Motion direction influences surface segmentation in stereo transparency

Ross Goutcher

Faculty of Natural Sciences, Psychology Division, University of Stirling, Stirling, Scotland, UK

To perceive multiple overlapping surfaces in the same location of the visual field (transparency), the visual system must determine which surface elements belong together, and should be integrated, and which should be kept apart. Spatial relations between surfaces, such as depth order, must also be determined. This article details two experiments examining the interaction of motion direction and disparity cues on the perception of depth order and surface segmentation in transparency. In Experiment 1, participants were presented with random-dot stereograms, where transparent planes were defined by differences in motion direction and disparity. Participants reported the direction of motion of the front surface. Results revealed marked effects of motion direction on perceived depth order. These biases interact with disparity in an additive manner, suggesting that the visual system integrates motion direction with other available cues to surface segmentation. This possibility was tested further in Experiment 2. Participants were presented with two intervals: one containing motion and disparity defined transparent planes, the other containing a volume of moving dots. Interplane disparity was varied to find thresholds for the correct identification of the transparent interval. Thresholds depended on motion direction: Thresholds were lower when disparities and directions in the transparency interval matched participants' preferred depth order, compared to conditions where disparity and direction were in conflict. These results suggest that motion direction influences the judgment of depth order even in the presence of other visual cues, and that the assignment of depth order may play an important role in segmentation.

Introduction

Natural environments often contain instances of surface transparency, where a surface is viewable behind another overlapping surface. Such surface transparency can arise due to surface translucency, where light passes through the front surface (e.g., glass, water) or from pseudotransparencies (Tsirlin, Alison, & Wilcox, 2008), where the rear surface is seen through gaps in the front surface (e.g., when looking through a chain-link fence, or through window blinds). When confronted with instances of surface transparency, the visual system must determine which overlapping points should be grouped together and which should be kept apart. To perceive pseudotransparency, the visual system performs these dual processes of surface integration and segmentation using differences in element motion, and/or binocular disparity-defined depth. These types of pseudotransparency are typically referred to as motion transparency and stereo transparency. In this article, we examine the interactions between these two cues.

For motion transparency, the visual system relies on differences in direction of motion and differences in spatial frequency content. In luminance-plaid stimuli, motion transparency is only perceived when component gratings differ in their directions of motion by more than 135° (Kim & Wilson, 1993; Wilson & Kim, 1994). Plaids whose component motions are closer together are perceived as moving in a single direction (Adelson & Movshon, 1982). For stimuli containing luminance gratings moving in opposite directions, the perception of transparency depends on relative spatial frequency content, with motion transparency only perceived once gratings differ in spatial frequency by a factor of four or more (Cavanagh & Mather, 1989; Levinson & Sekuler, 1975). Motion transparency may also be perceived for oppositely moving, same spatial frequency gratings, when one grating is luminance-defined, and the other contrast-defined (Goutcher & Loffler, 2009).

Effects of direction and spatial frequency content are also evident with random-dot stimuli. Spatially correlated random-dot patterns are not perceived as moving transparently, analogous to results with same spatial frequency gratings (Qian, Andersen, & Adelson, 1994a). Transparency may still be perceived in such patterns, however, if surfaces are separated in depth through the addition of binocular disparity (Qian et al.,

Citation: Goutcher, R. (2016). Motion direction influences surface segmentation in stereo transparency. *Journal of Vision*, *16*(15):17, 1–15, doi:10.1167/16.15.17.

doi: 10.1167/16.15.17 Received May 22, 2016; published December 22, 2016 ISSN 1534-7362

1994a). The addition of binocular disparity affects the perception of motion transparency in several other ways. Hibbard and Bradshaw (1999) found that signal-to-noise thresholds for the detection of motion transparency improved, and approached thresholds for the detection of a single direction of motion, if transparent surfaces differed in disparity. Similarly, Calabro and Vaina (2006) found that signal-to-noise thresholds for motion transparency discrimination were affected by the combination of disparity and differences in motion direction: When surfaces differed in disparity, motion transparency could be reliably detected with smaller differences in direction.

These findings on the interaction of disparity and motion processing in transparency have been taken to show that the measurement of disparity precedes the encoding of global motion patterns. In demonstrating the benefits of binocular disparity signals for the segmentation of multiple overlapping directions of motion, such research has, however, neglected another aspect of transparency perception—the ordering of surfaces in depth. While depth ordering has typically been investigated in stereo transparency (e.g., Tsirlin et al., 2008), such ordering is also a feature of pure motion transparency stimuli. In such stimuli one surface of dots is typically seen as lying behind the other surface. Although inherently ambiguous, several factors are known to affect depth ordering in motion transparency. Prior research has shown that element size, density, and relative velocity affect depth ordering, as do wavelength and duty cycle in periodic stimuli (Moreno-Bote, Shpiro, Rinzel, & Rubin, 2008; Schütz, 2011). Other factors shown to affect transparency depth ordering include motion adaptation (Schütz, 2011), surface usefulness (Chopin & Mamassian, 2011), and motion direction (Mamassian & Wallace, 2010).

For motion direction, Mamassian and Wallace (2010) found that observers maintain a bias to perceive particular directions of motion as "in front" in motion transparency displays. Such biases are idiosyncratic between observers, but remain consistent within each observer over extended periods of weeks, or even years (Mamassian & Wallace, 2010; Schütz, 2014; Wexler, Duyck, & Mamassian, 2015). This consistency is surprising, given that there would seem to be little reason for any particular direction of motion to either directly signal, or be indirectly correlated with, a particular ordering in depth. In other cases, such as effects of wavelength or element size, biases in depth ordering have been linked to the statistical regularities of natural scenes (Moreno-Bote et al., 2008). Directional biases have instead been attributed to the functioning of perceptual memory (Schütz, 2014; Wexler et al., 2015), with exposure to ambiguous stimuli shown to affect perceptual learning (Harrison & Backus, 2010; van Dam & Ernst, 2010).

In this article we examine the interactions between motion direction preference and the depth-ordering signals provided by binocular disparity. Our results suggest that the assignment of depth order, and the segmentation of surfaces in depth, is determined through the integration of binocular disparity and motion direction preference signals. As such, despite its uncertain foundations, the visual system appears to treat motion direction as if it provides a reliable signal for depth ordering. Together, these results suggest that the assignment of depth order plays an important role in surface segmentation.

Experiment 1: Motion and disparity signals interact in the assignment of depth order

Previous research has shown that preferences to perceive shorter wavelength surfaces as in front can counteract disparity-defined depth order (Moreno-Bote et al., 2008). This finding is surprising, since it suggests that perceptual preference in ambiguous stimuli can act against other cues, such as binocular disparity, that provide a more obvious signal for depth ordering. Two factors limit these previous findings. First, this work did not measure thresholds for depth ordering, instead using a fixed disparity signal to measure changes in perceptual preference. Second, the authors offered no means of assessing, or accounting for, the effect of wavelength on disparity depth ordering. Additionally, though not a limiting factor, Moreno-Bote et al. (2008) suggested that wavelength preferences for depth ordering may be related to natural scene statistics. The observed effect is therefore somewhat less surprising, as differences in surface wavelength may be considered as a true cue to depth. Here, we examine the interactions between disparity cues to depth ordering, and motion direction preferences, for which there is no readily apparent informational basis, and use disparity thresholds for depth ordering to provide a metric for the strength of the motion direction bias.

Methods

Participants

Seventeen participants completed Experiment 1, including author RG. All participants were staff or students of the University of Stirling, including six participants who completed the experiment in partial fulfilment of course requirements. All participants had normal or corrected-to-normal vision, with stereoacuity of <1 arcmin, as measured by the RanDot2 test (VAC, Elk Grove Village, IL). Participants were provided with details of the structure of the stimulus, but were unaware of specific experimental hypotheses and conditions. All participants gave written consent, and experimental procedures were subject to approval from the local psychology ethics board, in accordance with the guidelines of the British Psychological Society and the Declaration of Helsinki.

Stimulus and apparatus

Observers were presented with random-dot stereograms containing dots moving in opposing directions, resulting in the perception of motion transparency. Dot disparity and motion direction were varied under a method of constant stimuli in order to recover psychometric functions for the effects of disparity on perceived surface order, for each motion direction. Six directions were tested, including horizontal and vertical motions, in equal steps of 30° (Figure 1a). Disparity was added through offsetting of horizontal dot positions across a range of ±3.3 arcmin (i.e., at the maximum disparity, front and back surfaces were separated by a relative disparity of 6.6 arcmin). Dots were white circles of diameter 5.5 arcmin, and were presented against a black background at a dot density of 1.95 dots per degree^2.

Dots moved in opposing directions at a speed of 1.1°/s. There was no variability in motion direction, meaning that, on any given trial, each dot moved in one of only two possible directions. Stimuli were presented for 2 s each, resulting in a total displacement of 2.2° for each dot, on each trial. Dots were presented within a circular aperture of diameter 5.5°, and were surrounded by a circular zero disparity reference frame, comprised of overlapping white rectangles. The position of this aperture was fixed, and dots were randomly replaced as they moved beyond the aperture edge.

All stimuli were presented using a MacPro computer (Apple, Inc., Cupertino, CA), together with an Apple HD Cinema display with a resolution of 1920 × 1200 pixels and refresh rate 60 Hz, at a viewing distance of 76.4 cm. At this distance, each pixel measured 1.1 arcmin. The display was calibrated using a Spyder2Pro calibration device (ColorVision, DataColor, Lawrenceville, NJ) to ensure a linear luminance scale. Stimulus generation and presentation was controlled using Matlab (Mathworks, Inc., Natick, MA), together with the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997). Stereoscopic stimulus presentation was achieved using a modified Wheatstone stereoscope, calibrated to ensure no conflict between accommodation and angle of convergence. Stimulus presentation occurred in a darkened room with participants' head movements restricted using a Headspot chinrest (UHCO, Houston, TX).

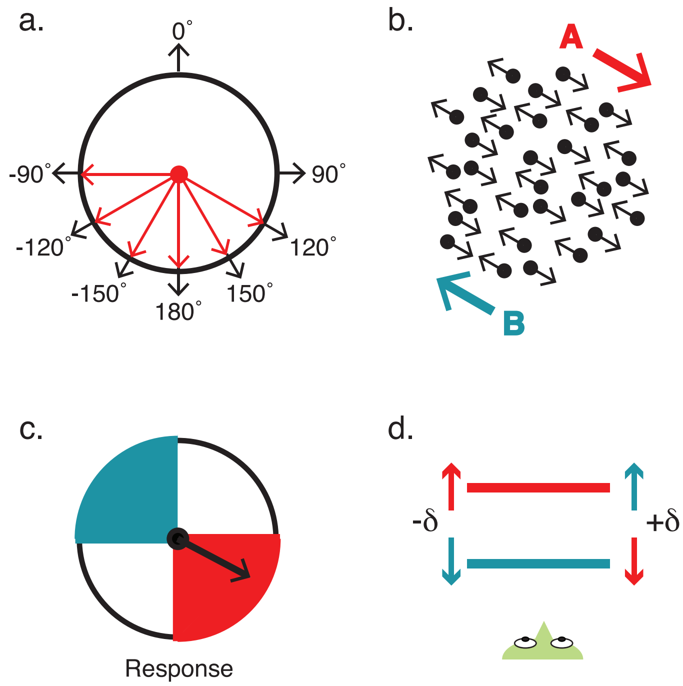

Figure 1. Illustration of the design for Experiment 1. (a) Participants were presented with motion transparency stimuli, containing opposing motion along six orientations. Motion directions were coded as the angular direction of one of the sets of dots, referred to as "Surface *A*." Surface *A* directions are shown by the red arrows. (b) Example of the opposing motions of a transparency stimulus, with color coding for Surface *A* (red) and opposing Surface *B* (blue). (c) When making a response, the participant dragged a line in the direction of the front surface. Responses were only accepted if they fell within ±45° of either of the true stimulus directions of motion. (d) Manipulations of disparity δ were such that positive δ values moved Surface *A* closer to the observer, while moving Surface *B* farther from the observer. Negative δ values operated in the opposite direction.

Design and procedure

Each participant was presented with seven levels of disparity, at six motion orientations, resulting in 42 experimental conditions. These were presented in random order to each participant, over multiple blocks, for a minimum of 20 repeated trials of each condition, for each participant. The presentation of each stimulus interval was preceded by the 500-ms presentation of a white 7.7 × 7.7 arcmin fixation cross, surrounded by the zero disparity reference frame, and followed by a blank screen. On each trial, participants were asked to report the direction of motion of the front surface by using a mouse pointer to draw a line from the center of this blank screen, in the direction the front surface dots were traveling (see Figure 1c). These responses were encoded under assumption that the participant's choice was the direction of motion closest to the selected orientation. We placed one restriction on this encoding rule; only orientation choices within ±45° of either

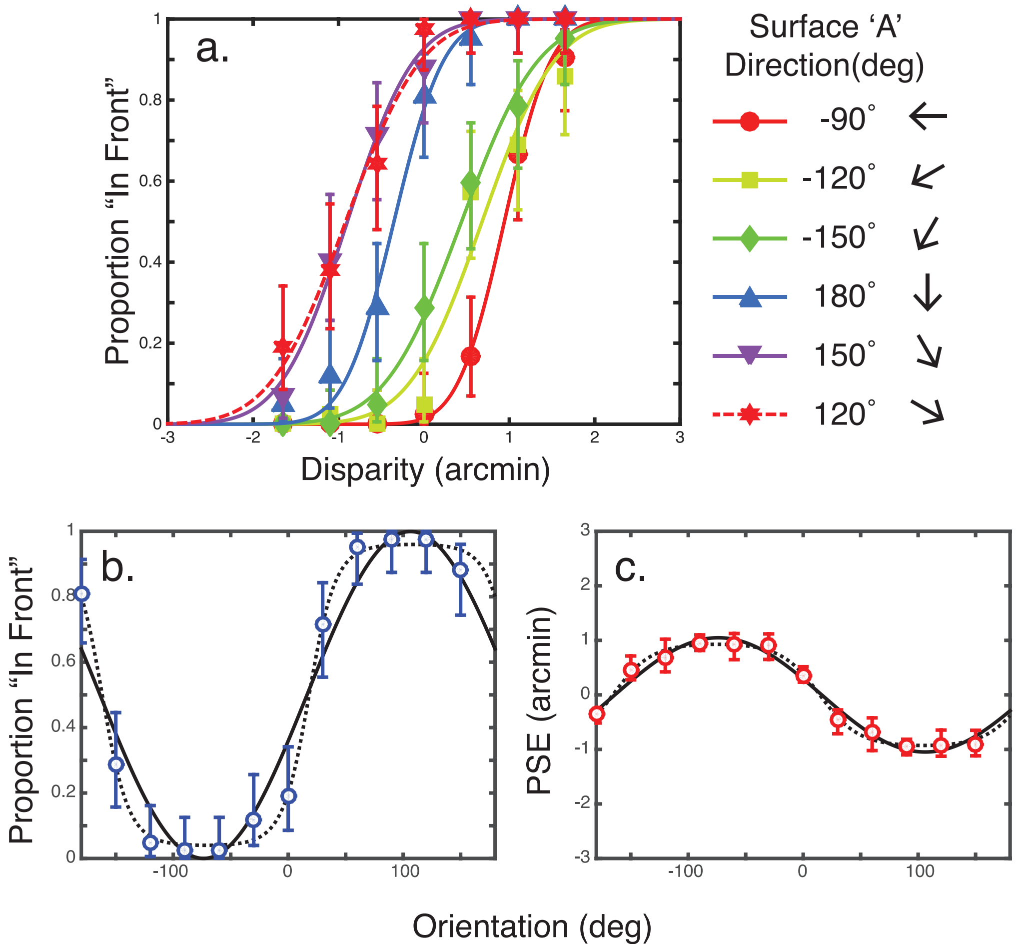

Figure 2. Results of Experiment 1 for an example participant. (a) Reported "in front" directions were used to obtain psychometric functions for the proportion of "Surface *A* in front" responses at each disparity and each direction of motion. Each psychometric function shows the fit for a single direction of motion (colors and symbols as per key). Error bars show binomial 95% CIs. (b) An example of the directional bias obtained for a single participant, using proportion "in front" responses when disparity was equal to zero. A cosine fit (Equation 1) is shown by the continuous black line, with the logit fit shown by the dashed black line. Error bars show binomial 95% CIs on the proportion of "in front" responses. The directional bias is defined as the peak of the fitted function. (c) Effects of direction on the location of the PSE, for the same participant, with fits obtained using the cosine fit in Equation 2, shown by the continuous black line, and the logit fit in Equation 3, shown by the dashed black line. Error bars here show bootstrapped 95% CIs on the location of the PSE. The directional bias is defined as the minimum of the fitted function.

presented direction of motion would be retained (Figure 1c). In practice no settings outside the $\pm 45°$ range were made, and thus the limiting rule did not need to be implemented.

Results and discussion

Psychophysical performance

Participants' choice of front surface on each trial was used to obtain measures of the proportion of trials in which each surface was perceived as being in front. For any given stimulus motion orientation, we defined each direction of motion as falling on one of two surfaces, *A* or *B*, and obtained the proportion of trials on which surface *A* was perceived as being in front. For manipulations of disparity, we defined positive values as those that moved surface *A* closer to, and surface *B* farther from, the observer, increasing the probability that surface *A* would be perceived as in front of surface *B*. Negative disparity values were those that moved surface *A* farther from, and surface *B* closer to, the observer, increasing the probability that surface *A* would be perceived as behind surface *B* (Figure 1d). This encoding of the data allowed for the plotting of six psychometric functions, describing the effects of disparity on perceived depth order for motion orientations from $-90°$ (surface *A* moving to the left), through 180° (surface A moving downward), to 120° (surface *A* moving down and to the right). These functions are shown, for an example observer, in Figure 2a. Note that functions for the complementary designation of surfaces A and B may be obtained by reversing the sign on the disparity axis and subtracting each proportion in front score from 1.

Two measures may be used to describe the scale and direction of any bias to perceive particular directions of motion as in front—the response bias for each direction of motion at zero disparity, and the disparity at the point of subjective equality (PSE) for each direction of motion. The PSE indicates the size and direction of disparity that must be added to counteract any direction of motion bias for depth ordering. These measures of directional bias are plotted for an example observer in Figure 2b and c. Direction-related changes in response bias, or PSE location, may be fitted with cosine functions defined by two free parameters a and b describing, respectively, the scale and direction of the bias. Functions describing these cosine relationships are given for zero disparity response biases in Equation 1, and for changes in PSE in Equation 2.

$$y = a\frac{\cos(x-b)}{2} + 0.5 \tag{1}$$

$$y = a\cos(x - [\pi + b]) \tag{2}$$

For zero disparity responses, $0 \leq a \leq 1$ and $-\pi \leq b \leq \pi$, where x defines the direction of motion in radians. For PSEs, $a \geq 0$. From these equations, the directional bias may be defined as the peak of the function for zero disparity responses, and as the function minimum for PSEs. The directional bias is thus the direction most likely to be perceived as in front for zero disparity responses, and, for PSEs, the direction for which the greatest opposing (i.e., negative) disparity must be added to counteract the directional bias.

In addition to these cosine fits of directional biases, logit function fits were also obtained, as in Mamassian and Wallace (2010) and Schütz (2014). For zero disparity responses, logit fits are defined by Equation 3:

$$y = c + a\frac{e^p}{1+e^p} \quad \text{where } p = d\cos(x-b). \tag{3}$$

Here a is a scaling parameter indicating the strength of the bias, b is the direction of the bias, c is a constant used to shift the curve on the y-axis, and d alters the shape of the curve, allowing for differences in observer sensitivity to small changes in motion orientation. Note that, unlike the cosine fit, the strength of directional bias for the logit fit is determined by a combination of a and d parameters. For PSEs b is replaced by $\pi + b$, as in Equation 2. Logit fits for zero disparity and PSE measures are shown, alongside cosine fits, in Figure 2b and c. R^2_{adj} values for logit fits were slightly higher than for cosine fits, with mean values of $R^2_{\text{adj}} = 0.86$ and 0.71, respectively, for zero disparity responses, and $R^2_{\text{adj}} =$ 0.78 and 0.75 for PSEs. Despite these improved fits, use of cosine fits has been maintained, as these provide a more readily interpretable measure of the strength of the biasing effect, via the single scaling parameter a (summarized in Table 1 for both zero disparity responses and PSEs).

Participant	$a_{(\text{Equation 1})}$	$a_{(\text{Equation 2})}$	a_{ori}	b_{ori}	n_{ori}
1	0.4	2.15	0.044	140.9	26.7
2	0.5	1.31	0.056	90.4	41.5
3	0.15	0.30	0.015	61.3	30.1
4	0.18	0.18	0.039	−0.9	12.5
5	0.18	1.59	0.017	124.0	26.4
6	0.16	0.45	0.047	−145.9	7.7
7	0.5	7.01	0.052	71.6	31.0
8	0.5	10.26	0.049	92.7	32.6
9	0.5	2.72	0.056	93.5	31.7
10	0.24	0.35	0.023	−109.8	27.5
11	0.43	0.56	0.051	87.3	30.0
12	0.34	0.89	0.042	−60.1	26.6
13	0.5	0.66	0.052	109.8	31.0
14	0.5	1.04	0.053	103.5	32.2
15	0.49	1.81	0.051	1.1	30.8
16	0.30	0.55	0.041	−167.7	27.4
17	0.19	0.58	0.019	136.9	26.6

Table 1. Scaling parameters a for cosine data fits, showing the strength of the directional bias effect in terms of zero disparity responses ($a_{(\text{Equation 1})}$) and PSEs ($a_{(\text{Equation 2})}$). *Notes*: Fits are also shown for the scaling (a_{ori}), biasing (b_{ori}), and reliability (n_{ori}) parameters of the beta distribution fit to zero disparity responses. Although values for a_{ori} are much smaller than for the other scaling parameters, this reflects the small shifts in the expected value of the beta distribution required to affect model responses.

The value of these scaling parameters varied from 0.15 to the theoretical maximum of 0.5 for zero disparity responses, and from 0.3 to 10.26 for PSEs. Higher values for the PSE scaling parameter reflect particularly strong directional biases for two participants (# 7 and 8), who showed ceiling and floor effects over some directions, but consistent effects of disparity on depth ordering at intermediate directions (i.e., directions between the extremes of their bias). The median scaling parameter for PSEs was 0.89, indicating that participants required surfaces to be separated by typical relative disparities of around 1.78 arcmin to counteract the effects of directional bias.

Plots of directional biases, defined by both zero disparity responses and PSEs, are shown in Figure 3, together with a plot of their correlation. We find that observers largely exhibit biases to see rightward motion as in front, although there are also notable biases in upward and down-left directions. That we find a strong correlation ($R^2_{\text{adj}} = 0.93$ for cosine fits and $R^2_{\text{adj}} = 0.94$ for logit fits) between the zero disparity and PSE measures of directional preference is unsurprising: If responses at zero disparity are biased, then PSEs must, by necessity, be nonzero. Instead, to gain a better understanding of

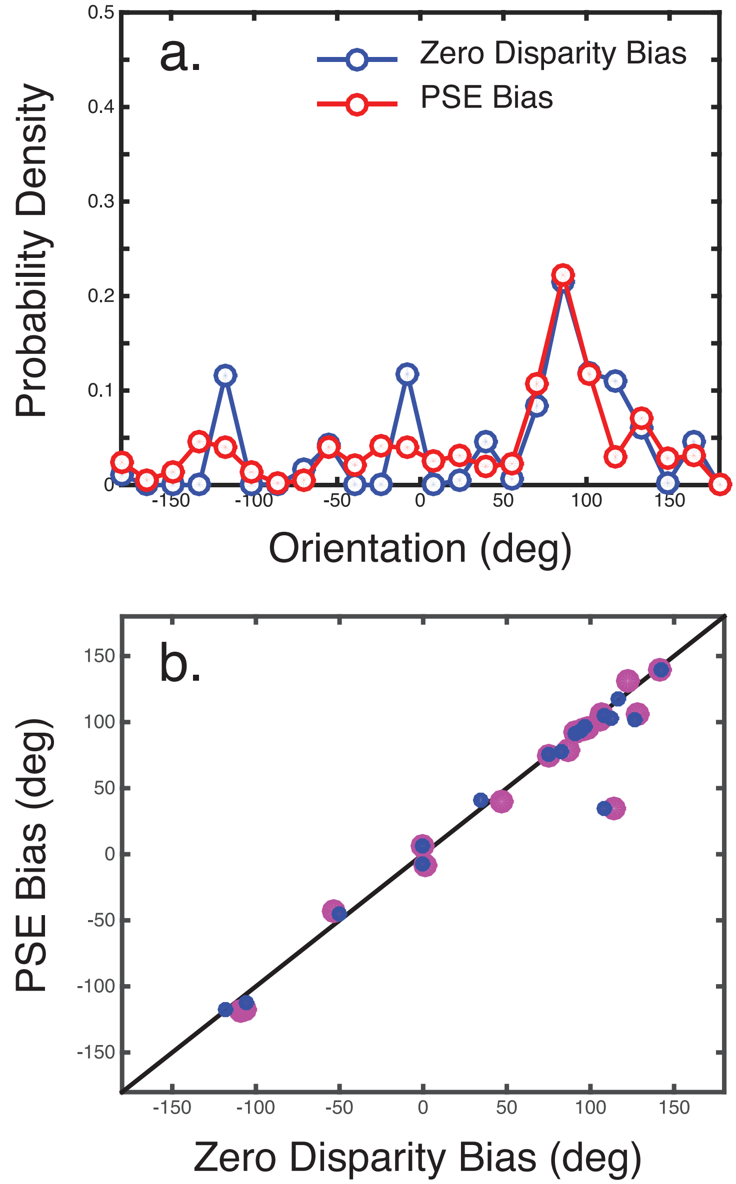

Figure 3. (a) Distribution of fitted directional biases for both zero disparity (blue) and PSE (red) measures. Distributions are derived from bootstrap resampling of responses in Experiment 1, with a total of 1,000 estimates of the location of the directional bias for each measure, obtained using fitted cosine functions. (b) Correlation of directional biases derived from zero disparity and PSE measures, for both cosine (magenta circles) and logit (blue circles) fits.

the interaction between disparity cues and directional preferences for depth ordering, one must examine the relationship between the strength of these two measures of bias. We turn to this issue in the following section.

Modeling depth order discrimination

To explore potential mechanisms underpinning the observed interactions of disparity and directional bias, we developed a model based on the Mixture of Bernoulli Experts (MBE) approach proposed by Backus (2009). In this approach, multiple cues for binary decisions are treated as beta distributions, defined over the interval [0, 1] by two parameters, θ and n. Under this definition, for any cue i, θ_i is the expected value of the distribution for that cue, and n_i is the assigned cue weight; as n_i increases in value, so too does cue reliability. The θ and n parameters together define the standard α and β parameters of the beta distribution where $\alpha = n\theta$ and $\beta = n(1-\theta)$. Note that, contrary to more typical weighted averaging models of signal integration (Landy, Maloney, Johnston, & Young, 1995), cue weights in the MBE models need not sum to one.

The MBE method is preferred here over standard linear averaging models, as it offers the benefit of defining cues in terms of their consistency with categorical responses, a property particularly suitable for the assessment of depth order. In other respects, the MBE approach provides predictions equivalent to the standard linear model: Interactions between cues in the MBE model are additive and are related to the relative reliability of each signal (Backus, 2009). We also compare the MBE model to a threshold model of depth ordering, which uses the same beta distribution formulation but, rather than integrating, selects between the disparity signal and the motion direction bias based on a fitted threshold level (see Supplementary Materials for full details of this approach). This threshold model acts as a minimal model for the role of directional bias; direction is treated as a response bias that is relied upon only in the absence of other cues to depth ordering.

Under the MBE approach, the beta distributions used to define each cue encode the probability that a given choice in a binary decision task accords with the true state of the world with a given regularity (i.e., the probability that a categorical response is successful). If we consider our example of depth ordering in motion transparency, then the beta distribution would encode the probability that, given some sensory signal (direction, disparity, or both) a given depth order decision accords with the true depth order of the stimulus with a given probability. To combine multiple sensory signals, α and β values are summed across each of N available cues, resulting in a posterior probability distribution $\pi(\theta)$, equivalent to a linear weighted average of available information sources (see Equation 4, or equation 1 from Backus, 2009).

$$\pi(\theta) = beta\left[\sum_{i=1}^{N} n_i\theta_i, \sum_{i=1}^{N} n_i(1-\theta_i)\right] \tag{4}$$

In addition to multiple signal cues, Backus (2009) included both a prior distribution and an unbiased decision noise term in his model. While the prior term is not essential in every case (and is not used here), the noise term allows for trial-to-trial response variations. Following Backus (2009) we model noise on each trial as a random, Gaussian-distributed value for θ, with $M = 0.5$ and $SD = 0.05$, and a fixed n value of $n = 10$. The posterior probability distribution $\pi(\theta)$ for the MBE model then becomes the sum of N available cues, including the random noise term. Here, we use the same deterministic decision rule as Backus (2009) to select between perceptual choices: the expected value of the posterior distribution $\hat{\theta}$ is calculated on each trial, where $\hat{\theta} = \alpha / (\alpha + \beta)$, and judged against the criterion of $\hat{\theta} \geq 0.5$.

To allow for the description of disparity and motion direction cues as beta distributions, changes in direction and disparity were restated as changes in θ. Changes in direction x were restated as sinusoidal changes in θ_{ori}, following Equation 1, fitting parameters a_{ori} and b_{ori} to define θ_{ori}. The parameter n_{ori} was also fitted to allow for a full definition of the beta distribution, conditional on orientation $\pi(\theta_{ori}|x)$ (Equations 5 and 6).

$$\pi(\theta_{ori}|x) = beta[n_{ori}\theta_{ori}, n_{ori}(1-\theta_{ori})] \tag{5}$$

where

$$\theta_{ori} = a_{ori}\frac{\cos[x - b_{ori}]}{2} + 0.5 \tag{6}$$

Together, these three parameters may be used to account for any bias in proportional responses when disparity is equal to zero. We fitted these three parameters for each participant, for zero disparity responses only, and used these fits for all subsequent analyses. Thus, while we used a model with three fitted parameters to account for each participant's directional bias, these parameters were fixed for all analyses of the interactions between directional bias and disparity. Bootstrapped 95% confidence intervals for R^2 values for these fits range from 0.88 to 0.93, with $M = 0.91$, indicating that the model of directional bias provides an excellent fit to the data (typically better than either cosine or logit fits). Note that, like the logit fit model, but unlike the standard cosine fit, fits derived from a beta-distributed representation of directional preference are able to account for variations in the extent to which small changes in orientation affect depth ordering. Summaries of the fitted parameters for the beta-distributed model of directional preference are provided in Table 1.

To investigate the interaction of disparity and directional bias in depth ordering, disparity values δ were restated as values of θ_{disp}, under the assumption that the two measures were linearly related, such that $\theta_{disp} = a_{disp}\delta + 0.5$. The full MBE model was then fit to each participant's data, using the previously fitted a_{ori}, b_{ori} and n_{ori} parameters for directional bias, together with free parameters a_{disp} and n_{disp}. Although free to vary between participants, a_{disp} and n_{disp} values were fixed across experimental conditions within participants. The results of this fitting process are summarized in Figure 4a through c. The full MBE model provided an excellent fit to participants' responses across all conditions and to measured PSEs. Bootstrapped 95% confidence intervals (CIs) for the R^2 values for these fits ranged from 0.948 to 0.963, with $M = 0.955$, for proportional responses, and from 0.898 to 0.968, with $M = 0.942$, for PSEs. Although performance was somewhat poorer on this measure, the full MBE model also provided a good fit to the slopes for each fitted psychometric function. R^2 values for this measure ranged from 0.493 to 0.802, with $M = 0.690$.

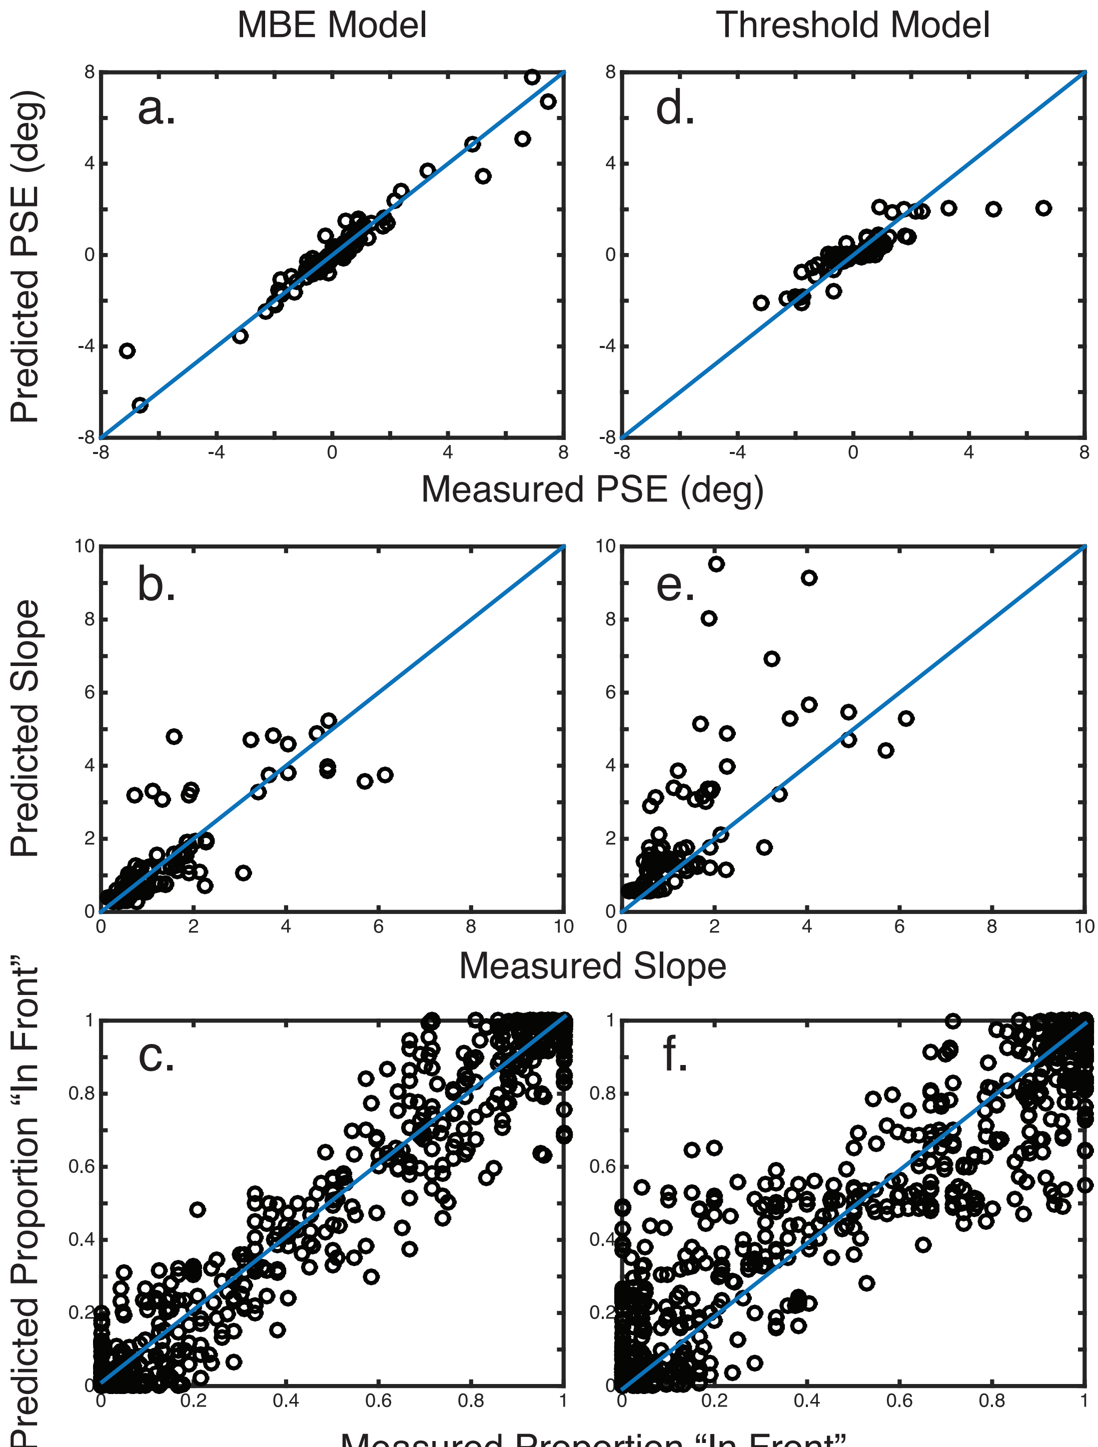

Figure 4. A comparison of the performance of the MBE (a–c) and threshold (d–f) models. (a, d) Plots of measured against predicted locations of the PSE. (b, e) Plots of measured against predicted psychometric function slopes. (c, f) Plots of measured against predicted proportion "in front" responses. In all cases, R^2 values for the MBE model are significantly higher than for the threshold model (see text).

The MBE model also provided a significantly better fit to the data than the threshold model (see Figure 4d through f, and Supplementary Materials). R^2 values for MBE model responses were, on average, 0.072 higher than for the threshold model (95% CIs ranged from 0.057 to 0.097). The threshold model also did a noticeably poorer job of fitting observer PSEs and slopes. R^2 values for fitted PSEs were, on average 0.628

higher for the MBE model (95% CIs ranged from 0.679 to 0.874), while values for fitted slopes were an average of 0.595 higher (95% CIs ranged from 0.311 to 0.783). As a further analysis, Akaike Information Criteria (AIC) and associated Akaike weights were calculated for each model (Akaike, 1974; Burnham & Anderson, 2002). AIC values were calculated using mean squared error values on the model fits, where $AIC = n \ln \sigma^2 + 2k$, k is the number of model parameters, σ^2 the mean squared error, and n the number of points at which model and observed responses are compared. Akaike weights for the MBE model approached 1, where weights were defined as $w_{MBE} = e^{(AIC_{min} - AIC_{MBE})/2} / \Sigma e^{(AIC_{min} - AIC_i)/2}$, and $\Sigma e^{(AIC_{min} - AIC_i)/2}$ indicates the sum of differences in AIC values across both models. This high weighting for the MBE model demonstrates the extent to which it provides an improved fit to the data, compared to the threshold model. The improved performance of the MBE model, relative to the threshold approach, suggests that participants' performance on the depth-ordering task reflects the action of mechanisms that integrate directional bias signals with other available cues to depth ordering.

Experiment 2: Motion direction affects surface segmentation in depth

The results of Experiment 1 established that the visual system integrates motion direction depth-order preferences with binocular disparity signals in a linear fashion. These results suggest that, just as binocular disparity may aid the segmentation of surfaces in motion transparency (Calabro & Vaina, 2006; Greenwood & Edwards, 2006a, 2006b; Hibbard & Bradshaw, 1999; Qian et al., 1994a), so too may motion signals aid, or hinder, the segmentation of surfaces in stereo transparency. We address this question directly in Experiment 2 by examining whether, when placed in conflict with disparity, directional depth-ordering biases affect thresholds for the discrimination of stereo transparency.

Methods

Participants

Five participants completed Experiment 2, including author RG. Each had previously completed Experiment 1, allowing for the measurement of directional preferences for depth ordering in motion transparency. With the exception of author RG, all participants were naive as to the purpose of the experiment, and had no explicit knowledge of their directional preference for depth ordering. Written consent was obtained for all participants and approval obtained from the local ethics board, in accordance with the guidelines of the British Psychological Society and the Declaration of Helsinki.

Stimulus and apparatus

Experiment 2 was completed using the same equipment set up as Experiment 1. Participants were presented with combined motion and stereo transparency stimuli, identical to those used in Experiment 1, with motion orientation constrained to fall along the axis of each participant's preferred orientation for depth ordering. With the addition of disparity, stimuli were made to move with depth order and direction either in accord with a participant's directional bias, or with depth ordering in opposition to that bias. Disparities up to ± 7 arcmin were added to each surface, resulting in relative disparities between front and back surfaces of up to 14 arcmin. Directional biases for the five participants were 3.6°, 75.4°, 95.6°, 104.1°, and 106.3°.

On each trial, a combined motion and stereo transparency stimulus was paired with a second stimulus, in a two-interval forced choice procedure. This additional stimulus had identical directions of motion, but had disparities drawn at random from a uniform distribution, with maximum and minimum set to match its transparency partner (see Figure 5). Dot motion in these stereoscopic volume stimuli was constrained such that all crossed disparity dots moved in the same direction, while all uncrossed disparity dots moved in the opposite direction. Disparity sign was opposition to dots moving in the motion and stereo transparency stimulus that moved in the same direction (e.g., crossed disparity dots moving rightward in the transparency condition were matched with uncrossed disparity dots moving rightward in the stereoscopic volume condition).

Design and procedure

The arrangement of paired moving stereo transparency and stereo volume stimuli, described above, resulted in two main experimental conditions. In the first, participants viewed a transparency stimulus where disparity-defined depth order was in agreement with the participant's directional bias for depth order (i.e., the direction preferentially seen in front had a crossed disparity, while the rear surface had an uncrossed disparity). In this preferred direction condition, the transparency stimulus was accompanied by a stereo volume stimulus where dots with uncrossed disparities moved in the direction preferentially seen as in front, while crossed disparity dots moved in the antipreferred direction.

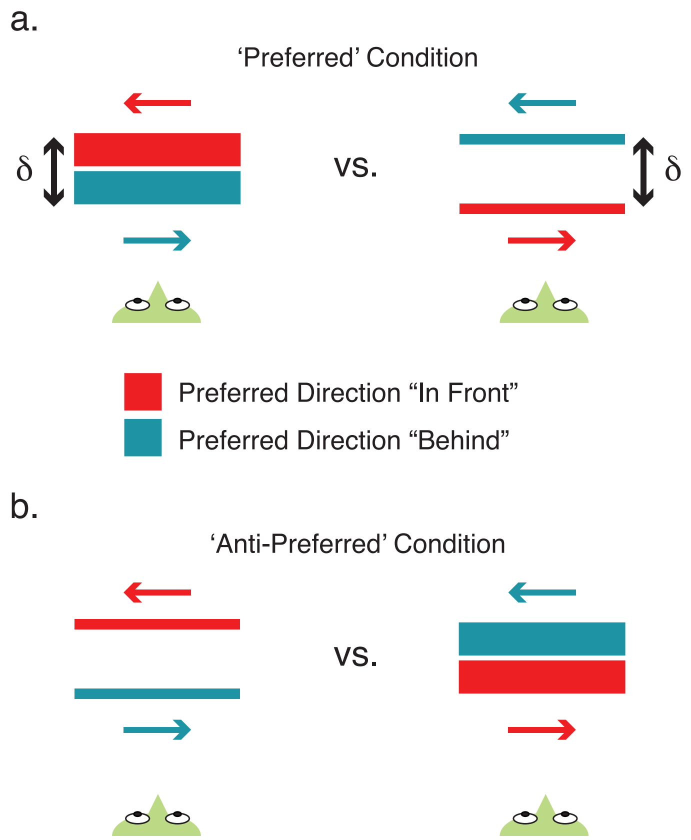

Figure 5. Illustration of the design of Experiment 2. (a) In the preferred condition, the transparent stimulus (top right) contains dots with crossed disparities moving in the direction preferentially reported as in front by each participant, and dots with uncrossed disparities moving in the opposite direction. The disparity volume stimulus (top left) contains the opposite arrangement of disparity and direction of motion. These arrangements are intended to encourage the processing of surface segmentation in the transparent interval, and inhibit segmentation for the volume stimulus (b) In the antipreferred condition, the arrangement of disparity and motion direction in each stimulus is reversed such that crossed disparities in the volume stimulus (bottom right) move in the preferred direction, with uncrossed disparities moving in the antipreferred direction. For the transparent stimulus (bottom left) crossed disparity dots move in the antipreferred direction, while uncrossed disparity dots move in the preferred direction. In this case surface segmentation processes are intended to be encouraged for the volume stimulus and inhibited for the transparent stimulus.

In the second antipreferred direction condition, disparity-defined depth order in the transparency stimulus was in conflict with the participant's directional bias for depth order (i.e., dots moving in the antipreferred direction had a crossed disparity, while those moving in the preferred direction had an uncrossed disparity). The accompanying stereo volume stimulus contained crossed disparity dots moving in the participant's preferred direction and uncrossed disparity dots moving in the opposite direction.

The logic of this experimental design is as follows: If directional depth ordering biases affect surface segmentation processes, then segmentation should be easier when surfaces move in their preferred directions relative to depth order, than when they move in the antipreferred direction. Any effect of directional depth order biasing should, therefore, mean that segmentation in the preferred direction condition occurs at smaller disparities than in the antipreferred direction condition, allowing transparent surfaces to be more easily discriminated from disparity-defined volumes (i.e., disparity thresholds should be lower in the preferred direction condition). In addition, the reversal of depth order and directional preference in these volumes are intended to further enhance any such effect. If segmentation is made more difficult when directional preference and depth order are in conflict, then discriminating between transparent surfaces and disparity-defined volumes should be easier in the preferred direction condition, compared to the antipreferred direction condition. Counter arrangements of directional preference and depth order within conditions should therefore encourage segmentation in one interval while inhibiting it in the other.

In each condition, stimulus pairs were presented, in random order, in a two-interval forced choice procedure. Participants' task was to select the interval containing the transparent planes. The disparity between front and back surfaces (maximum and minimum disparities in the stereo volume interval) was varied under a method of constant stimuli to obtain psychometric functions for the proportion of correctly identified transparent intervals. Participants were presented with nine levels of disparity for a minimum of 20 repeated trials, with the maximum disparity set for each participant after an initial practice block. Maximum disparities ranged from 2 to 7 arcmin. If directional depth order preferences have an effect on participants' ability to segment surfaces in depth, then thresholds for discriminating stereo transparency should be higher when disparity is in conflict with directional preference, compared to the condition where disparity and directional preference are in agreement.

Results and discussion

Disparity thresholds for the 75% correct identification of stereo transparency are shown, for each participant and each experimental condition, in Figure 6. Thresholds for the identification of stereo transparency were elevated in the antipreferred condition, compared to thresholds in the preferred direction condition. The effects of direction were such that no

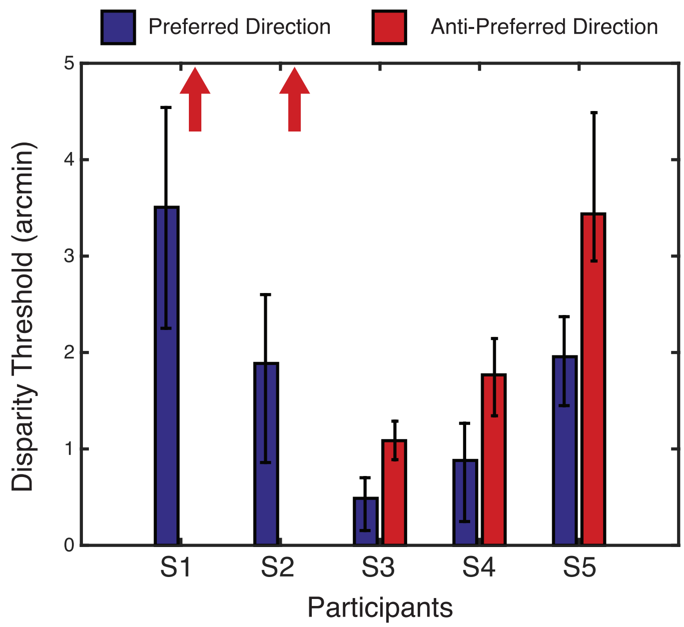

Figure 6. Results of Experiment 2. Graph shows disparity thresholds for the 75% correct identification of the transparent interval. Thresholds for the preferred condition are shown in blue, while red bars show thresholds for the antipreferred condition. Error bars show bootstrapped 95% CIs. Participants S1 and S2 were unable to reach threshold performance in the antipreferred condition. Thresholds for these participants are therefore displayed as arrows to indicate that they lie beyond measureable ranges.

antipreferred threshold could be fitted for two of the five participants. Responses for the antipreferred condition for these participants instead remained consistently at chance level. Bootstrapped 95% CIs for the difference in thresholds show that the disparity required for surface segmentation in the antipreferred condition was significantly higher for each remaining participant. CIs for antipreferred–preferred thresholds ranged from 0.29 to 0.94 arcmin, 0.26 to 1.57 arcmin, and 0.80 to 2.59 arcmin for these participants. These results show that the ability to segment moving surfaces in depth may be impaired simply through a reversal of disparity sign. This suggests that directional biases for depth ordering in motion transparency impact upon surface segmentation tasks that require the use of binocular disparity segmentation cues, even when motion cues are, in and of themselves, irrelevant.

General discussion

The experiments detailed in this article examined how the directional depth-ordering biases observed by Mamassian and Wallace (2010) interact with disparity cues to depth ordering in the perception of motion and stereo transparency. Our results show that this directional bias affects both the disparity required to determine depth order, and the disparity necessary for the segmentation of transparent surfaces in depth. Depth order in transparency is determined by the combination of disparity and motion direction signals, and the strength of this depth-order signal affects the processing of surface segmentation.

Our results both contrast with, and compliment, earlier findings from Moreno-Bote et al. (2008), who showed that wavelength-related effects in motion transparency depth ordering were sufficient to override depth ordering from disparity cues. While the data presented by Moreno-Bote et al. (2008) show that wavelength affects disparity-defined depth ordering, our results quantify the effects of directional bias, providing a measure of its strength in terms of the disparity required to counteract it. In addition, where Moreno-Bote et al. (2008) proposed that their measured effects of wavelength could be linked to statistical regularities of the natural environment, the idiosyncratic nature of directional biases (cf. Mamassian & Wallace, 2010; Schütz, 2011, 2014; Wexler et al., 2015) suggests that such signals do not possess a similar origin.

The quantification of the interaction between disparity cues and directional bias allowed us to model possible interactive processes underlying our results. This modeling used the MBE approach proposed by Backus (2009), and examined whether the observed effects of direction on depth ordering were consistent with a weighted linear integration of multiple cues. Our results are consistent with such processing, suggesting that the visual system integrates directional bias signals much as it would any standard cue to depth order. Below, we discuss a possible basis for this directional depth-ordering cue.

Direction of motion as a cue to depth order

Previous research has highlighted several potential cues to depth ordering in motion transparency, with such cues typically linked to statistical properties of the natural environment (Moreno-Bote et al., 2008; Schütz, 2011). The directional depth-ordering cue examined here, and previously reported by Mamassian and colleagues (Chopin & Mamassian, 2011; Mamassian & Wallace, 2010), does not seem to fit with this account. There seems little reason to suppose that a particular direction of motion is more likely associated with a given depth order in the world. The apparent idiosyncratic nature of directional bias also serves to undermine the possibility of a statistical basis of this kind. To account for this bias, one may instead consider how other cues to depth ordering may interact with perceived direction of motion, allowing for the possibility of directional cues to depth ordering being acquired through perceptual learning (Schütz, 2014; Wexler et al., 2015).

Backus and colleagues (Backus & Haijiang, 2007; Haijiang, Saunders, Stone, & Backus, 2006; Harrison & Backus, 2010) have provided demonstrations of several cases where the visual system acquired new associations that affected the perception of ambiguous visual stimuli. Similarly, other researchers (Adams, Graf, & Ernst, 2004) have shown that sensory integration can help recalibrate prior assumptions for the use of visual depth cues. The MBE approach proposed by Backus (2009), and used here, provides a suitable framework for the modeling of such perceptual learning, since it allows for the development of new cue associations, where signal and perceptual response are suitably correlated. Backus (2009) used the MBE model to show that new perceptual associations could be learned, under the assumption that different signals are conditionally independent (i.e., that a new signal depends on the world state, not on other available sensory signals). Such associative learning could account for the observed idiosyncratic directional biases, if initial responses to ambiguous motion transparency stimuli depend on noise in other known depth-ordering cues (e.g., relative element density). Initial responses could then be used to facilitate bootstrapped learning of any directional bias. Such a possibility is consistent with the findings of Harrison and Backus (2010), who found that ambiguous stimuli provided an effective input for perceptual learning. One possible means of testing this idea, other than replicating the methods of Harrison and Backus (2010), could be to examine whether motion adaptation, which is known to affect depth ordering in motion transparency (Schütz, 2011), could be used to engender a persistent directional bias that, as observed in the current study, interacts with other depth ordering signals.

Effects of depth ordering on surface segmentation

In Experiment 2, participants' thresholds for discriminating disparity-defined transparent surfaces from a disparity-defined volume were affected by the direction of motion of stimulus dots, relative to each participant's preferred direction for depth ordering. Discrimination thresholds were lower when dots on the transparent surface moved in accord with the depth order of participants' directional bias, compared to the case where the depth order of dots on the transparent surface was in opposition to this bias.

One possible mechanism to account for this effect would be for motion direction to play a quantitative role in the processing of depth from disparity. Under this explanation, the effects of directional bias on surface segmentation would arise due to directional signals increasing or decreasing the encoded separation between front and back surfaces. Such effects of ordinal depth cues on perceived quantitative depth have been previously reported. Burge, Peterson, and Palmer (2005) showed that figure-ground segmentation cues affect perceived depth from binocular disparity, with Burge, Fowlkes, and Banks (2010) suggesting that such effects may be based on the natural statistics of image contours. Our present results, however, provide no means of determining whether our directional biases in depth order display such quantitative effects.

An alternative possibility is that, rather than acting as quantitative cues to depth, motion direction and disparity may interact as cues to surface segmentation. In this case, the magnitude of the disparity signal is important only to the extent that it signals a separation of surfaces in depth, with motion direction, and its associated depth-ordering signal, important for a similar reason.

This surface organization account of the interaction of disparity and motion direction preference presents an intriguing possibility. In Experiment 2, participants' directional biases were used to place depth-ordering signals into conflict, reducing the strength of the depth-ordering signal in one stimulus, and increasing it in the other. The results of Experiment 2 therefore show that a manipulation of the strength of the depth-ordering signal affects the ability to segment surfaces in depth. If our results depend upon the interaction of disparity and motion direction signals as cues to surface organization, then the results of Experiment 2 suggest that the ordering of transparent surfaces in depth is a critical step in segmentation processing. Thus, rather than the assignment of depth order following a decision to segment signals, such an assignment may actually precede and inform segmentation processes. This idea is consistent with the importance of image features such as *T* and *X* junctions in image parsing (Adelson, 1993; Dresp, Durand, & Grossberg, 2002; Kawabe & Miura, 2006; Metelli, Da Pos, & Cavedon, 1985), but has not typically been considered in the processing of motion transparency, where explanations have concerned themselves with the integration of local motion signals differing in direction, orientation, and spatial frequency (Curran, Hibbard, & Johnston, 2007; Kanai, Paffen, Gerbino, & Verstraten, 2004; Qian et al., 1994b; Raudies & Neumann, 2010; Smith, Curran, & Braddick, 1999; Snowden & Verstraten, 1999). Interestingly, Schütz (2012) found that perceived numerosity in transparency stimuli is affected by depth order, but not by disparity magnitude. As with our findings, this suggests that the assignment of depth order is itself a critical processing step in the segmentation of multiple surfaces. Further research is required on the mechanisms for, and consequences of, depth order assignment.

Neurophysiological considerations

At the neural level, cells in area MT have been shown to respond to specific combinations of motion direction and disparity in rotating cylinder stimuli (Bradley, Chang, & Andersen, 1998; Parker, Krug, & Cumming, 2002). In addition, joint motion/disparity-tuned neurons also show responses that are correlated with task-related perceptual decisions, even when the direction of rotation of a stimulus is ambiguous (i.e., when the stimulus contains no nonzero disparities; Bradley et al., 1998). More recent electrical stimulation studies have also shown that these neurons are directly involved in determining perceptual decisions with such ambiguous stimuli (Krug, Cicmil, Parker, & Cumming, 2013). Since neurons of this kind show increased response with increasing disparity, in principle the strength of their responses could also be affected by other cues that appear to alter the strength of the depth-ordering signal, such as the motion direction cue studied here. Increased responses from these cells could operate either to increase perceived depth, or to increase the strength of any surface segmentation signal.

Conclusions

The experiments reported in this article have provided a demonstration of the interaction between an idiosyncratic motion direction bias on depth ordering and binocular disparity cues to depth. These signals interact in tasks involving both the assignment of depth order and the segmentation of surfaces in depth. These results indicate that, despite apparently lacking an environmental basis, directional depth-order biases are treated by the visual system in the same way as more readily identifiable cues to depth, and suggest that depth-ordering processes may play a key role in establishing surface segmentation.

Keywords: motion transparency, stereo transparency, cue integration, individual differences, segmentation

Acknowledgments

Commercial relationships: none.
Corresponding author: Ross Goutcher.
Email: ross.goutcher@stir.ac.uk.
Address: Faculty of Natural Sciences, Psychology Division, University of Stirling, Stirling, Scotland, UK.

References

Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the "light-from-above" prior. *Nature Neuroscience*, *7*(10), 1057–1058.

Adelson, E. H. (1993). Perceptual organisation and the judgment of brightness. *Science*, *262*, 2042–2044.

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

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control AC*, *19*, 716–723.

Backus, B. T. (2009). The mixture of Bernoulli experts: A theory to quantify reliance on cues in dichotomous perceptual decisions. *Journal of Vision*, *9*(1):6, 1–19, doi:10.1167/9.1.6. [PubMed] [Article]

Backus, B. T., & Haijiang, Q. (2007). Competition between newly recruited and pre-existing visual cues during the construction of visual appearance. *Vision Research*, *47*, 919–924.

Bradley, D. C., Chang, G. C., & Andersen, R. A. (1998). Encoding of three-dimensional structure-from-motion by primate area MT neurons. *Nature*, *392*, 714–717.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.

Burge, J., Fowlkes, C. C., & Banks, M. S. (2010). Natural scene statistics predict how the figure-ground cue of convexity affects human depth perception. *Journal of Neuroscience*, *30*(21), 7269–7280.

Burge, J., Peterson, M. A., & Palmer, S. E. (2005). Ordinal configural cues combine with metric disparity in depth perception. *Journal of Vision*, *5*(6):5, 534–542, doi:10.1167/5.6.5. [PubMed] [Article]

Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodal inference: A practical information-theoretic approach*. New York: Springer-Verlag.

Calabro, F. J., & Vaina, L. M. (2006). Stereo motion transparency processing implements an ecological smoothness constraint. *Perception*, *35*, 1219–1232.

Cavanagh, P., & Mather, G. (1989). Motion: The long and the short of it. *Spatial Vision*, *4*, 103–129.

Chopin, A., & Mamassian, P. (2011). Usefulness influences visual appearance in motion transparency depth rivalry. *Journal of Vision*, *11*(7):18, 1–8, doi:10.1167/11.7.18. [PubMed] [Article]

Curran, W., Hibbard, P. B., & Johnston, A. (2007). The visual processing of motion-defined transparency. *Proceedings of the Royal Society B: Biological Sciences*, *274*, 1049–1056.

Dresp, B., Durand, S., & Grossberg, S. (2002). Depth perception from pairs of overlapping cues in pictorial displays. *Spatial Vision*, *15*, 255–276.

Goutcher, R., & Loffler, G. (2009). Motion transparency from opposing luminance and contrast modulated gratings. *Vision Research*, *49*, 660–670.

Greenwood, J. A., & Edwards, M. (2006a). Pushing the limits of transparent-motion detection with binocular disparity. *Vision Research*, *46*, 2615–2624.

Greenwood, J. A., & Edwards, M. (2006b). An extension of the transparent-motion detection limit using speed-tuned global-motion systems. *Vision Research*, *46*, 1440–1449.

Haijiang, Q., Saunders, J. A., Stone, R. W., & Backus, B. T. (2006). Demonstration of cue recruitment: Change in visual appearance by means of Pavlovian conditioning. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(2), 483–488.

Harrison, S. J., & Backus, B. T. (2010). Uninformative visual experience establishes long term perceptual bias. *Vision Research*, *50*, 1905–1911.

Hibbard, P. B., & Bradshaw, M. F. (1999). Does binocular disparity facilitate the detection of transparent motion? *Perception*, *28*, 183–191.

Kanai, R., Paffen, C. L. E., Gerbino, W., & Verstraten, F. A. J. (2004). Blindness to inconsistent local signals in motion transparency from oscillating dots. *Vision Research*, *44*, 2207–2212.

Kawabe, T., & Miura, K. (2006). Recovering depth-order from orientation-defined junctions. *Psychological Research*, *70*, 375–383.

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

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *Perception*, 36, ECVP Abstract Supplement 14.

Krug, K. Cicmil, N., Parker, A. J., & Cumming, B. G. (2013). A causal role for V5/MT neurons coding motion-disparity conjunctions in resolving perceptual ambiguity. *Current Biology*, *23*, 1454–1459.

Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, *35*, 389–412.

Levinson, E., & Sekuler, R. (1975). The independence of channels in human vision selective for direction of movement. *Journal of Physiology*, *250*, 347–366.

Mamassian, P., & Wallace, J. M. (2010). Sustained directional biases in motion transparency. *Journal of Vision*, *10*(13):23, 1–12, doi:10.1167/10.13.23. [PubMed] [Article]

Metelli, F., Da Pos, O., & Cavedon, A. (1985). Balanced and unbalanced: Complete and partial transparency. *Perception & Psychophysics*, *38*, 354–366.

Moreno-Bote, R., Shpiro, A., Rinzel, J., & Rubin, N. (2008). Bi-stable depth ordering of superimposed moving gratings. *Journal of Vision*, *8*(7):20, 1–13, doi:10.1167/8.7.20. [PubMed] [Article]

Parker, A. J., Krug, K., & Cumming, B. G. (2002). Neuronal activity and its links with the perception of multi-stable figures. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *357*, 1053–1062.

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

Qian, N., Andersen, R. A., & Adelson, E. H. (1994a). Transparent motion perception as detection of unbalanced motion signals. I. Psychophysics. *Journal of Neuroscience*, *14*(12), 7357–7366.

Qian, N., Andersen, R. A., & Adelson, E. H. (1994b). Transparent motion perception as detection of unbalanced motion signals. III. Modeling. *Journal of Neuroscience*, *14*(12), 7381–7392.

Raudies, F., & Neumann, H. (2010). A model of neural mechanisms in monocular transparent motion perception. *Journal of Physiology – Paris*, *104*, 71–83.

Schütz, A. C. (2011). Motion transparency: Depth ordering and smooth pursuit eye movements. *Journal of Vision*, *11*(14):21, 1–19, doi:10.1167/11.14.21. [PubMed] [Article]

Schütz, A. C. (2012). There's more behind it: Perceived depth order biases perceived numerosity/density. *Journal of Vision*, *12*(12):9, 1–16, doi:10.1167/12.12.9. [PubMed] [Article]

Schütz, A. C. (2014). Interindividual differences in preferred directions of perceptual and motor decisions. *Journal of Vision*, *14*(12):16, 1–17, doi:10.1167/14.12.16. [PubMed] [Article]

Smith, A. T., Curran, W., & Braddick, O. J. (1999). What motion distributions yield global transparency and spatial segmentation? *Vision Research*, *39*, 1121–1132.

Snowden, R. J., & Verstraten, F. A. J. (1999). Motion transparency: Making models of motion perception transparent. *Trends in Cognitive Sciences*, *3*(10), 369–377.

Tsirlin, I., Alison, R. S., & Wilcox, L. M. (2008). Stereoscopic transparency: Constraints on the perception of multiple surfaces. *Journal of Vision*, *8*(5):5, 1–10, doi:10.1167/8.5.5. [PubMed] [Article]

van Dam, L. C. J., & Ernst, M. O. (2010). Preexposure disrupts learning of location-contingent perceptual biases for ambiguous stimuli. *Journal of Vision*, *10*(8):15, 1–17, doi:10.1167/10.8.15. [PubMed] [Article]

Wexler, M., Duyck, M., & Mamassian, P. (2015). Persistent states in vision break universality and time invariance. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(48), 14990–14995.

Wilson, H. R., & Kim, J. (1994). A model for two-dimensional motion perception: Coherence and transparency. *Visual Neuroscience*, *11*, 1205–1220.