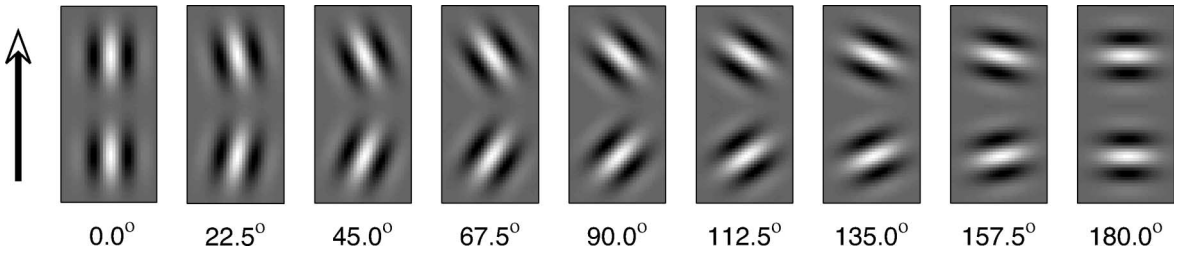

Fig. 2. Examples of stimuli from the inner (bottom) and outer (top) rings of displays like the one shown in Fig. 1C. Each pair is described in terms of the orientation difference between the micropatterns relative to upward motion. Lines parallel to the true 2D flow have a 0° orientation difference and those perpendicular to the true 2D flow have a 180° orientation difference.

2. Experiment 1a: Orientation dependence of the Pinna illusion

The question addressed in Experiment 1a is whether identical patterns of retinal stimulation elicit different "orientation tuning curves" when produced by self-motion or screen-motion. The experiment employed the method of relative salience judgments used by Gurnsey et al. (2002). Participants were presented with all possible pairings of nine orientation differences and the relative strength of illusion produced by each orientation difference was determined by the number of times it was chosen as producing the stronger illusion. The tuning curves produced in this way for self-motion were compared to those associated with screen-motion.

2.1. Method

2.1.1. Stimuli

Fig. 3A provides an example stimulus used in Experiment 1a. All stimuli comprised 24 Gabor patches in the inner ring and 30 in the outer ring so that micropattern density was the same within the two rings. Each Gabor patch was defined as

$$g(x,y) = \mathrm{e}^{-\frac{(x^2+y^2)}{2\sigma^2}} \cos\left[\frac{2\pi}{\lambda}(x\cos\theta - y\sin\theta)\right], \quad (1)$$

where λ is the wavelength in pixels, θ is the orientation and σ is the spread of the Gaussian envelope. The nine orientation differences ranged from 0° to 180° in 22.5° steps. Nominally, 0° indicated that adjacent micropatterns in the inner an outer rings were both oriented parallel to the line that connects their centres to the centre of the display, and 180° indicated that adjacent micropatterns in the inner an outer rings were both perpendicular to the line that connects their centres to the fixation dot.

In the screen-motion condition each micropattern moved back and forth along a notional line emanating from the centre of the display. This motion simulated the expansions and contractions that would occur if the participant viewed a stationary stimulus while moving towards the screen (starting at 67 cm and moving to 47 cm) then away from the screen (i.e., back to 67 cm). Because self-motion cannot be initiated or stopped instantly, the stimulus size on each frame was proportional to the value of a cosine function with a wavelength of 40 cm (i.e., the cosine function goes through half a cycle in 20 cm). There were $N = 30$ frames in the forward sequence so the simulated distance from the display on each frame of forward motion is described as

$$\text{distance}_i = \cos(\pi * \mathit{frame}_i / N_{\text{frames}}) * 10 + 57. \quad (2)$$

In this way the change in simulated distance from frame to frame depended on the frame number; the size change was greatest at $\text{frame}_{N/2}$ and smallest at frame_1 and frame_{N-1}. Thirty frames were created and played forwards and backwards to create one cycle of expansion and contraction; frame_N was presented only once so there were $N * 2 - 1 = 59$ images in the sequence. Each image was presented for two video refreshes on a monitor having a refresh rate of 85 Hz. Therefore, each image played for $2 * 1000/85 = 23.53$ ms for a total sequence duration of $59 * 23.53 = 1388.27$ ms.

The size of each micropattern (determined by λ and σ) was scaled with simulated distance to the screen. At a simulated distance of 57 cm, each micropattern had $\lambda = 9.57$ pixels and $\sigma = 4.78$ pixels. The centres of the inner and outer rings were 143.6 and 169.1 pixels from screen centre. At each simulated distance these standard values were scaled by $57/\text{distance}_i$. The two rings of micropatterns were surrounded by "guides" (see Fig. 3A) whose Gaussian cross-sections were scaled with simulated distance in the same way as the Gabor patches.

2.1.2. Participants

There were eight participants (six females and two males) recruited from the Concordia community. Two participants had previously participated in similar psychophysical experiments and the remaining six were naïve to the experimental procedures.

2.1.3. Apparatus

The experiments were conducted using a Macintosh G4 connected to a 21-in. multiscan colour monitor with display resolution set at 1024×768 pixels. Pixel width was .37 mm and the screen refresh rate was 85 Hz. The gamma correction software available in the Psychtoolbox (Brainard, 1997) was used to linearize the screen luminance and a Minolta CS-100 photometer was used to find the absolute luminance levels. Stimuli were created and exper-

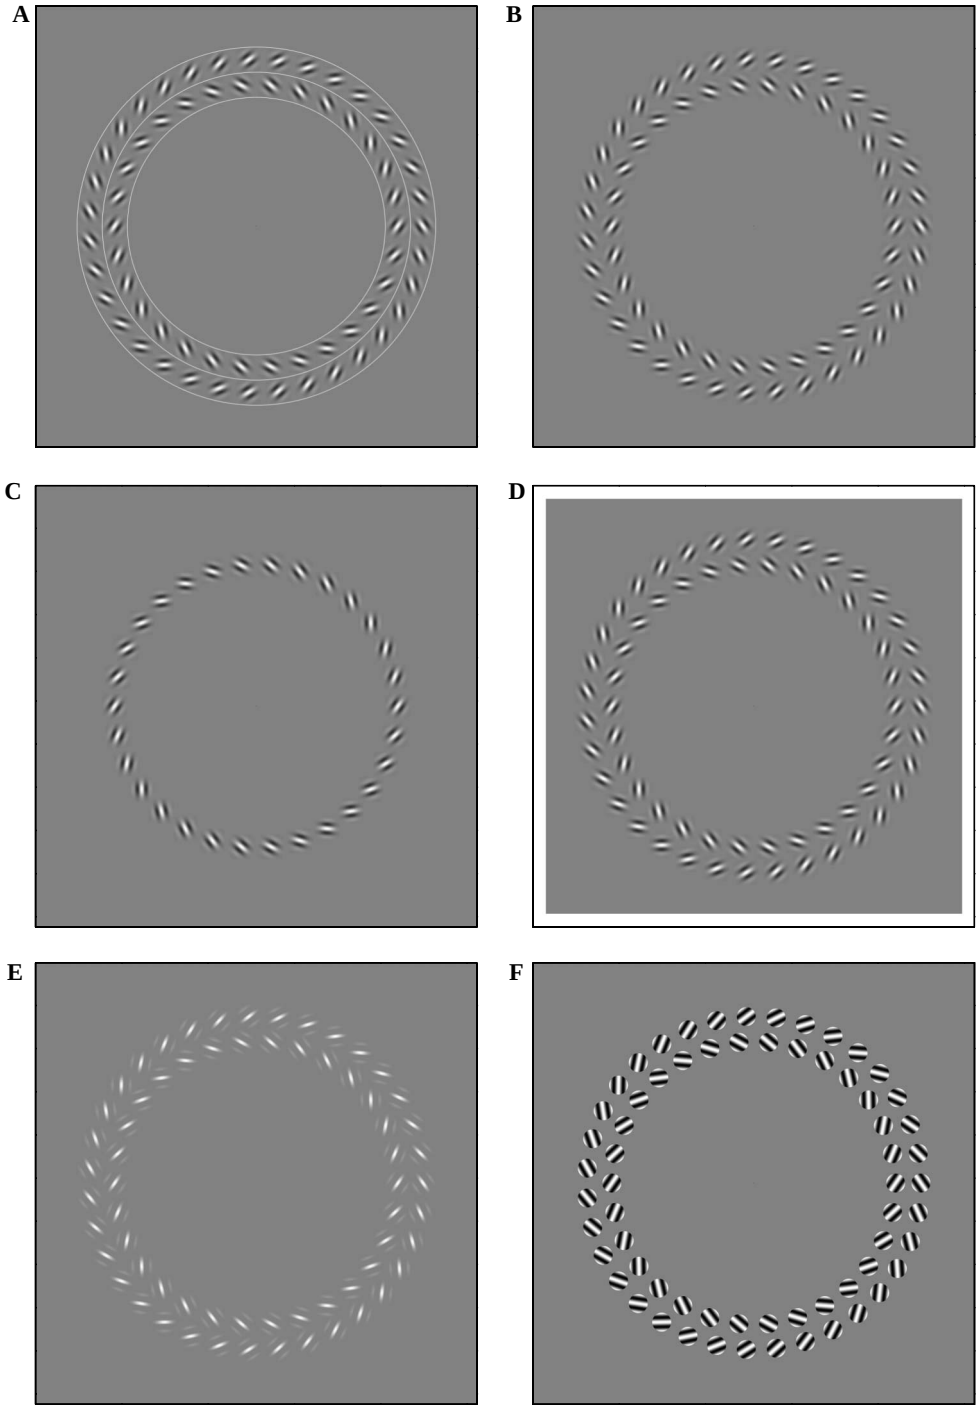

Fig. 3. Examples of the stimuli used in Experiments 1–5. (A) Two rings of micropatterns with "guides." (B) Two rings of micropatterns without guides. (C) One ring of micropatterns. (D) Two rings, no guides, surround by a square frame. (E) Half-wave rectified Gabors. (F) Cosine gratings within a circular, step-edge aperture.

iments were run in the MATLAB (Mathworks) environment using functions in the Psychtoolbox (Brainard, 1997) that provide high level access to the routines of the VideoToolbox (Pelli, 1997).

2.1.4. Procedure

2.1.4.1. Screen-motion conditions. The participants viewed the screen from a distance of approximately 67 cm and a chin rest was used to maintain this viewing distance. On each trial the participant viewed a fixation dot in the centre of the screen and two stimuli were presented with an inter-stimulus interval of 250 ms. The two stimuli differed only in the orientation difference between the micropatterns in the inner and outer rings (see Fig. 2). After each trial the participant indicated whether the first or second stimulus produced the stronger sense of motion by clicking once or

twice on the mouse. The next trial began immediately following a valid response. (A link to a screen-motion trial is provided in the Appendix A. The file is in Quicktime format.)

There were nine stimuli (each representing a specific orientation difference) so there were ${}_9C_2 = 36$ possible pairings of stimuli. Within a block of trials each of the 36 pairs was presented once. The order of presentation of the pair elements was chosen randomly on each trial. There were three blocks of trials and each began with an instruction screen that reminded the participant to click once if the first stimulus produced the more salient motion and twice if the second produced the more salient motion. The block was initiated by a mouse-click.

2.1.4.2. Self-motion conditions. At the beginning of each self-motion trial the participant sat at 67 cm from the screen. This was achieved by adjusting a neck-strap connected to a chin rest that was fixed at 47 cm from the screen. Prior to each trial a "motion calibration" screen appeared which comprised just the three guide rings of Fig. 3A expanding and contracting exactly as they did in the screen-motion condition. In addition, the central fixation dot (5×5 pixels) changed from white to black to white over the same period as the expansion and contraction of the guides (1388.27 ms). Two hundred and fifty milliseconds following the offset of the guides and fixation dot a stimulus appeared along with the fixation dot. The participant moved back and forth from the rest position (67 cm) to the chin rest (47 cm) and back to the rest position in synchrony with the luminance variation of the fixation dot. Two hundred and fifty milliseconds after the offset of the fixation dot the second stimulus appeared and the self-motion was repeated. At the end of the trial the participant indicated which of the two stimuli produced the stronger sense of illusory motion.

2.2. Results

The dependent measure was the number of times that each stimulus (i.e., orientation difference) was judged to produce the stronger sense of illusory motion. Within each block, each stimulus was presented once with each of the other eight stimuli and so it could be chosen a maximum of eight times as producing the more salient motion and a minimum of 0 times. There were three blocks of trials so the maximum "salience score" was 24 and the minimum was 0. Fig. 4A shows plots of the relative salience measure (expressed as a proportion of the maximum score) as a function of orientation difference for the screen-motion and self-motion conditions. (Note that data from one participant were excluded from the statistical analyses because they seemed to show that on a large percentage of trials she reversed the response options.) It is clear that the pattern of relative salience judgments does not change as a function of how the retinal motion was produced.

For each participant the relative salience curve was fit with a third degree polynomial and the peak of the curve was found. The averaged performance peak was at an orientation difference of 65.17° in the self-motion condition and at 73.46° in the screen-motion condition. The self-motion and screen-motion peaks for each participant were submitted to a paired t test and were found to be not statistically different [$t(6) = -1.73$, $p > .05$]. The results of Experiment 1a indicate that the relative salience judgments are little affected by the manner in which otherwise identical patterns of retinal motion are produced.

3. Experiment 1b: Relative salience of illusory motion produced by self-motion and screen-motion

It is possible for the relative salience measures to be identical and yet differ substantially in absolute terms. To directly test the idea that self-motion produces a stronger illusion we had participants judge the salience of illusions arising from self-motion and screen-motion. On each trial participants were presented with the same stimulus under screen-motion and self-motion conditions and judged which condition produced the stronger illusion.

3.1. Method

3.1.1. Stimuli/apparatus/participants

The stimuli, apparatus and participants were the same as in Experiment 1a.

3.1.2. Procedure

A trial consisted of the two presentations of the same retinal motion, once produced by self-motion and once produced by screen-motion. Within a block each of the nine orientation differences was presented once. There were 12 blocks and in half the blocks screen-motion was presented first and in half self-motion was presented first.

3.1.2.1. Screen-motion-first trials. At the beginning of each block an instruction screen was presented to remind participants that a screen-motion trial would be followed by a self-motion trial. A trial began with a stimulus presentation exactly like the screen-motion condition of Experiment 1a followed by a stimulus presentation exactly like the self-motion condition of Experiment 1a; i.e., the participant moved forwards and backwards in synchrony with the fixation dot. There was no ISI between the two presentations. Following the self-motion part of the trial the participant clicked once or twice to indicate which interval produced the more salient illusory motion.

3.1.2.2. Self-motion-first trials. At the beginning of each block an instruction screen was presented to remind participants that a self-motion trial would be followed by a screen-motion trial. A trial began with a stimulus presentation exactly like the first part of a self-motion condition of Experiment 1a in which a "guides only" sequence preceded

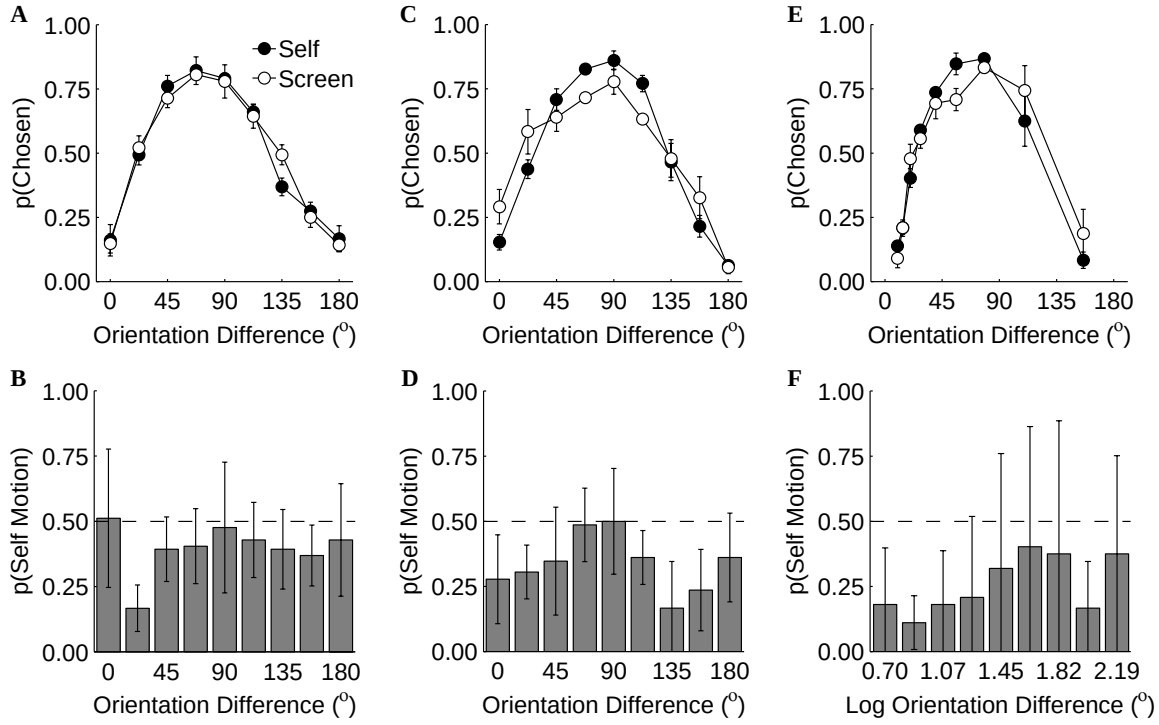

Fig. 4. Results of Experiment 1. (A) Proportion of times each orientation difference was chosen as producing the stronger illusion. $N = 7$, error bars = ±SEM. Filled circles represent self-motion conditions and unfilled circles represent screen-motion conditions. (B) Probability of reporting that self-motion produced the more salient illusion as a function of orientation difference. $N = 7$, error bars = 95% CI. Results of Experiment 2. (C and D) Same as (A and B) for three experienced observers. Results of Experiment 3. (E and F) Same as (C and D) for logarithmically spaced orientation differences.

the presentation of the first self-motion stimulus. The self-motion component of the trial was followed immediately by a stimulus presentation exactly like the screen-motion condition of Experiment 1a. Therefore, in both the screen-motion-first and self-motion-first conditions, the self-motion stimulus was always preceded by a motion stimulus that acted as an exemplar to help the participants co-ordinate their motions. As in the screen-motion condition there was no ISI between the two presentations. Following the self-motion part of the trial the participant clicked once or twice to indicate which interval produced the more salient illusory motion. Participants were given as many practice blocks as they required to get comfortable with the procedure.

3.2. Results

The dependent measure was the number of times that each motion type (i.e., self-motion vs. screen-motion) was chosen as producing the stronger sense of illusory motion for each orientation difference. Because there were 12 blocks of trials the maximum "salience score" was 12 and the minimum was 0. Fig. 4B shows plots of the proportion of times screen-motion was chosen as more salient for each orientation difference. If self-motion conditions produce stronger illusions then proportions should be greater than .5. In fact, there is a general tendency for screen-motion to be preferred. For the orientation differences of 22.5° and 157.5° the 95% confidence intervals did not include .5 indicating a statistically significant preference for screen-motion. Although the confidence intervals around the remaining conditions include .5, in 7 of 7 conditions in which the micropatterns in the inner and outer rings differed in orientation, the proportions shown in Fig. 4B were less than .5; this has a probability of 0.008 under the binomial distribution assuming a 50% chance of success. Therefore, there is no evidence that self-motion produces a stronger illusory motion. To the contrary, the available evidence suggests a preference for screen-motion.

4. Experiments 2 and 3: Replication with experienced observers

Experiment 2 was a replication of Experiment 1 except that three experienced observers served as participants. Experiment 3 differed from Experiment 2 only in terms of the orientation differences used. Rather than nine equal linear steps between 0° and 180° there were nine equal logarithmic steps between 10° and 156°. Experiment 3 was run mainly for comparison with results of Experiment 4, described below.

The results of Experiment 2 are summarized in Figs. 4C and D. Again, there is little evidence of different tuning curves for the screen- and self-motion conditions (Fig. 4C); in fact the average peak orientation differences were almost identical (76.2° and 76.3°). And, when the salience of the

screen- and self-motion is directly compared (Fig. 4D) the preference (when there is one) is for screen-motion.

The results of Experiment 3 are summarized in Figs. 4E and F. The general pattern of results is the same as for Experiments 1 and 2. The tuning curves in Fig. 4E are shifted slightly to the left relative to those in Figs. 4A and C. For self-motion the peak was at 58.1° and for screen-motion it was at 64.9°. This shift occurs because of the asymmetric distribution of orientation differences presented to the participants.

The results shown in Figs. 4B, D, and F indicate that there is no preference for screen-motion or self-motion for orientation differences near the peaks of the tuning curves in Figs. 4A, C, and D. The preference for screen-motion over self-motion appears for non-optimal orientation differences. In other words, when the illusion is strongest there is no preference for screen-motion over self-motion but as the illusion weakens participants show a preference for screen-motion over self-motion.

5. Experiment 4a: Measuring the absolute strength of the illusion

In the previous experiments we assessed how judgments about the relative salience of stimuli varies as a function of the orientation difference between micropatterns in the inner and outer rings. Although these experiments provide useful answers to certain questions they do not provide an absolute measure of the strength of a particular illusory motion. A time-honoured way of assessing the absolute strength of a percept is by measuring the amount of opposing stimulation needed to null it (e.g., Hurvich & Jameson, 1957). Thus, an obvious way to assess the absolute strength of any version of the Pinna illusion is to measure the amount of physical rotation counter to the illusory rotation that is needed to null the illusion.

There were two reasons for conducting experiments using the nulling procedure. The first is to compare the results with those of the relative salience procedure to determine whether the two measures produce consistent results. The second is to use an objective methodology to answer a number of questions about the display characteristics that determine the strength of the illusion.

5.1. Method

5.1.1. Stimuli

The stimuli were similar in most respects to those used in the screen-motion condition of Experiment 1 with the following exceptions. During a trial, stimuli only expanded on the screen (as opposed to expanding and contracting). As well, the step size from frame to frame was constant instead of following a cosine function. Again there were nine versions of the basic stimulus with the orientation difference between micropatterns in the inner and outer rings as the independent variable. Unlike the stimuli in Experiments 1 and 2 the orientation differences ranged from 10° to 156° in equal logarithmic steps. Logarithmic steps were used because the tuning curves from Experiments 1 to 3 seemed quite broad and we wanted to measure the effect of smaller gradations of orientation differences (below the peak) without increasing the number of intervals.[1]

For each of the nine orientation differences there were nine levels of counter-rotation ranging in linear steps from 0 to 11.35 angular deg/s. For the inner ring the counter-rotation was in the counter-clockwise direction and in the outer ring it was clockwise. (A link to a file showing a stimulus near the PSE is provided in the Appendix A. The file is in Quicktime format.)

5.1.2. Participants

One of the authors (GP) and two other experienced psychophysical observers (PY and SM) served as participants. All had normal or corrected to normal vision and had participated in many psychophysical tasks.

5.1.3. Apparatus

The apparatus was exactly as in Experiment 1.

5.1.4. Procedure

On each trial a single stimulus was presented and the participant judged whether the *inner ring* appeared to rotate in a clockwise or counter-clockwise direction. In all conditions the illusory motion was clockwise. A single mouse click indicated clockwise and two clicks indicated counter-clockwise. For each orientation difference, each of the nine counter-rotations was presented once within each of 20 blocks of trials. Thus, for each orientation difference there were 180 trials. The nine conditions were run in a different random order for each of the three participants.

5.2. Results

A single click (indicating clockwise) was scored as a 1 and a double click (indicating counter-clockwise) was score as a -1. Means across the 20 replications of each counter-rotation were computed. Means for all orientation differences were generally close to 1 when the counter-rotation was 0 (i.e., participants experienced the illusory motion) and dropped to -1 when the counter-rotation was 11.35 (i.e., participants experienced the true rotary motion). To compute this point of subjective equality (PSE) for a particular orientation difference the means across the nine counter-rotations were fit with a Gaussian integral scaled to the range $[-1, 1]$. The mean of the best-fitting Gaussian corresponds to the PSE. Fig. 5A shows plots of PSE as a function of orientation difference for each of the three participants.

[1] Of course, the trade-off here is a lack of precision as the orientation differences increase. However, this imprecision does not limit our ability to answer questions about the relationship between the nulling and relative-salience tasks or about the stimulus factors influencing the strength of the illusion.

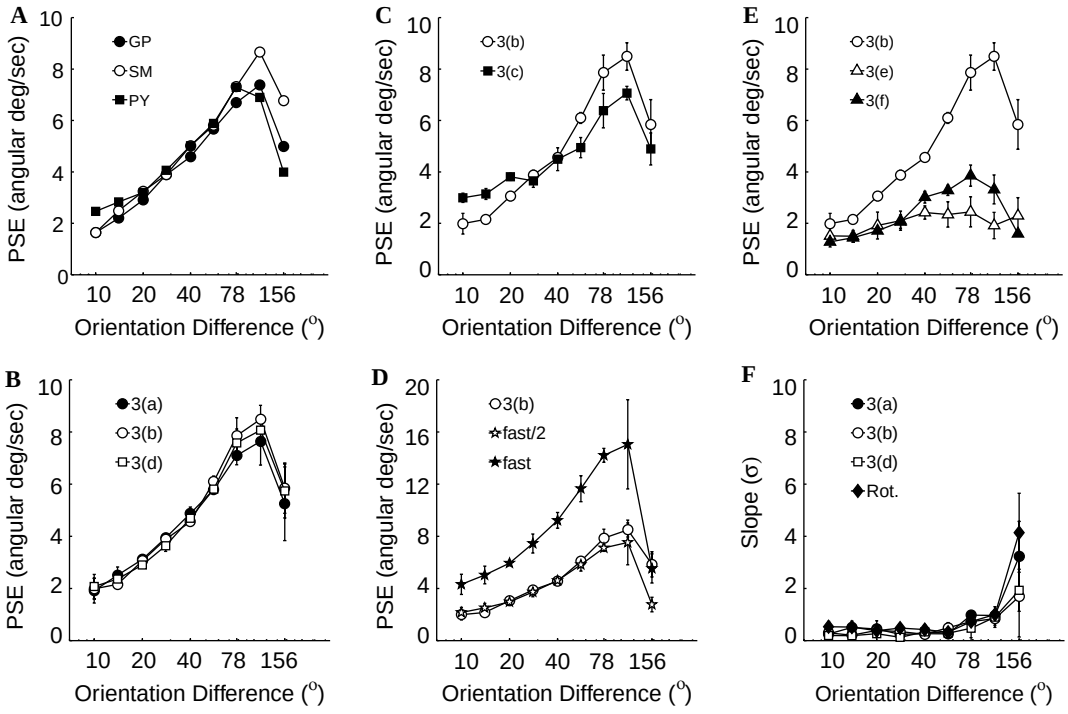

Fig. 5. Results of Experiments 4a–4f and Experiment 5. Please note that legends refer to the stimuli shown in Fig. 3. (A) Results of Experiment 4a. Stimuli were two rings and guides (Fig. 3A) for three observers. (B) Results of Experiments 4a (filled circles: two rings with guides, Fig. 3A), 4b (unfilled circles: two rings, Fig. 3B), and 4d (unfilled squares: two rings with frame, Fig. 3D); $N = 3$, error bars = ±SEM. (C) Results of Experiment 4b (unfilled circles: two rings, Fig. 3B) and 4c (filled squares: one ring, Fig. 3C); $N = 3$, error bars = ±SEM. (D) Results of Experiment 4b (unfilled circles: two rings, Fig. 3B) and 4f. The filled stars represent PSEs when the stimuli from Fig. 3B were presented at double speed. The unfilled stars show the "double-speed PSEs" divided by two. $N = 3$, error bars = ±SEM. (E) Results of Experiment 4b (unfilled circles: two rings, Fig. 3B) and 4e. The unfilled triangles represent the PSEs obtained using the rectified stimulus shown in Fig. 3E and the filled triangles represent the PSEs obtained with the broadband stimuli shown in Fig. 3F. $N = 3$, error bars = ±SEM. (F) Standard deviations of the best fitting psychometric functions obtained in Experiments 4a, 4b, 4d, and 5. In all cases the average standard deviation of the best fitting Gaussian integral is plotted as a function of orientation difference. $N = 3$, error bars = ±SEM.

The results plotted in Fig. 5A show a very similar pattern for the three participants. The average orientation difference yielding the strongest illusion was 94.2° (92.8°, 103.9°, and 86.1° for GP, PY, and SM, respectively). (Peaks were obtained using the best-fitting third degree polynomials.) On the log scale there is a linear increase in the strength of the illusion with increasing orientation difference and then an abrupt drop once the peak is reached. It is worth noting that the average peak of 94.2° is slightly lower than the average peak at 112° reported by Gurnsey et al. (2002) in an experiment that used the method of relative salience judgments. (Note that when the data reported by Gurnsey et al. are fitted with third order polynomials, the average performance peak is at 101.4°.) However, a peak at 94° peak is substantially larger than the average peak of 64.9° found in Experiment 3, which also employed the method of relative salience judgments. We discuss reasons for this divergence in Section 8. For now we note that the nulling procedure produces reliable results and thus provides a tool for addressing questions about the factors that determine the strength of the illusion.

6. Experiments 4b–4f: Factors influencing the strength of the illusion

Experiments 4b–4f used exactly the same methodology and participants as Experiment 4a to assess factors that determine the strength of the Pinna illusion.

6.1. Experiment 4b: Importance of "guide rings"

In all experiments so far we employed stimuli in which the Gabor patches were surrounded by rings, which we have called guides. These guides were also used by Gurnsey et al. (2002) because informal observations suggested that they enhanced the strength of the illusion. The results of a theoretical and empirical analysis of the Pinna illusion by Bayerl and Neumann (2002) suggest that the illusion is enhanced when local normal velocities equal to the true 2D image velocity neighbour the Gabor patches. Bayerl and Neumann demonstrated this effect using displays much like ours. In one case they superimposed on each Gabor patch a second Gabor that was higher in spatial frequency and oriented perpendicular to the line connecting it to the

centre of the display. Thus, when the high frequency Gabor translates away from the centre of the display its normal velocity equals the true, local 2D image velocity. The guides in our stimuli provide exactly the same kind of information but are broadband rather than narrowband. Therefore, the analysis and results of Bayerl and Neumann suggest that eliminating the guides should diminish the strength of the illusion. To assess this suggestion we replicated Experiment 4a in all respects except that we eliminated the guides from the displays (see Fig. 3B).

6.2. Experiment 4c: Number of rings

Bayerl and Neumann (2002) also argued that the strength of the Pinna illusion involves interactions between the two rings of micropatterns and this lead them to expect that eliminating one of the rings should reduce the strength of the illusion. However, their results show only modest differences between one and two ring versions of the experiment. To assess the importance of having one vs. two rings we replicated Experiment 4b in all respects except that only the inner ring of Gabor patterns was present (see Fig. 3C).

6.3. Experiment 4d: Frame of reference

Experiment 1b showed no evidence that self-motion yields a stronger illusion than screen-motion; in fact the opposite seemed to be true. However, our screen-motion condition differs somewhat from the case in which a stimulus is printed on a piece of paper and moved towards the eyes. In the former situation, the expansion and contraction of the stimulus occurs within a fixed frame of reference (i.e., the outline of the computer monitor) whereas in the latter the figure expands on the retina while a frame (the piece of paper) expands as well. It might be argued that the expanding frame of reference provided by the outline of the paper provides information to the visual system that the display is part of an object moving towards the observer. This information might weaken the strength of the illusion. To examine this possibility we replicated Experiment 4b but surrounded the figure with a white square frame (shown Fig. 3D) that expanded and contracted with the concentric rings.

6.4. Experiment 4e: Stimulus speed

It is clear that the Pinna illusion is orientation tuned but it is not clear why. And, it is not clear whether there is a link between the orientations of the micropatterns in the displays and the speed at which they are moving. It might be that the strength of the illusion depends on the speed of the normal velocity of each micropattern. So, if the illusion is optimal for a particular speed then its orientation tuning should change as a function of stimulus speed. On the other hand if it is just the orientation difference in the inner and outer rings that affects the illusion there should be no effect of stimulus speed on the *shape* of the orientation tuning-curve. To examine this issue we used the two-ring display (Fig. 3B) that expanded at twice the speed used in Experiments 4a–4d.

6.5. Experiment 4f: Stimulus bandwidth

An interesting difference between the Gabor versions of the Pinna illusion used here and the original examples used by Pinna and Brelstaff (2000) [and later by Bayerl and Neumann (2002)] concerns the bandwidth of the stimuli. Gabor patches are narrowband in the spatial frequency domain and the micropatterns in Fig. 1A are broadband. Yet, in both cases strong illusions are produced. This suggests that the spatial frequency bandwidth of the stimuli has no effect on the strength of the illusion so that minor modifications to the spatial structures of the Gabor displays that change the stimulus bandwidth should have little effect on performance. To assess this implication, two modified versions of the standard stimulus used in Experiment 4b were created. Fig. 3E shows a *half wave rectified* display in which all luminance values below the mean luminance are set to the background luminance. This rectification manipulation seems relatively minor but it substantially increases the spatial frequency bandwidth of the stimulus. The Gabor patches were modified in a second way by replacing the Gaussian window with a step-edged circular aperture (Fig. 3F). Again, this manipulation substantially increases the spatial frequency—and orientation—bandwidth of the patch. With these broadband stimuli we repeated the conditions of Experiment 4b. Rather than using counter-rotations of 0 to 11.35 angular deg/s, we used counter-rotations of −2.88 to 8.64 (where −2.88 indicates rotations in the same direction as the illusory rotation) because pilot testing indicated that this manipulation substantially reduced the illusion. Therefore, the counter-rotations of −1.44 and −2.88 angular deg/s were included to ensure that at least some of the displays appeared to rotate in the illusory direction.

6.6. Results

6.6.1. Experiments 4b and 4d

Fig. 5B summarizes the results of Experiments 4b (two rings and no guides: unfilled circles, see Fig. 3B) and 4d (two rings surrounded by a frame: unfilled squares, see Fig. 3D) averaged over the three participants, along with the averaged data of Experiment 4a (two rings and guides: filled circles, see Fig. 3A). (Please note that in all panels of Fig. 5 the legend entries refer to the stimuli depicted in the indicated panels of Fig. 3.) All three curves are essentially identical. Therefore, the suggestion by Gurnsey et al. (2002) that the "guides" might enhance the salience of the illusion is not supported by the present data. As well, these data do not seem to support the idea that normal velocities consistent with the radial flow enhance the illusion as suggested by Bayerl and Neumann (2002).

Surrounding the display with a square frame that expands and contracts at the same rate (Experiment 4d) clearly has no effect on the strength of the illusion.

6.6.2. Experiment 4c

Fig. 5C plots the results of Experiment 4c (one ring: filled squares, see Fig. 3C) along with the results of Experiment 4b (two rings: unfilled circles, see Fig. 3B). Using only one ring of micropatterns produced a moderate increase in the illusion strength (relative to Experiments 4a, 4b, and 4d) at small orientation differences and a moderate decrease in the illusion strength at large orientation differences. We have no explanation for the small increase in illusion strength for small orientation differences. The modest decrease in illusion strength at larger orientation differences (i.e., 56–156°) suggests possible synergistic effects between biased velocity estimates of opposite directions in neighbouring regions (Bayerl & Neumann, 2002). Perhaps the resulting motion contrast enhances the illusion. It must be said, however, that these differences are relatively small compared to the differences resulting from other manipulations to be described below.

6.6.3. Experiment 4e

The results of Experiment 4e (two rings presented at double speed: filled stars, see Fig. 3B) are plotted in Fig. 5D along with the results of Experiment 4b (two rings: unfilled circles, see Fig. 3B). The expansion of the stimulus in Experiments 4a–4d required .7 and .35 s in 4e. Doubling the speed of the stimulus expansion doubles the amount of counter-rotation required to null the illusion. This can be seen when the PSEs from Experiment 4e are divided by 2 (Fig. 5D, unfilled stars). In this case the results of Experiment 4b and 4e largely overlap except at the 156° orientation difference. There is no indication that the faster curve is shifted to the left as might be expected if the illusion depended on an optimal speed of motion normal to the micropattern orientation. Rather, the illusion seems to be largely orientation dependent.

6.6.4. Experiment 4f

The most dramatic effects in Experiment 4 were produced by the bandwidth manipulation in Experiment 4f. Fig. 5E compares the results of Experiment 4b (two rings: unfilled circles, see Fig. 3B) with those of the two broadband conditions in Experiment 4f. Increasing the bandwidth of the stimulus effectively eliminated the illusion for the rectified patterns (two rings rectified Gabors: unfilled triangles, see Fig. 3E) and dramatically reduced it for the hard-edged apertures (two rings of micropatterns within step-edge apertures: filled triangles, see Fig. 3F). These results may reveal important principles of direction selectivity in the human visual system. (A link to a file showing the effect of rectification is provided in the supplemental material at the end of this document. The file is in Quicktime format.)

7. Experiment 5: Orientation tuning discrepancies between Experiments 3a and 4a

Experiment 3a (see Fig. 4E) and Experiment 4a (see Fig. 5B) provide two measures of the orientation tuning of the Pinna illusion for three experienced psychophysical observers; the results are clearly discrepant. If the results of the nulling procedure revealed the mechanisms underlying the relative salience judgments then the peak orientation difference should be the same in both cases. A comparison of Figs. 4E and 5B indicates that the peaks are not the same. Furthermore, Fig. 5B shows that the PSE for the largest (156°) orientation difference is about 6 angular deg/s which is greater than the PSEs for orientation differences of 40° and less. This implies that if the mechanisms subserving relative salience judgments are the same as those subserving the nulling procedure, then an orientation difference of 156° should be judged as more salient than orientation differences of 40° and less. Fig. 4E shows this is clearly not the case.

The discrepancy between the two tasks may be explained by differences in task requirements. In Experiment 3, participants had only to compare two percepts and decide which was more salient. Experiment 4, although apparently simpler in some sense, might actually be a less direct measure of the strength of the illusion. In the nulling procedure the strength of the illusion is inferred from the amount of counter-rotation required to cancel the illusion. However, the orientations of the Gabors composing the stimulus may affect both the probability of seeing illusory motion (in the case of no counter-rotation) and the ability to detect a counter-rotation. Therefore, an orientation difference that produces a relatively weak illusion in a relative salience task might produce a rather large PSE if it is difficult to detect its physical counter-rotation.

One indication of these two contributions to observed PSE can be seen in the standard deviations of the psychometric functions used to compute the PSEs. As mentioned, the PSEs were computed by fitting Gaussian integrals to the direction judgment data in Experiment 4. Each Gaussian has a mean (μ) and standard deviation (σ). (The μ value gives the PSE.) Fig. 5F plots the average σ values as a function of orientation for Experiments 4a, 4b, and 4d, all of which involved two rings of micropatterns (see Figs. 3A, B, and D) presented at the normal speed. It is clear that the standard deviations of the functions increase as the orientation difference increases. Fig. 5B shows that orientation differences of 56° and 156° (in Experiments 4a, 4b, and 4c) both elicit PSEs of about 6 angular deg/s but it can be seen in Fig 5F that the standard deviations of the psychometric functions are much smaller for 56° than 156°. This difference in standard deviations means that there is more uncertainty for judgments around the PSE at 156° than at 56°. This may indicate difficulty detecting the counter-rotation of the stimuli. A final experiment was conducted to assess the possibility that PSEs may be affected by difficulty detecting the physical counter-rotation of the stimulus.

7.1. Method

The participants, apparatus and procedure were essentially identical to those of Experiment 4b with the following exceptions. On each trial the two rings of the stimulus rotated by did not expand. The participants' task was to report the direction of rotation of the inner ring (clockwise or counter-clockwise). There were 9 levels of rotation of the inner rings ranging from −2.88 (counter-clockwise) to 2.88 (clockwise) angular deg/s.

7.2. Results

In all cases the PSEs were very close to 0 ($M = .06$ angular deg/s) however there was some bias at 156° ($M = -.5$ angular deg/s) indicating that the PSE was a slight clockwise rotation. There were substantial differences in the σ values across orientation differences. These are shown as filled diamonds in Fig. 5F. For 156° in particular, the standard deviations are extremely large indicating a great deal of uncertainty around the PSE. In other words, participants become less sensitive to rotations as the orientation difference between the rings increases. This is consistent with the idea that the large PSEs found at large orientation differences in Experiment 4 were not a pure measure of the strength of illusory motion but arise in part because of the difficulty participants have detecting physical rotations as the orientations of the Gabor patches get closer to parallel with the direction of motion.

7.3. A simple model

As mentioned in Section 1, a number of recent papers have pointed out that optical flow estimates may be biased towards the strongest normal velocity within the region of integration when noise contaminates the local velocity signals (e.g., Weiss & Fleet, 2002) or the spatiotemporal derivatives (Fermüller et al., 2000). To relate these arguments to the present results we implemented a simple optical flow model and examined its responses to drifting columns of Gabor patches in the presence of varying degrees of Gaussian noise.

The optical flow constraint equation (Horn & Schunck, 1981) tells us that

$$f_x u_x + f_y u_y + f_t = 0, \tag{3}$$

where u_x and u_y are velocity components, f_x and f_y are spatial derivatives, and f_t is the temporal derivative, at image point $I(x, y, t)$. More concisely

$$\mathbf{f}_s \mathbf{u} + f_t = 0, \tag{4}$$

where $\mathbf{u} = (u_x, u_y)$ and $\mathbf{f}_s = (f_x, f_y)$. In our model f_x, f_y and f_t were computed within 5^3 pixel volumes around image point $I(x, y, t)$. If derivatives are computed at N locations around $I(x, y, t)$ we can let $\mathbf{I}_s = ([f_{x,1}, f_{y,1}], \ldots, [f_{x,N}, f_{y,N}])$ and $\mathbf{I}_t = (f_{t,1}, \ldots, f_{t,N})$; i.e., $\mathbf{I}_s$ is an $N \times 2$ matrix and $\mathbf{I}_t$ is an N dimensional vector. The least squares estimate of 2D image flow at position (x, y, t) is obtained by

$$\hat{\mathbf{u}} = -(\mathbf{I}_s' \mathbf{I}_s)^{-1} \mathbf{I}_s' \mathbf{I}_t. \tag{5}$$

In our model $\hat{\mathbf{u}}$ was estimated at position (x, y) from the responses in a 32×32 region centred on (x, y). To calculate the response of the model to a translating stimulus we averaged the $\hat{\mathbf{u}}$ values estimated obtained at each (x, y) location through which the stimulus travelled.

The column of Gabors in Fig. 6A was chosen to approximate a segment of a pattern such as shown in Fig. 3C. Each Gabor patch was defined within a 32×32 pixel window and the stimulus translated to the right over five frames at a rate of 2 pixels per frame. We were interested in the effects of Gabor orientation and noise on estimates of image flow. The Gabor orientations were those used in Experiment 4 and the noise at each pixel (Weiss & Fleet, 2002) was drawn from a zero mean Gaussian distribution. The standard deviation of the noise (σ) was the same for x, y, and t. The values of σ were chosen somewhat arbitrarily to range from 0 to 1, with 1 being the maximum value in the Gabor patches. For each orientation and noise level, 10 flow estimates were computed and then averaged.

Selected results of the simulations are summarized in Fig. 6E, which plots error in the estimated direction of flow

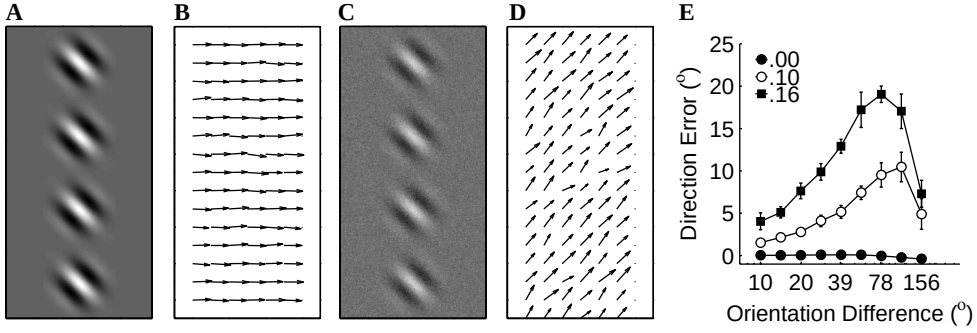

Fig. 6. (A) An example of the stimulus used in the simulations. The stimulus contained no noise and translated to the right over five frames at a rate of 2 pixels per frame. (B) An example flow field recovered from (A). (C) Same stimulus as shown in (A) but with added Gaussian noise. (D) An example flow field recovered from (C). (E) Plot of orientation bias as a function of orientation difference for three levels of added noise. Please note that although only one column of micropatterns was used in simulations the x-axis is labelled to be consistent with Fig. 5. (See text for further details.)

as a function of stimulus orientation and noise level. When there is no noise in the signal (e.g., Fig. 6A) the direction of flow is recovered correctly for all orientations (e.g., Fig. 6B). When noise is added to the stimulus (e.g., Fig. 6C) the direction estimate is biased towards the velocity normal to the orientation of the Gabor (e.g., Fig. 6D). For a given noise level the degree of bias depends on the orientation of the Gabor. In fact, for noise = .1 the magnitude of the direction bias across orientations shows a pattern remarkably similar to the pattern of results found in Experiment 4 (cf. Fig. 5B). As noise level increases direction errors increase for each orientation difference and the orientation difference producing the maximum errors shifts to the left.

The results of the simulations show that when an appropriate level of noise is added to stimuli similar to those used in the experiments, the direction errors induced in a relatively standard optical flow algorithm nicely parallel the strength of the Pinna illusion as measured using the nulling procedure. This suggests that in large part the Pinna illusion can be understood as a direct consequence of intrinsic noise within the spatiotemporal filters that recover optical flow within the human visual system.

Although there is a very nice correspondence between the behavior of the model and human performance for the examples shown in Fig. 6, there is a substantial divergence between the two in the case of broadband stimuli. When the rectified stimuli (Fig. 3E) are submitted to the model, the results are identical to those obtained with narrowband stimuli. Thus, the model is not affected by stimulus bandwidth in the same way that human performance is. Possible reasons for this difference are considered in the following section.

8. General discussion

8.1. Screen-motion vs. self-motion

Experiments 1, 2, and 3 do not support the notion that extra-retinal inputs (such as efference copy or vestibular signals) play an important role in determining the salience of the Pinna illusion. In fact, when there is a difference, it is the screen-motion condition that yields the more salient illusion. These results raise two questions: (a) why might self-motion produce a stronger illusion than "paper-motion" and (b) why does screen-motion produce a stronger illusion than self-motion in certain conditions of Experiments 1–3? The answer to both of these questions might have to do with the *smoothness* of the retinal motions produced under conditions of "paper-motion," self-motion and screen-motion because these three conditions can be ordered with respect to the smoothness of the retinal motions they produce. When one holds a piece of paper and moves it towards the eyes, motion jitter might arise from paper flutter (caused by air resistance), slight hand tremors and error prone muscular contractions required to move the paper. Self-motion on the other hand involves no issues of flutter or hand tremors and the resulting motion is more ballistic with fewer sources or error. Screen-motion eliminates all sources attributable to muscle noise and hence produces the smoothest retinal motion. Thus, simply assuming that the illusion is strongest when retinal motion is smoothest reconciles claims that the illusion is strongest under self-motion conditions (Gurnsey et al., 2002; Morgan, 2002) with the present data, which show the opposite. The fact that screen-motion produces an illusion that is as strong or stronger than self-motion is fortunate because it means that various aspects of the illusion can be studied in conditions over which we have the greatest control.

8.2. Origins of the illusion

Experiments 4a–4d showed that many structural (i.e., global) aspects of the stimuli had little or no effect on its strength, as measured by the nulling procedure. The strength of the illusion is not affected by guide rings, or a moving frame and only modestly affected by the number of rings present. Combined with the results of Experiments 1–3, the results of Experiments 4a–4f suggest that the characteristics of micropatterns making up the display are the main determinants of the illusion's strength. Furthermore, we showed that a very simple model to compute optical flow is in excellent agreement with the data from the nulling experiments. Central to this model is the important role of noise. In a noiseless signal the model always recovers the true 2D image flow. But, when noise is added the model produces flow estimates that are biased towards the strongest normal velocities in the display. The degree of this bias is related to the amount of noise in the system (either internal or external) and the orientations of the micropatterns. Therefore, we conclude that the illusion is essentially attributable to a failure of the early visual system to solve the aperture problem in the presence of noise. We see no need to posit the involvement of high-level feedback or extra-retinal inputs to explain the general features of the illusion. Therefore, we think a model somewhat simpler than that offered by Bayerl and Neumann (2002) might be sufficient to capture the essentials of the illusion.

8.3. Orientation tuning

All experiments reported here measured the orientation tuning of the Pinna illusion, as did Experiment 1 in Gurnsey et al. (2002). Experiments 1a, 2a, and 3a measured orientation tuning-curves for screen-motion and self-motion conditions using very similar procedures. The mean tuning curve peaks for the self- and screen-motion conditions of Experiment 1a were 65.2° and 73.5°, respectively. For Experiment 1b these means were 76.2° and 76.3°, respectively, and, for Experiment 1c these means were 58.1° and 64.9°, respectively. Therefore, these means covered an 18° range with a mean of 69°. A reanalysis of Gurnsey et al. (2002), Experiment 1, revealed a mean peak of 101.4°, which is substantially larger than the peaks found in the present study. This difference is difficult to explain. Of course the participants were different in the two cases so

it could be a simple matter of sampling error. As well, there were stimulus differences between the two studies. At a viewing distance of 57 cm the micropatterns in the Gurnsey et al. (2002) stimuli were 6.67° and 7.85° from fixation whereas in the present experiment they were 5.3° and 6.2° from fixation. In the Gurnsey et al. study the one cycle of expansion and contraction took 2 s whereas in the present experiment on cycle took 1.4 s. Finally, there were 30 micropatterns in the inner ring of the stimuli used by Gurnsey et al. and 24 in the present experiment. Therefore, it might be that stimulus position, speed and density conspire to move the peak of the tuning curves to higher orientation differences.

The peaks of the tuning curves in Experiments 1a, 2a, and 3a (relative salience judgments) were consistently lower than those of Experiments 4a, 4b, 4c, 4d, and 4e (nulling experiments), which were about 96° on average. It was shown in Experiment 5 that the difference in orientation tuning between the nulling experiments and the relative-salience experiments could be attributable to task differences. In particular, the increased peaks in the nulling experiments may reflect difficulty detecting the physical counter-rotation thereby artificially inflating the estimated strength of the illusion.

8.4. The issue of bandwidth

Experiment 4f showed that relatively minor changes to the micropatterns produce very large changes in the strength of the illusion. In fact, a half-wave rectification of the stimulus (Fig. 3E) essentially eliminates the illusion. This finding is unexpected and probably has important implications for a general understanding of how optical flow is computed in the human visual system. The bandwidth effect shows that the visual system fails to solve the aperture problem only when the micropatterns making up the stimulus are narrowband; or, perhaps more precisely, when the predominant normal velocities producing the illusion are restricted to a narrow band of relatively low frequencies, as in the original Pinna illusion (see Figs. 1A and B). We consider several explanations for the bandwidth effect.

In a very recent study of the Ouchi illusion Ashida et al. (2005) showed that stimuli comprising thin lines (rather than the more usual rectangles or gratings) elicit salient illusory motion. However, the salience of the illusion decreases monotonically with increasing positional jitter of the lines. Both the jittered and non-jittered stimuli are broadband and have equivalent energy, but the energy dispersion is greater in the case of the jittered stimuli. Ashida et al. point out that jittering the stimuli decreases the energy at the fundamental frequency of the display (and at higher harmonics) and they attribute the reduction in the illusion's strength to these energy reductions. However, in this case energy reductions were confounded with increased energy dispersion (e.g., see Ashida et al., Fig. 3) and so cannot provide either a necessary or sufficient account of the jitter effect. Nevertheless, the effect of rectification in Experiment 4f also has the effect of reducing energy at the peak frequency of each Gabor micropattern and might explain the bandwidth effect (at least for the rectified stimulus) in Experiment 4f. This explanation can be assessed informally by looping the Appendix A (''broadband.mov'') while reducing the screen contrast. Even large reductions in contrast have little effect on the illusion and certainly do not weaken the illusion induced by the standard stimulus to the point of that induced by the full contrast, rectified stimulus. Therefore, we think that energy dispersion is a more likely account of reduced salience of the both the Ouchi- and Pinna-illusions.

Given the dependence of bias on the presence of noise we might conclude that broadband signals have lower net internal noise than narrowband signals. It is likely that optical flow is computed from the outputs of a large number of filters each of which may be contaminated by independent internal noise. It may be that when information is combined across scales the effect is to average out these independent sources of noise. This seems an unlikely explanation because the effect of noise should be similar across scales. If optical flow computations were carried out within a single spatial frequency band the result would always be a bias in the direction of the dominant normal velocities. Therefore, it seems unlikely that combining information from different scales would produce the desired ''averaging-out'' effect. In any case, being broadband is sufficient to eliminate the illusion; the original Pinna stimuli were broadband and elicited striking illusions. The same is true of the original Ouchi- and Kitaoka-illusions.

In broadband signals (e.g., Fig. 3E) there are phase-coincidences across scales (Marr & Hildreth, 1980; Morrone & Burr, 1988). These may be taken as non-accidental events that have the same underlying physical cause. If such coincidences normally form the input to an optical flow computation they might effectively attenuate noise in the system. Because there is a higher probability that a phase coincidence will arise from a common physical cause (e.g., a translating surface-, luminance- or material-discontinuity) than from an accidental coincidence of structure at different scales (arising from independent noise sources), internal responses that are unrelated to the moving stimulus (i.e., noise) would be eliminated.

Our model shows that direction bias occurs in a standard optical flow algorithm only when there is noise in signals. Therefore, a simpler explanation for the bandwidth effect is that high-frequency channels have less internal noise. In fact, Gurnsey et al. (2002) showed that in judgments of relative salience high frequency Gabors produced a weak illusion. Gurnsey et al. also showed that very low frequency Gabors produced a weak illusion. In either case, the bandwidth effect may not reveal the consequences of integration across scales but differences in internal noise at different scales. That is, as the bandwidth of a signal increases (or as the energy dispersion increases) the effect may be to introduce significant energy into channels having

less internal noise, permitting a less biased estimate of image flow.

The present results are not sufficient to distinguish these different possibilities. However, from the preceding discussion it is clear that relatively straightforward follow-up studies may shed light on which are more or less plausible.

8.5. Relation to motion transparency

The Pinna illusion would seem to bear some relationship to transparent motion plaids (e.g., Adelson & Movshon, 1982; von Grünau & Dubé, 1993). When two drifting gratings (sine wave or square wave) are superimposed they will appear either as a single coherent grating or two transparent gratings, one sliding over the other. Therefore, transparency may also be seen as a failure to integrate two local motions to produce a common motion much like the Pinna illusion reflects a failure to integrate local motions into a coherent local flow.

The tendency to see transparency or coherence depends on many factors, including the direction difference (normal velocities) between the two gratings; as the direction difference increases transparency is increasingly likely to be seen (e.g., Adelson & Movshon, 1982; Gurnsey & von Grünau, 1997; Kim & Wilson, 1993). The largest normal velocity direction difference in the present stimuli is associated with the smallest orientation difference; i.e., for an orientation difference of 10° micropatterns in the inner and outer rings are oriented ±5° from the direction of true 2D flow, producing a normal velocity direction difference of 170°. As the orientation difference increases the direction difference decreases. Over the range of 10–156° orientation difference (see Fig. 4), the Pinna illusion shows the opposite dependence on direction difference; as the orientation difference increases—and direction difference decreases—the illusion becomes stronger. Therefore, the principles governing motion transparency seem to be different from those governing the Pinna stimuli.

9. Conclusions

- We find no evidence that self-motion produces a more compelling illusion than screen-motion. It seems likely that the smoothness of the retinal motion (rather than extra-retinal inputs) determines the strength of the illusion.
- The nulling procedure shows that many global/structural changes to the stimulus displays have little effect on the illusion. The strength of the illusion is not affected by the number or rings, the presence of "guides" of a frame that expands and contracts along with the stimulus, nor is the orientation tuning of the illusion affected by the speed of expansion.
- The largest effects we observed arise from changes to the micropatterns comprised in the display. In particular broadband signals produce a far weaker illusions than narrowband signals. Therefore, the strength of the illusion seems to depend on local factors rather than global/structural factors.
- The Pinna illusion can be understood as an "orthogonal bias" that arises from internal noise in the mechanisms that compute optical flow.
- The relative-salience and nulling procedures reveal different dependencies on orientation differences. The nulling procedure may misrepresent the actual perceptual salience of the illusion because it confounds the strength of the biased illusory motion signal with sensitivity to physical counter-rotation.
- The reasons for the bandwidth effect are unclear. It is interesting to note that noise-dependent biases in flow estimates from a standard optical flow model are similar for narrowband stimuli (Fig. 3B) and broadband stimuli (Fig. 3E). Apparently, this is not the case for psychophysical observers, leaving an intriguing question to pursue.

Acknowledgment

This research was supported by an NSERC Research Grant to Rick Gurnsey.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.visres.2005.09.011.

References

Adelson, E. H., & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. *Nature, 300*, 523–525.

Ashida, H. (2002). Spatial frequency tuning of the Ouchi illusion and its dependence on stimulus size. *Vision Research, 42*, 1413–1420.

Ashida, H., Kitaoka, A., & Sakurai, K. (2005). A new variant of the Ouchi illusion reveals fourier-component-based processing. *Perception, 34*, 381–390.

Bayerl, P., Neumann, H., 2002. Cortical mechanisms of processing visual flow—Insights from the Pinna–Brelstaff illusion. In Proceedings of the 4th Workshop Dynamische Perzeption, Bochum. http://www.informatik.uni-ulm.de/ni/mitarbeiter/PBayerl/homepage/.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 443–446.

Conway, B. R., Kitaoka, A., Yazdanbakhsh, A., Pack, C. C., & Livingstone, M. S. (2005). Neural basis for a powerful static motion illusion. *Journal of Neuroscience, 25*, 5651–5656.

Faubert, J., & Herbert, A. M. (1999). The peripheral drift illusion: A motion illusion in the visual periphery. *Perception, 28*, 617–621.

Fermüller, C., & Malm, H. (2004). Uncertainty in visual processes predicts geometrical optical illusions. *Vision Research, 44*, 727–749.

Fermüller, C., Pless, R., & Aloimonos, Y. (2000). The Ouchi illusion as an artifact of biased flow estimation. *Vision Research, 40*, 77–96.

Fraser, A., & Wilcox, K. J. (1979). Perception of illusory movement. *Nature, 281*, 565–566.

Gurnsey, R., Sally, S. L., Potechin, C., & Mancini, S. (2002). Optimizing the Pinna–Brelstaff illusion. *Perception, 31*, 1275–1280.

Hine, T., Cook, M., & Rogers, G. T. (1997). The Ouchi illusion: An anomaly in the perception of rigid motion for limited spatial frequencies and angles. *Perception & Psychophysics, 59*, 448–455.

Horn, B. K. P., & Schunck, B. (1981). Determining optical flow. *Artificial Intelligence, 17*, 185–203.

Hurvich, L. M., & Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review, 64*, 384–404.

Khang, B. G., & Essock, E. A. (1997a). Apparent relative motion from a checkerboard surround. *Perception, 26*, 831–846.

Khang, B. G., & Essock, E. A. (1997b). A motion illusion from two-dimensional periodic patterns. *Perception, 26*, 585–597.

Kim, J., & Wilson, H. R. (1993). Dependence of plaid motion coherence on component grating directions. *Vision Research, 33*, 2479–2489.

Marr, D., & Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society London B, 207*, 187–217.

Mather, G. (2000). Integration biases in the Ouchi and other visual illusions. *Perception, 29*, 721–727.

Morgan, M. (2002). Square up the circle. The Guardian Thursday, 24 January 2002. http://www.guardian.co.uk/Archive/Article/0,4273,4341518,00.html.

Morrone, M. C., & Burr, D. C. (1988). Feature detection in human vision: a phase-dependent energy model. *Proceedings of the Royal Society London B, 235*, 221–245.

Naor-Raz, G., & Sekuler, R. (2000). Perceptual dimorphism in visual motion from stationary patterns. *Perception, 29*, 325–335.

Ouchi, H. (1977). *Japanese optical and geometrical art*. New York: Dover.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437–442.

Pinna, B., & Brelstaff, G. (2000). A new illusion of relative motion. *Vision Research, 40*, 2091–2096.

Spillmann, L., Heitger, F., & Schüller, S. (1986). Apparent displacement and phase unlocking in checkerboard patterns. Paper presented at the 9th European Conference on Visual Perception, Bad Nauheim.

Spillmann, L., Tulunay-Keesey, U., & Olson, J. (1993). Apparent floating motion in normal and stabilised vision. *Investigative Ophthalmology and Visual Science, 34*(4), Abstract 1611, p. 1031.

Spillmann, L., Pinna, B., Stürzel, F., & Werner, J. S. (2003). Extraretinal factors required for visual illusions [Abstract]. *Journal of Vision, 3*(9), 656, http://journalofvision.org/3/9/656/, doi:10.1167/3.9.656.

Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience, 5*, 598–604.

Weiss, Y., & Fleet, D. J. (2002). Velocity likelihoods in biological and machine vision. In R. P. N. Rao, B. A. Olshausen, & M. S. Lewicki (Eds.), *Probabilistic models of the brain: Perception and neural function* (pp. 77–96). Cambridge: MIT Press.