Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research 45 (2005) 469–483
Vision Research
www.elsevier.com/locate/visres

# Selective biasing of stereo correspondence in an ambiguous stereogram

Ross Goutcher *, Pascal Mamassian

*Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8QB, UK*

Received 27 November 2003; received in revised form 26 July 2004

**Abstract**

In spite of numerous studies in stereoscopic perception, it is still not clear how the visual system matches features between the two eyes. One reason is that these previous studies used stimuli that presented little perceptual ambiguity, so the correspondence problem had only one solution. We present here a novel stimulus that presents a more complex correspondence problem. This stimulus is inspired by "wallpaper" stimuli and was specifically designed to put into conflict two possible constraints underlying stereo correspondence matching. These constraints are the nearest neighbour matching rule—that biases surfaces towards the horopter—and the nearest disparity rule—that biases surfaces to be smooth. By varying the contrast of adjacent image features in this stimulus, we were able to reveal and quantify a preference for nearest disparity matching. The magnitude of this preference is dependent upon the magnitude of possible disparities in the scene and is consistent with the idea that the visual system seeks to minimise local differences in disparity. We discuss these results with regard to the use of prior constraints in models of stereo matching.


*Keywords:* Stereopsis; Correspondence problem; Prior constraints; Wallpaper illusion

## 1. Introduction

The perception of depth from binocular disparity depends upon the correct matching of corresponding features between the left and right eyes' images. In complex scenes the visual system may be confronted with multiple candidate features for matching and must reduce the number of possible correspondences in order to attain a stable, unified representation of the scene. The resolution of this correspondence problem for stereo vision has been a topic of near constant interest for researchers in the 40 years since Julesz' popularisation of the random dot stereogram (Julesz, 1964). Many computational models of the correspondence matching process have been proposed (e.g. Jones & Malik, 1992; Marr & Poggio, 1976, 1979; Pollard, Mayhew, & Frisby, 1985; Prazdny, 1985; Prince & Eagle, 2000; Qian & Zhu, 1997; Read, 2002a, 2002b; Sato & Yano, 2000; Tsai & Victor, 2003). To resolve the correspondence problem, such models must limit possible matches with a series of constraints or rules. Models often differ in the constraints they use and the extent to which these are employed in an explicit (e.g. Marr & Poggio, 1976, 1979; Pollard et al., 1985) or implicit (e.g. Prince & Eagle, 2000; Qian & Zhu, 1997; Read, 2002a, 2002b) manner. Constraints on matching include feature similarity, matching to the nearest neighbour or nearest disparity, and considering only epipolar matches (for an extensive review of proposed matching rules, see Howard & Rogers, 2002).

In this paper, we concentrate on the visual system's adherence to the solutions provided by nearest neighbour, nearest disparity and contrast similarity matching rules. *Nearest neighbour* matches (Arditi, Kaufman, &

* Corresponding author.
*E-mail address:* ross@psy.gla.ac.uk (R. Goutcher).

0042-6989/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2004.08.025

Movshon, 1981) minimise the absolute disparity of image features. That is, they select the correspondence solution that places the image feature closest to the horopter. In contrast, the *nearest disparity* rule (Marr & Poggio, 1976, 1979; McKee & Mitchison, 1988; Mitchison & McKee, 1987a, 1987b) minimises the relative disparity of image features, giving the correspondence solution that minimises the difference in disparity between nearby points. As such, the nearest disparity rule has been thought of as a 'smoothness' constraint and is often referred to as a continuity or cohesiveness constraint. The contrast similarity rule is one of a series of constraints—including also, orientation similarity—concerned with feature similarities. Under this constraint, matches are made between features of maximally similar contrast (Smallman & McKee, 1995).

Despite the suggestion of so many constraints in the literature, very little research has been conducted to examine the competition between matching rules. There is precious little empirical data showing which solution the visual system adheres to when confronted with multiple plausible matches (i.e. multiple matches that satisfy one or more matching constraint). Zhang, Edwards, and Schor (2001) recently investigated this issue. Using a periodic stimulus consisting of a one-dimensional luminance Gabor flanked, above and below, by two similar Gabors, they found that matching tends towards the solution that minimises the disparity between adjacent surfaces; that is, the solution that minimises relative disparity. This finding was particularly interesting since their stimulus put nearest neighbour and nearest disparity matching rules into conflict. Their experiments thus suggest that the process of correspondence matching is concerned more with finding solutions that satisfy the nearest disparity constraint than those that satisfy the nearest neighbour constraint.

One important characteristic of the study of Zhang et al. (2001) is that their stimuli consisted of three isolated objects rather than a single surface, so their result can be interpreted as a contextual effect. Furthermore, the stimuli used by Zhang et al. (2001) contain a potential confound between local changes in disparity and the total change in disparity across the scene. We here define this maximum change in disparity across the scene as the *global relative disparity*. This distinction between local and global relative disparity is clearer if one considers the relative disparities that arise, at a global and local level, with different stimuli. Consider the stimuli depicted in Fig. 1. Fig. 1a illustrates a single, fronto-parallel surface in depth, with two local areas—$x$ and $y$—highlighted. At area $x$ relative disparity is zero, since all points are at the same depth. However, at area $y$, relative disparity is determined by the difference in disparity between the stimulus and a zero disparity surround. Thus, although relative disparity is zero across much of the image, the relative disparity across the entire image—the *global relative disparity*—is determined by those few areas containing a difference in disparity between stimulus and surround. Figs. 1b–d illustrate increasingly complex stimuli, where the presence of local areas with zero relative disparity is increasingly scarce. In such stimuli the global relative disparity is determined by the largest change in disparity across the entire scene. For example, in the case of the squarewave illustrated in Fig. 1b, the global relative disparity is the peak-to-trough disparity of the waveform.

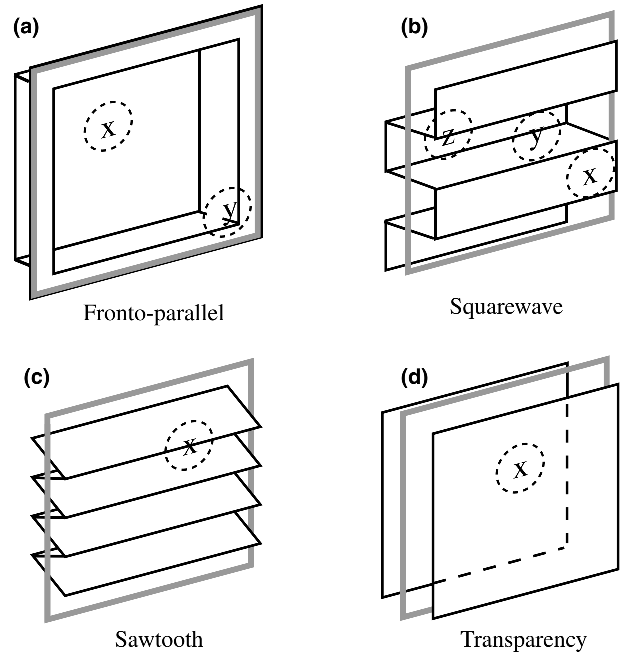


Fig. 1. Distinction between locally and globally defined relative disparities. All four figures (a–d) have identical global relative disparity—defined as the largest change in disparity across the entire scene—but different local relative disparity structure. (a) Illustration of a fronto-parallel surface located behind a frame. Highlighted areas $x$ and $y$ show points where relative disparity is zero ($x$) and non-zero ($y$). Global relative disparity is determined by the disparity between the surface and the surround ($y$). (b) Illustration of a squarewave modulation in depth. The squarewave contains many areas where relative disparity is zero ($x$ and $y$), though fewer than (a). Global relative disparity is determined by the peak-to-trough relative disparity of the waveform ($z$). (c) Illustration of a sawtooth modulation in depth, which contains no areas with a relative disparity of zero (e.g. area $x$). Global relative disparity is determined by the disparity at the sharp depth transitions. (d) Illustration of two overlapping transparent surfaces in depth. Local relative disparity is never zero since both surfaces are present within any local area ($x$). Global relative disparity is determined by the disparity between front and back surfaces.

Readers should note that a stimulus with a small global relative disparity may, locally, have a great deal of variation in disparity. Consider the transparent surfaces depicted in Fig. 1d. There are no local areas containing zero relative disparity in such a stimulus since, over a local area, both surfaces are visible. However, the global relative disparity—determined by the disparity between front and back surfaces—may be small if the separation between the surfaces is small. Conversely, the global relative disparity of a single fronto-parallel surface may be large if the disparity between stimulus and surround is large.

In the stimulus used by Zhang et al. (2001) participants' responses could have been driven by a preference to minimise either local or global relative disparity. In this paper, we detail a novel ambiguous stimulus that puts nearest neighbour and nearest disparity constraints into conflict. By varying the disparity of surfaces arising from the nearest disparity or nearest neighbour match, the global relative disparity of both matches may be equated. This allows us to investigate whether the nearest disparity bias reported by Zhang et al. (2001) is due to a locally or globally implemented nearest disparity constraint.

In addition to the manipulation of disparity, manipulating the contrast similarity of image features in the stimulus can selectively bias one of the two percepts (see also Anderson & Nakayama, 1994 and Smallman & McKee, 1995 for similar uses of contrast to bias matching). Such selective biasing may be used to examine the relative strength of the solutions provided by distinct matching constraints. To pre-empt our results, we confirm and quantify the preference for nearest disparity matching found by Zhang et al. (2001) and show that the tendency to match to the nearest disparity is affected by the depth resulting from nearest neighbour and nearest disparity solutions. These results demonstrate a locally implemented bias for the nearest disparity, rather than a bias concerned with minimising the global relative disparity. We further suggest that the stimulus and method detailed here could be used to reveal the impact of numerous factors on stereo matching.

Our stimulus makes use of the stereo 'wallpaper illusion' (Brewster, 1844) and takes the form of a modified periodic 'wallpaper' dot stereogram. Wallpaper stimuli repeat an identical pattern multiple times from left to right, thereby increasing the ambiguity of the correspondence between features of the left and right eyes. Assuming no strong idiosyncratic preference for crossed or uncrossed disparities, the typical wallpaper stimulus is resolved to elicit the perception of a single plane either in front of, or behind fixation. Such percepts correspond to the solutions given by both the nearest disparity and nearest neighbour constraints. By adding small, systematic offsets to image features we may put these constraints into conflict. When in conflict, two qualitatively different percepts result. The first percept consists of a single opaque fronto-parallel surface and the second consists of two semi-transparent fronto-parallel surfaces. The perception of a single, fronto-parallel surface indicates adherence to a nearest disparity constraint, since, in such a stimulus, relative disparity is zero across almost the entire image. Conversely, the perception of stereo transparency corresponds to the nearest neighbour match, since, in the stimulus, such surfaces lie closer to the point of fixation than the single surface match. Note that the global relative disparity of these two percepts may be equated if the relative disparity between the single surface and fixation is the same as the relative disparity between front and rear transparent surfaces. Fig. 2a and b provide examples of similar stimuli that lead to such different percepts.

The dominance of either the nearest disparity or the nearest neighbour constraint can thus be studied by monitoring the perceptual preference for a single opaque or two transparent surfaces. In an attempt to quantify this dominance further, we also manipulate the similarity of the features across eyes in order to create a stimulus that will equally likely be perceived as opaque or transparent. In experiment 1, we describe the construction of our stimuli in detail, and provide results from the manipulation of the contrast similarity of the image features for two stimulus durations. In experiments 2 and 3, we detail the effects of varying disparity on the resolution of our stimulus and discuss the implications of these findings for models of stereo correspondence matching.

## 2. Experiment 1

In this experiment, we seek to assess the bias for nearest disparity matching found by Zhang et al. (2001). By varying the contrast between image features in an ambiguous stereogram we bias matching between conflicting nearest neighbour and nearest disparity solutions. We assess the bias for nearest disparity matching for two presentation durations. In experiment 1a stimuli were presented for 2 s in a raised temporal cosine window. In experiment 1b stimuli were presented for 80 ms in a square-wave temporal window. Short presentation durations were used to discount any potentially confounding effects of vergence eye movements.

### 2.1. Stimuli

#### 2.1.1. Structure of the stimulus

We first describe the wallpaper pattern that forms the basis of our stimulus, and then the specific way in which this pattern is modified. The wallpaper pattern consisted of 16 rows of about 14 dots within an area of $7.48° \times 7.48°$, at a viewing distance of 80 cm (i.e. dot density was 4.0 dots per degree squared). Each dot was a small circular Gaussian blob measuring $10.5'$. The vertical separation between two rows was $27.1'$ and the horizontal distance between two consecutive dots was $\alpha = 30.1'$ (see Fig. 3a). In order to avoid a regular grid, the first dot of each row was randomly shifted horizontally (by an amount $\alpha_0$ between 0 and $\alpha$). The vertical edges of

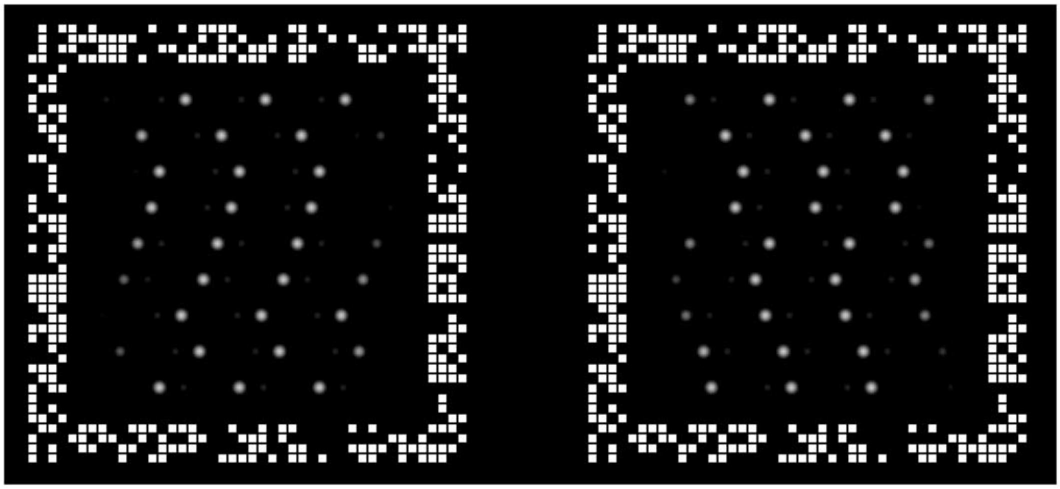

(a)

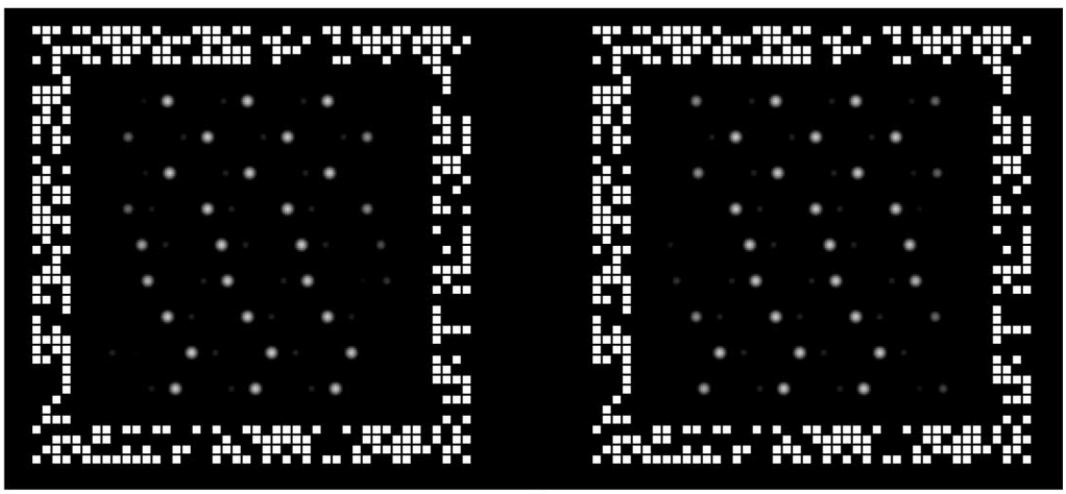

(b)

Fig. 2. An example of the ambiguous stereogram stimuli used in the experiment. The upper stereogram (a) shows the experimental stimulus with contrast ambiguity biased towards the transparent—nearest neighbour—percept. The lower stereogram (b) shows the stimulus with a single surface—nearest disparity—biased contrast ambiguity level. Both examples show the stimulus with the same dot spacing as in the experiment, but with half the number of dots (i.e. stimulus is half the size of that shown to participants, with the same dot density).

this pattern were then smoothed to diminish the visibility of unmatched dots between left and right half-images. The smoothing was obtained by multiplying the pattern with a rectangular window attached to two half-Gaussian distributions on either side (the centres of the Gaussians were at 1.87° and 5.61° and the standard deviation was 0.63°). Note that this smoothing operation had an effect only on the left and right sides of the stimulus, leaving the luminance constant across rows of dots. When fused, this wallpaper stimulus would lead to a single fronto-parallel plane either in front of or behind fixation.

The basic wallpaper stimulus can now be modified to also allow transparent percepts. Within each row, dots were shifted alternately by an amount $+\delta$ or $-\delta$ in one half-image, and by $-\delta$ or $+\delta$ in the other half-image (here, $\delta$ was set to $6.02'$: see Fig. 3b). Finally, the left half-image was shifted by $\varepsilon$ and the right half-image by $-\varepsilon$ to give the stimulus a pedestal disparity of $-2\varepsilon$ ($\varepsilon$ was set to $3.01'$: see Fig. 3c).

### 2.1.2. Dominant percepts

As with any stereogram, there are multiple possible interpretations of the adapted wallpaper stimulus shown in Fig. 2. However, only two percepts are dominant here. The first of these corresponds to the nearest neighbour solution while the second corresponds to the nearest disparity solution. We describe these two types of percept in turn.

When the dots are matched according to the nearest neighbour constraint, dot A in Fig. 3c will match with dot C, and dot B with dot D. From that figure, one can readily see that these two pairs of dots will lead to two different disparities, consistent with two

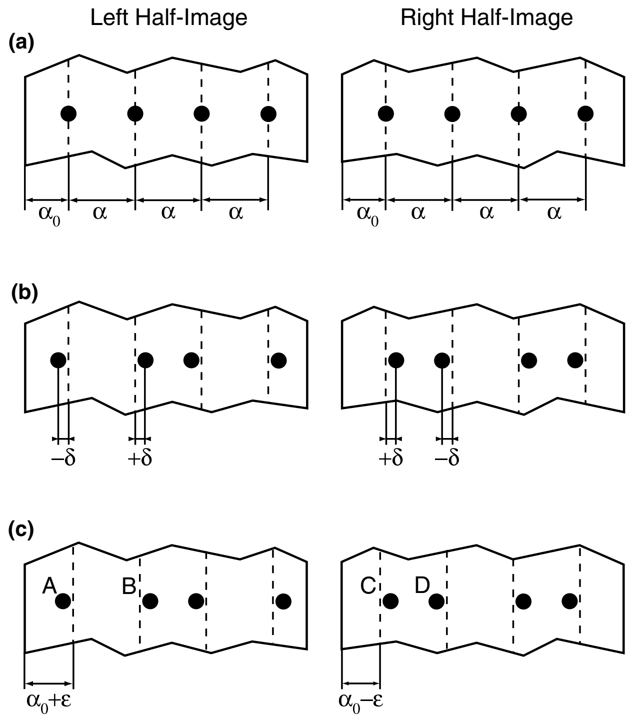


Fig. 3. Illustration of the steps involved in the construction of the stimulus. (a) The basis of our stimulus is a wallpaper pattern where dots along a particular row are equally spaced. (b) Small positive and negative displacements are then introduced in alternate dots, and in the opposite direction in the other half-image. (c) Finally, a pedestal disparity is introduced, shared between the left and right images. Dot A in the left image can be matched to multiple dots in the right image, including dots C and D.

fronto-parallel surfaces overlaid in depth. These two disparities $d_1$ and $d_2$ are given by

$$\begin{cases} d_1 = (\alpha_0 + \delta - \varepsilon) - (\alpha_0 - \delta + \varepsilon) = 2\delta - 2\varepsilon \\ d_2 = (\alpha_0 + \alpha - \delta - \varepsilon) - (\alpha_0 + \alpha + \delta + \varepsilon) = -2\delta - 2\varepsilon \end{cases} \tag{1}$$

In other words, the transparent percept will be such that the rear surface has an uncrossed disparity $d_1$ (equal to $+6.02'$) and the front surface has a crossed disparity $d_2$ (equal to $-18.06'$).

The nearest disparity solution is found by making a next-to-nearest neighbour match. This solution results in the perception of a single fronto-parallel plane whose disparity is either crossed or uncrossed. Uncrossed disparity will be obtained if dot A in Fig. 3c is matched with dot D, while crossed disparity will be obtained when dot B is matched with dot C. These two disparities $d_3$ and $d_4$ are given by

$$\begin{cases} d_3 = (\alpha_0 + \alpha - \delta - \varepsilon) - (\alpha_0 - \delta + \varepsilon) = \alpha - 2\varepsilon \\ d_4 = (\alpha_0 + \delta - \varepsilon) - (\alpha_0 + \alpha + \delta + \varepsilon) = -\alpha - 2\varepsilon \end{cases} \tag{2}$$

In absolute values, disparity $d_3$ (equal to $+24.08'$) is much smaller than disparity $d_4$ (equal to $-36.12'$), so the uncrossed disparity solution will be preferred and the single surface will be perceived behind the fixation plane.

Note that the values of $\alpha$, $\delta$ and $\varepsilon$ were chosen such that the global relative disparity between the front and back surfaces in the transparent percept (i.e. $4\delta = 24.08'$) equals the disparity of the surface in the single surface percept (i.e. $d_3 = \alpha - 2\varepsilon$). Thus, only a matching procedure concerned with local relative disparities—rather than the global relative disparity in the scene—may underlie any observed nearest disparity bias.

### 2.1.3. Manipulation of contrast

So far, because all the dots are identical, the stimulus we have generated is truly ambiguous and will lead to a transparent percept if the nearest neighbour constraint is adopted, and to a single surface if the nearest disparity constraint is adopted instead. In this paper, we look beyond such classification of perceptions and attempt to selectively bias matching towards the percept given by one or the other constraint. The biasing of ambiguous stereograms has previously been used to demonstrate the importance of half-occlusions (Anderson & Nakayama, 1994) and fixation depth (McKee & Mitchison, 1988) on disparity computation. In general, correspondence biasing is used to manipulate the perception of an ambiguous stereogram and thus demonstrate the importance of the biasing factor. However, in this paper such biasing is used as a probe to measure the relative preference for the matching solutions underlying two distinct percepts. Here, biasing was obtained by manipulating the contrast of dot pairs as detailed below. Inter-ocular contrast is known to greatly affect stereopsis both in terms of stereoacuity (e.g. Halpern & Blake, 1988; Legge & Gu, 1989) and correspondence matching (Anderson & Nakayama, 1994; Smallman & McKee, 1995).

To illustrate the contrast manipulation procedure, let us assume that dots A, B, C and D in Fig. 3c have luminances $L_A$, $L_B$, $L_C$ and $L_D$ that can vary between 0 and 1 (the luminances will then be scaled by the maximum luminance of the display, i.e. 19.2 cd/m$^2$). If we want to selectively impair the nearest disparity constraint, we can simply increase the luminance difference between $L_A$ and $L_D$ while preserving the same luminances for A and C. Similarly, if we want to selectively impair the nearest neighbour constraint, we can increase the difference between $L_A$ and $L_C$ while preserving the same luminances for A and D. These two manipulations can be grouped together by creating a continuum that we call the *contrast ambiguity level* and denote $\phi$. Negative values (between $-1$ and 0) of this parameter correspond to an impairment of the nearest disparity constraint, and positive values (between 0 and +1) an impairment of the nearest neighbour constraint. The absolute value of $\phi$ corresponds to the difference in luminance between A and D (or A and C). Therefore, values of $\phi$ close to $-1$ will strongly favour the nearest neighbour constraint (transparent percept), while values close to +1 will strongly favour the nearest disparity constraint (single surface percept). When $\phi$ is zero, all dots have the same luminance and, at least as regards contrast similarity, the two constraints are equally supported. More precisely, the luminances of A, B, C and D are defined as

$$\begin{cases} L_A = (1-\phi)/2 \\ L_B = (1+\phi)/2 \\ L_C = (1+|\phi|)/2 \\ L_D = (1-|\phi|)/2 \end{cases} \tag{3}$$

Fig. 4a shows a composite stimulus where contrast ambiguity is manipulated across rows, whilst Fig. 4b illustrates different values of the contrast ambiguity scale. As can be seen in this composite stimulus, at extreme values of $\phi$, low luminance dots approach threshold visibility. However, since the ordering of high and low luminance dots is randomised across rows, no stimulus contains a surface comprised entirely of dots at threshold contrast. Readers should also note, however, that the variation of luminance and thus visibility may produce a confound between inter-ocular contrast and another dot property, that of perceived size. From the composite Fig. 4a it is evident that low luminance dots appear smaller than high luminance dots. This is a consequence of the Gaussian profile of the dots. As such, observers could base their matching on similarity of size rather than similarity of contrast. However, for the purpose of the experiments detailed here, this distinction is unimportant. The precise property of feature similarity that influences matching is of secondary interest to the underlying biases towards nearest neighbour or nearest disparity matching that it reveals.

### 2.2. Procedure

Stimuli were presented on a Sony Trinitron monitor with a refresh rate of 75 Hz and a resolution of $1152 \times 870$ pixels, powered by an Apple PowerMac computer with a G4 processor running at 867 MHz. Stimuli were generated and presented using Matlab™ combined with the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997). Observed stimuli were as described above, with the addition of a narrow zero disparity frame comprised of randomly positioned, regularly spaced squares to help maintain vergence. In experiment 1a (long duration) stereograms were presented for 2 s in a raised-cosine temporal window. In experiment 1b (short duration) stereograms were presented for 80 ms in a square-wave temporal window. All stimuli were viewed in a darkened room. The presentation of a stimulus was preceded by the presentation of the zero disparity frame and a zero disparity fixation cross for 0.5 s, whilst a mask of zero disparity, random pixel noise followed stimulus presentation. In experiment 1b participants were asked to align nonius lines presented above and below the fixation cross before initiating each trial. Stereo fusion was obtained using a split-screen Wheatstone stereoscope.

Participants were initially presented with examples of the stimulus at extreme levels of $\phi$ in order to familiarise them with the two dominant percepts. Following this familiarisation, participants were presented with stimuli at 19 different levels of $\phi$ in the long duration experiment (1a) and 9 different levels of $\phi$ in the short duration experiment (1b), ranging from −0.9 to 0.9. In experiment 1a participants responded to 40 repeated measures at each level of $\phi$ over the course of five blocks. In

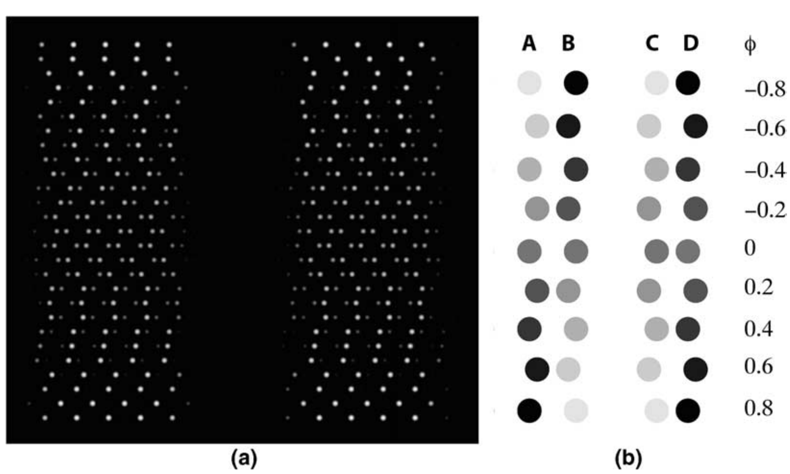


Fig. 4. Manipulation of the contrast ambiguity level. (a) Composite stimulus showing the effect of varying the contrast ambiguity level across rows. Contrast ambiguity level varies from −0.9 on the top row of the stimulus to 0.9 on the bottom row. As such, readers should note the change between the transparent percept at the top of the image and the single surface percept at the bottom of the image. Readers may also wish to judge where they see the switch between these percepts. (b) Illustration of the variation in contrast across rows. Readers should note the change in the dot corresponding to the nearest contrast following the point of uniform dot contrast. Values for the contrast ambiguity between dots are given in the final column of the figure.

experiment 1b 24 repeated measures were recorded over three blocks. Responses were made through key presses where the participants' task was to state whether they perceived one or two surfaces in depth. That is, did participants see the single surface, nearest disparity percept, or the two-surface, nearest neighbour percept?

### 2.3. Participants

Five observers (2 male, 3 female) participated in experiment 1a, including author RG. All had normal or corrected-to-normal vision. Three of the five participants were experienced psychophysical observers (BC, PW and RG) and, with the exception of the author, all participants were naïve. Five observers (2 male, 3 female) participated in experiment 1b, including author RG. All had normal or corrected-to-normal vision. Two of the five participants were experienced psychophysical observers. Only two of the participants (RG and LO) also participated in experiment 1a. With the exception of author RG, all participants were naïve. Participants were all postgraduate or postdoctoral researchers at the University of Glasgow and gave written consent of their participation.

### 2.4. Results

In reporting the experimental results we consider experiments 1a and 1b separately. We begin by reporting results for experiment 1a (long duration). Subsequently, we detail the results of experiment 1b (short duration) and compare the data obtained for the two experiments.

#### 2.4.1. Experiment 1a: long stimulus duration

Fig. 5 shows the proportion of "Single Surface" responses at each level of the contrast ambiguity scale $\phi$ for the 2 s presentation duration. Psychometric functions (cumulative Gaussians) were fitted to the proportion of "Single Surface" responses for each observer. These fits were very similar for the four naïve participants and markedly different for the author RG, both in terms of thresholds and slopes. These differences may be due to an idiosyncratic bias for observer RG, a greater prior exposure to the stimulus—that is, a familiarity with the transparent percept—or a greater sensitivity to noise in contrast matching. Consequently, we report below mean data averaged across all five participants and also mean data restricted to the four naïve participants (excluding RG).

The main result of the experiment can readily be seen by looking at the response bias when the stimulus was truly ambiguous, that is when the contrast ambiguity level was zero. On such occasions, naïve observers almost exclusively saw a single surface; that is, they relied strongly on the nearest disparity constraint. Across all five observers, the single surface was perceived on average 85.3% (SEM ±12.6%) of the trials (97.8 ± 1% if RG is excluded).

Manipulating the contrast ambiguity allowed us to make the task more interesting for the four naïve participants who would otherwise have almost always pressed the same key. It also provided a different measure of the bias for the nearest disparity constraint. The point of subjective equality (PSE, i.e. the 50% point) was extracted from each psychometric function. PSEs were biased towards negative $\phi$ values for the four naïve observers. Taking the average PSE across all participants (including RG) gives a tightly bound bias towards negative $\phi$ values (mean −0.272, SEM 0.105), which is significant on a one-sample $t$-test ($t_4 = 2.59$, $p < 0.05$ on a one-tailed test). Removing author RG from the mean data results in an even greater bias towards negative $\phi$ values. The mean $\phi$ value in this instance is −0.371 (SEM 0.048, $t_3 = 7.69$, $p < 0.01$). 95% confidence intervals for these cases are ± 0.18 and ± 0.08, respectively. The finding of such an extensive bias in PSE (i.e. beyond three nearest neighbour biased levels of $\phi$) indicates that, within limits, the nearest disparity bias is able to override the bias to match to similar contrasts. Values of $\phi$ may also be restated as values of Michelson contrast between horizontally adjacent dots. We may therefore say that a Michelson contrast of 0.272 (or 0.371 excluding RG) in favour of nearest neighbour matching results in the single and transparent surface perceptions being equally probable.

#### 2.4.2. Experiment 1b: short stimulus duration

Fig. 6 shows the data obtained for five observers in the short duration experiment (1b). Results are plotted as the proportion of "Single Surface" responses against each level of the contrast ambiguity scale $\phi$, and fitted with a psychometric function. There are two features to these data. First, the slopes of the fitted psychometric curves are, with the exception of participant DM, generally shallower than in experiment 1a. This may simply be due to the increased difficulty of the task given the brief presentation duration. Second, as with the 2 s presentation duration, participants are generally biased towards the nearest disparity solution. With the exception of MT, all participants have a negative PSE. Although one participant (MT) did not present a strong bias towards nearest disparity matching, she also did not show a strong bias towards the nearest neighbour solution. Instead, MT exhibited relatively unbiased matching. A one-sample $t$-test showed that the mean bias in the PSE was significantly different from zero ($t_4 = 2.14$, $p < 0.05$ on a one-tailed test), with a 95% confidence interval of ±0.158. Thus, observers are more likely to match to the nearest disparity solution when stimulus dots are all of identical contrast. The mean bias in the PSE for the 80 ms presentation was −0.18. This was

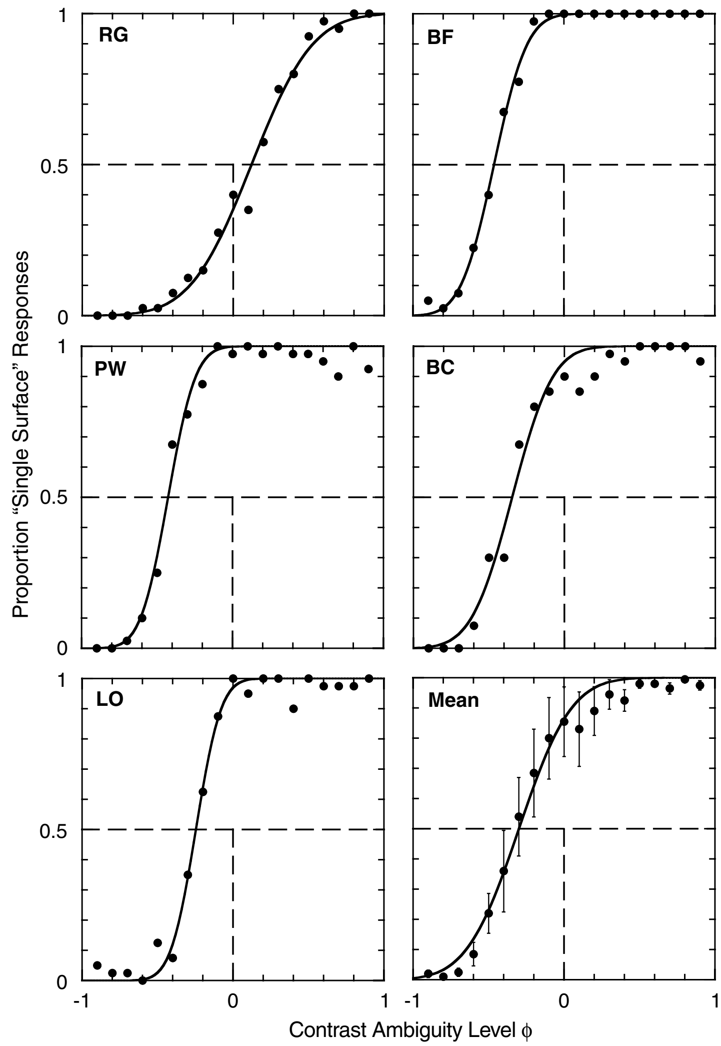


Fig. 5. Results of manipulating the contrast ambiguity level for five observers in experiment 1a (long duration). Results are plotted as the proportion of "Single Surface" responses against the contrast ambiguity level $\phi$ for each observer and are fitted with a cumulative Gaussian distribution (least squared fit). Negative contrast ambiguities indicate that contrast favours the nearest neighbour, transparency match, whilst positive contrast ambiguities indicate that contrast favours the nearest disparity, single surface match. Mean results at each level of $\phi$, across all five observers, are also shown. Error bars on mean results show the standard error on the mean.

not significantly different from the bias of $-0.272$ observed when the stimulus was presented for 2 s.

### 2.5. Discussion

Irrespective of the manner in which these results are reported, it remains clear that there is a strong bias towards matching to the nearest disparity solution. Such a result is consistent with previous empirical data (e.g. Zhang et al., 2001). However, since global relative disparity is the same for both percepts, we also show that the bias towards nearest disparity matching depends upon local differences in disparity rather than a preference for minimising the global relative disparity. The importance of local differences in disparity—and a matching strategy that seeks to avoid them—is consistent with the idea that the successful resolution of the correspondence problem is dependent upon the cohesiveness and piecewise smoothness of objects in the environment (Marr & Poggio, 1976, 1979). In experiments 2 and 3, we go on to examine the impact of varying the parameters of the stereogram on this tendency to resolve the stimulus in line with the nearest disparity solution.

The results of experiment 1b discount the possibility that the observed bias for nearest disparity matching is the result of vergence eye movements. Convergence of the eyes at the depth of the nearest disparity match eliminates the competition between nearest neighbour and nearest disparity constraints by presenting observers with the correspondence solution of a single surface containing no absolute or relative disparities. That is, a single, fronto-parallel surface at fixation. However, since

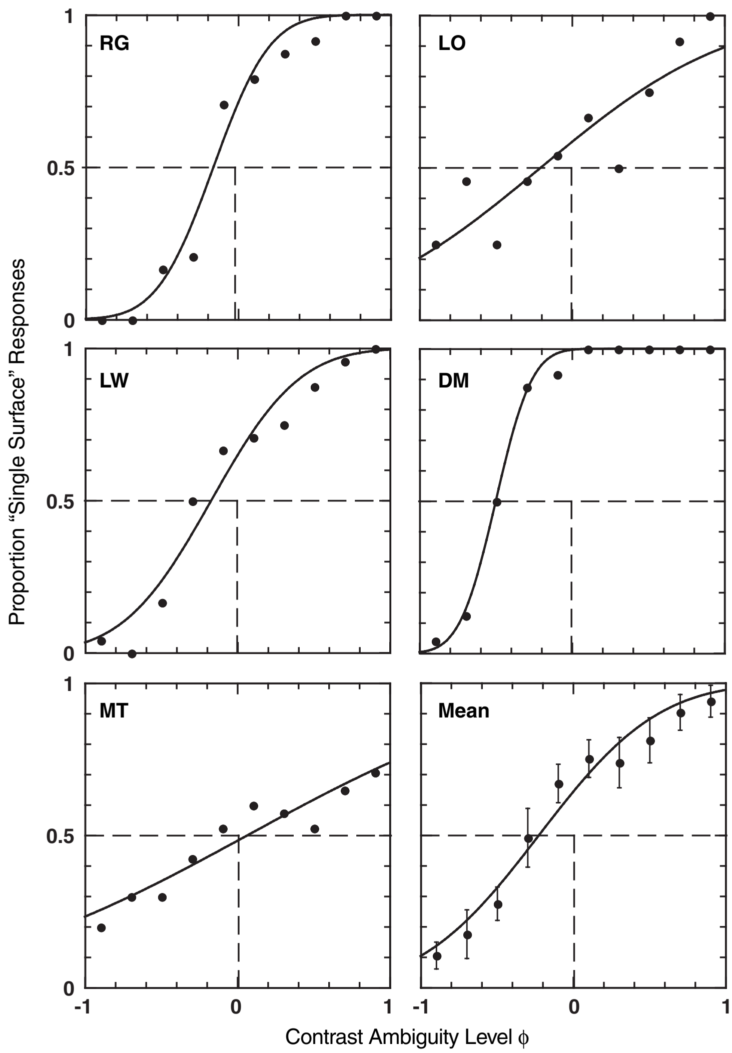


Fig. 6. Results of manipulating the contrast ambiguity level for five observers in experiment 1b (short duration). Results are plotted as the proportion of "Single Surface" responses against the contrast ambiguity level $\phi$ for each observer and are fitted with a cumulative Gaussian distribution (least squared fit). Negative contrast ambiguities indicate that contrast favours the nearest neighbour, transparency match, whilst positive contrast ambiguities indicate that contrast favours the nearest disparity, single surface match. Mean results at each level of $\phi$, across all five observers, are also shown. Error bars on mean results show the standard error on the mean.

the presentation duration of the stimulus in experiment 1b was too short to allow for convergence of the eyes, such movements cannot underlie the bias towards the single surface, nearest disparity solution. It should also be noted that an explanation of the nearest disparity bias based on vergence eye movements presupposes the detection of a disparity signal that can trigger such eye movements. As such, a vergence explanation of the nearest disparity bias presupposes the, at least partial, resolution of the correspondence problem.

## 3. Experiment 2

In experiment 1, we were able to selectively bias the correspondence matching process between solutions that fulfilled either the nearest neighbour or nearest disparity matching constraints. This biasing revealed a tendency for the visual system to resolve the experimental stimulus in a manner consistent with the nearest disparity constraint. However, other factors may also influence the correspondence matching process. In this experiment, we vary the depths of the two surfaces obtained from the nearest neighbour match by varying the value of $\delta$, the nearest neighbour offset. Readers should note from Eq. (2) that the variation of $\delta$ has no effect on the depth of the single surface given by the nearest disparity match.

Varying $\delta$ produces several changes in the experimental stimuli that may affect the resolution of the correspondence problem. First, it alters the magnitude of local differences in disparity between neighbouring points. Second, it alters the disparity between the nearest neighbour match and the horopter, thereby possibly altering the response of any mechanisms biased towards the nearest neighbour match. Third, the variation of $\delta$ alters the disparity between the two transparent surfaces as a whole by changing the global relative disparity in the scene (i.e. the disparity between front and rear surfaces). As has already been noted, a strategy of minimising relative disparity across the entire scene could not account for the nearest disparity bias observed in experiment 1. However, it does not necessarily follow that such a global process has no impact on correspondence matching. Sato and Yano (2000) have argued that global surface structure is important for the resolution of the correspondence problem, whilst Zhang et al. (2001) have argued that the correspondence matching process seeks to minimise the disparity between adjacent surfaces. As such, the variation of $\delta$ may alter the resolution of the correspondence process due to its effect on global differences in disparity. Given these possibilities, we hypothesise that reducing the value of $\delta$ will decrease the bias for nearest disparity matching, whilst increasing the value of $\delta$ will increase this bias.

### 3.1. Stimuli

Experimental stimuli were identical in structure and presentation to the stimuli used in experiment 1 except that the value of $\delta$ was varied. Three values of $\delta$ were used: $\delta_1 = 4.52$, $\delta_2 = 6.02$ and $\delta_3 = 7.53'$. Note that the global relative disparity between front and rear surfaces (i.e. $4\delta$) is now less that the global relative disparity of the single surface (i.e. $d_3 = \alpha - 2\varepsilon$) in the first instance and greater than the global relative disparity of this percept in the third instance.

### 3.2. Procedure

Stimulus generation and presentation was as in experiment 1. Over the course of five blocks, participants viewed 40 repeated measures of 9 levels of $\phi$ between values of $-0.8$ and 0.8. As in experiment 1, participants were presented with stimulus exemplars at extreme levels of $\phi$ in order to familiarise them with the dominant percepts elicited by the stimuli.

### 3.3. Participants

Five observers (3 male, 2 female) with normal or corrected-to-normal vision participated in experiment 2, including author RG. All participants, excluding RG, were naïve as to the experimental stimuli and hypothesis. Four of the participants were experienced psychophysical observers (RG, EG, PW and BC). Four observers had previously participated in experiment 1 (RG, PW, BC and LO). Participants were all postgraduate or postdoctoral researchers at the University of Glasgow and gave written consent of their participation.

### 3.4. Results and discussion

Fig. 7 shows the proportion of "Single Surface" responses for each participant at each level of the contrast ambiguity scale $\phi$, for each value of the nearest neighbour offset $\delta$. Psychometric functions (cumulative Gaussians) were fitted for each $\delta$ and the PSE extracted. The variation of $\delta$ produces small, but reliable, shifts in the PSE, in the order hypothesised ($F_{2,8} = 6.805$, $p = 0.019$). Pairwise comparisons showed that PSEs obtained for the $\delta_1$ and $\delta_3$ conditions were significantly different ($p < 0.05$ on a Newman–Keuls test). Mean PSEs for each value of $\delta$ were $-0.07$, $-0.16$ and $-0.25$. Thus, as the value of $\delta$ increased—increasing the disparity between front and rear surfaces in the nearest neighbour match—the tendency to match to the nearest disparity increased. Furthermore, although there was a significant shift in PSE due to the manipulation of $\delta$, matching was still generally biased towards the nearest disparity solution, as indicated by the sign of the PSE. Readers should also note that varying the value of $\delta$ had little effect on the slope of the obtained psychometric functions ($F_{2,8} = 0.577$, $p = 0.583$). That is, varying $\delta$ had no effect on the visual system's adherence to the contrast similarity constraint. Instead, varying $\delta$ is simply altering the extent of the preference for nearest disparity matching.

These results show that the correspondence matching process is sensitive to the stimulus changes produced through the manipulation of the nearest neighbour offset $\delta$. However, several mechanisms could be responsible for the observed change in matching preference, as discussed above. In experiment 3, we address further the issue of possible matching mechanisms.

## 4. Experiment 3

The results of experiment 2 show that manipulating the value of the nearest neighbour offset $\delta$ can alter the preference for nearest disparity matching. This variation in the preference for nearest disparity matching could be due to fluctuations in local mechanisms responsible for the measurement of disparity (i.e. a greater response for smaller absolute and locally defined relative disparities). However, these variations could also be due to changes at a more global level affecting any preference to minimise global relative disparity.

In order to address the issue between possible global and local mechanisms, we varied the value of $\alpha$ in the experimental stimulus. Readers should note from Eqs. (1) and (2) that varying $\alpha$—the standard horizontal dot separation—alters the depth of the single surface, nearest disparity match, without changing the depth of the

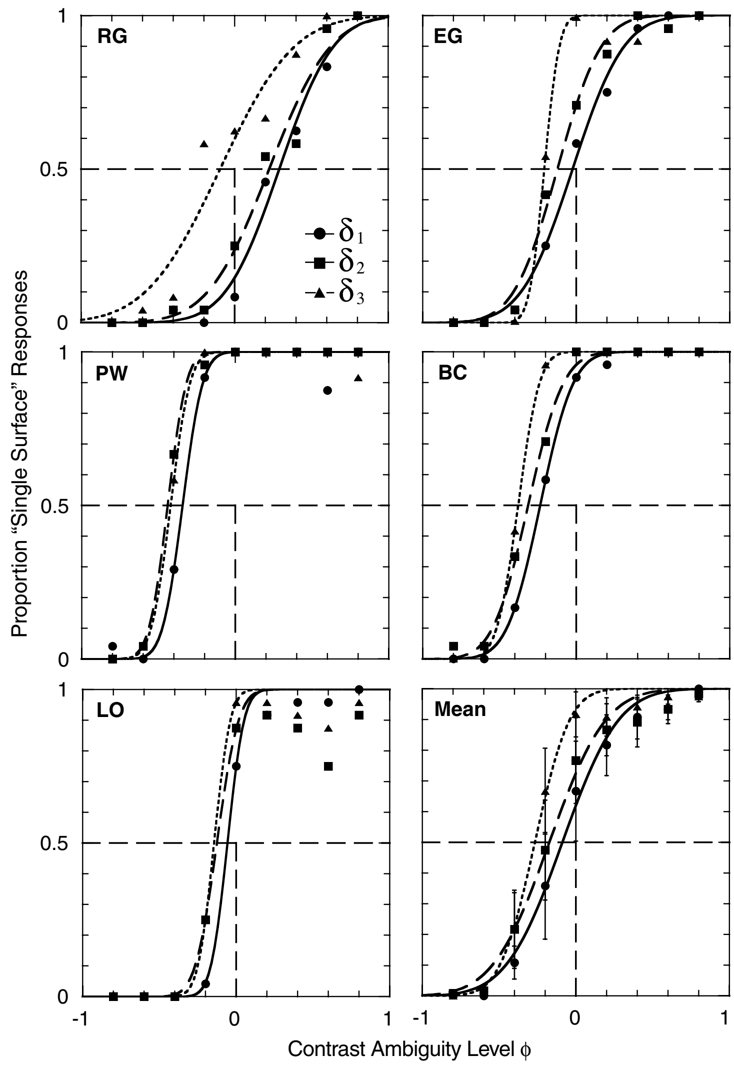


Fig. 7. Results of manipulating the contrast ambiguity level $\phi$ and nearest neighbour offset $\delta$ for five observers in experiment 2. Results are plotted as the proportion of "Single Surface" responses against contrast ambiguity level $\phi$ for each observer and are fitted with cumulative Gaussian distributions (least squared fit). Filled circles and continuous lines show the results for $\delta_1$. Filled squares and dashed lines show the results for $\delta_2$. Filled triangles and dotted lines show the results for $\delta_3$. Negative contrast ambiguities indicate that contrast favours the nearest neighbour, transparency match, whilst positive contrast ambiguities indicate that contrast favours the nearest disparity, single surface match. Mean results at each level of $\phi$, across all five observers, are also shown. Error bars on mean results show the standard error on the mean.

surfaces in the nearest neighbour match. As such it is the complementary manipulation to the change in $\delta$ used in experiment 2. Importantly, varying $\alpha$ produces the same change in the global relative disparity of the single surface percept as the variation of $\delta$ does for the two-surface percept. If the effects of varying the nearest neighbour offset $\delta$ are due to a preference for the minimal global relative disparity then varying the standard dot separation $\alpha$ should produce identical shifts in the PSE.

### 4.1. Stimuli

Stimuli were identical to those used in experiment 1, except that the value of $\alpha$ was varied. Three values of $\alpha$ were used: $\alpha_1 = 24.1$, $\alpha_2 = 30.1$ and $\alpha_3 = 36.14'$. The depth between the front and rear surfaces of the nearest neighbour match (i.e. $4\delta$) is less than the depth of the single surface (i.e. $d_3 = \alpha - 2\varepsilon$) for the lowest value of $\alpha$, but greater than the depth of the single surface for the highest value of $\alpha$. Readers should note that the disparity resulting from these values of alpha is equivalent to the disparity between the front and rear surfaces of the nearest neighbour match for the values of $\delta$ used in experiment 2.

### 4.2. Procedure

Stimulus generation and presentation was identical to that of experiments 1 and 2. As in experiment 2, participants viewed 40 repeated measures of 9 levels of $\phi$ between values of $-0.8$ and 0.8, over the course of five blocks. Prior to running the experiment, participants viewed exemplar stimuli to familiarise themselves with the dominant percepts upon which their responses are based.

### 4.3. Participants

Five observers (2 male, 3 female) with normal or corrected-to-normal vision participated in the experiment, including author RG. With the exception of RG, all participants were naïve. Three participants (RG, BC and EG) were experienced psychophysical observers. A further three observers had previously participated in both experiments 1 and 2 (RG, BC and LO), whilst one observer had previously participated in experiment 1a only (BF). The remaining observer had previously participated in experiment 2 only (EG). Participants were all postgraduate or postdoctoral researchers at the University of Glasgow and gave written consent of their participation.

### 4.4. Results and discussion

Fig. 8 shows the proportion of "Single Surface" responses for each participant at each level of the contrast ambiguity scale $\phi$, for each value of the standard dot separation $\alpha$. Psychometric functions (cumulative Gaussians) were fitted for each participant, for each value of

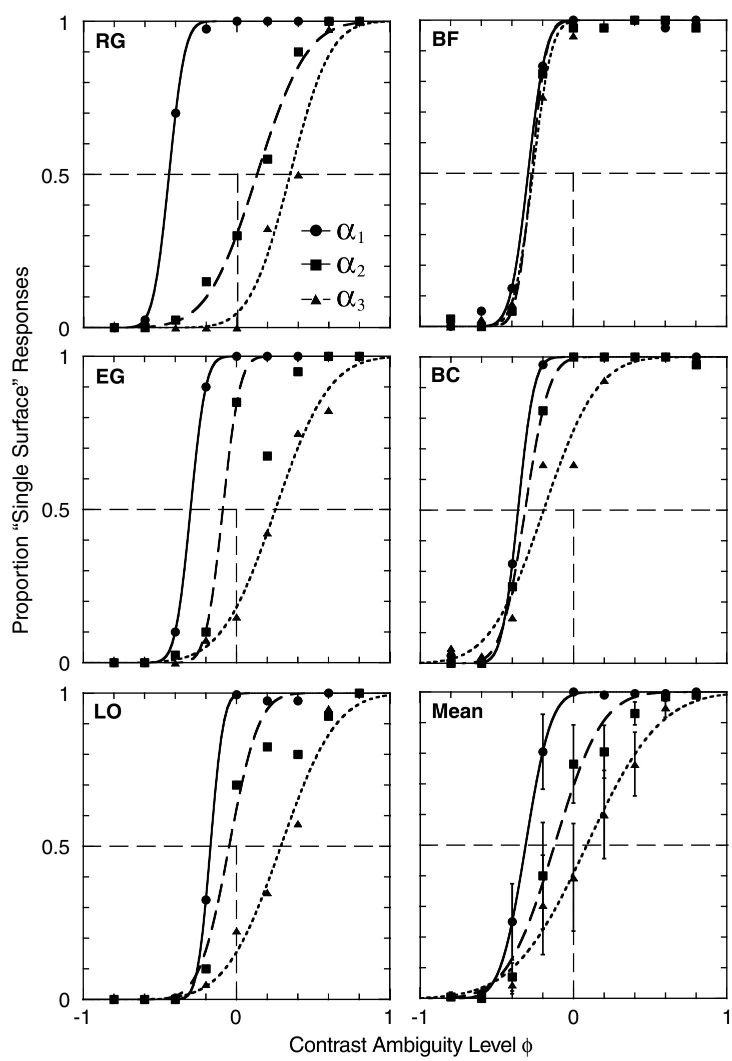


Fig. 8. Results of manipulating the contrast ambiguity level $\phi$ and the standard horizontal dot separation $\alpha$ for five observers in experiment 3. Results are plotted as the proportion of "Single Surface" responses against contrast ambiguity level $\phi$ for each observer and are fitted with cumulative Gaussian distributions (least squared fit). Filled circles and continuous lines show the results for $\alpha_1$. Filled squares and dashed lines show the results for $\alpha_2$. Filled triangles and dotted lines show the results for $\alpha_3$. Negative contrast ambiguities indicate that contrast favours the nearest neighbour, transparency match, whilst positive contrast ambiguities indicate that contrast favours the nearest disparity, single surface match. Mean results at each level of $\phi$, across all five observers, are also shown. Error bars on mean results show the standard error on the mean.

$\alpha$, and PSEs extracted. Varying $\alpha$ produces sizeable shifts in the PSE for four of the five observers. There also appears to be some change in the slope of the psychometric functions for these observers. In the remaining observer the shift due to $\alpha$ is much smaller (although still present and in the expected order). Mean PSEs at each level of $\alpha$ were −0.32, −0.12 and 0.09. Pairwise comparisons showed the shift between $\alpha_1$ and $\alpha_3$ to be significant ($F_{2,8} = 7.483$, $p = 0.015$; $p_{\alpha 3-\alpha 1} < 0.05$ on a Newman–Keuls test). Thus, as the value of $\alpha$ increased—increasing the depth of the single surface percept—the tendency to match to the nearest disparity decreased, in favour of the nearest neighbour percept. On average, the slope of the fitted psychometric function also varied between conditions. Pairwise comparisons showed the difference in slope between $\alpha_1$ and $\alpha_3$ to be significant ($F_{2,8} = 7.106$, $p = 0.017$; $p_{\alpha 3-\alpha 1} < 0.05$ on a Newman–Keuls test). Thus, unlike the manipulation of $\delta$, varying $\alpha$ also affected the influence of contrast similarity on matching.

If the resolution of the correspondence problem is affected by the difference in global relative disparity between possible matching solutions then the shift in PSEs due to the manipulation of $\alpha$ and $\delta$ should be identical. One may make this prediction since the manipulation of $\alpha$ and $\delta$ result in the same change in the global disparity differences between nearest neighbour and nearest disparity solutions. However, the shift in PSE due to the manipulation of $\alpha$ is far greater than the shift due to the manipulation of $\delta$. This suggests that the mechanisms affecting the change in tendency to match to the nearest disparity may be operating at a more local level.

## 5. General discussion

When confronted with a complex stereogram, the visual system has multiple options to match features between the left and right images. We have designed an ambiguous stimulus that allowed us to determine the strategy that is preferred by the visual system. In particular, we could alter the contrast of neighbouring dots in the image so as to bias matching towards solutions that conformed to either the nearest neighbour or nearest disparity constraint. When dots were of equal contrast, we found that most observers were strongly biased to perceive a single surface rather than two transparent surfaces. That is, they favoured the solution given by the nearest disparity over the nearest neighbour constraint. This bias for nearest disparity matching was also evident when inter-ocular dot contrast was manipulated, indicating that a preference to match to the nearest disparity can override the preference to match to similar contrasts. We also found that the tendency to perceive a single surface may be altered by varying the disparity of either the single surface or the two-surface percept.

Zhang et al. (2001) demonstrated a similar bias for nearest disparity matching, which they interpret as a rule for matching to the minimum relative disparity. However, the stimulus used by Zhang et al. (2001) does not allow them to distinguish between locally and globally computed minimisation of relative disparity. Such distinctions are important if we are to understand the mechanisms used by the brain to compute binocular disparity. In experiment 1, we demonstrate that the visual system possesses a locally implemented bias for nearest disparity matching. In experiments 2 and 3, we further demonstrate that changes in the nearest disparity bias—elicited by changing the disparity of alternative matching solutions—are also due to local disparity computations that prefer to minimise relative disparity. These findings suggest that the visual system is primarily concerned with integrating information within a local area in order to make judgements about depth in the scene. However, the neural mechanisms and computational strategies responsible for this local integration of information remain unclear.

In considering the neural mechanisms and computational strategies that underlie the observed bias for nearest disparity matching one must take care to avoid the reification of notional matching constraints. Here, experimental results have been discussed in terms of the preference for 'nearest disparity matching' as opposed to 'nearest neighbour matching'. However, the visual system may not be organised in a manner that is explicitly concerned with the adherence of potential correspondence matches to such rules. Instead, matching biases may arise due to mechanisms that implicitly prefer specific patterns of disparity. For example, an implicit nearest neighbour bias may arise due to the use of Gaussian envelopes in model disparity energy neurones (Qian & Zhu, 1997) or the re-weighting of a population of such neurones to attenuate large disparities (Prince & Eagle, 2000; Read, 2002a, 2002b). Below, we discuss how implicit nearest disparity biasing may occur and why this may be useful in resolving the correspondence problem.

### 5.1. Integration, segmentation and the nearest disparity bias

Anderson and Nakayama (1994) have argued that the resolution of correspondence ambiguity is best achieved by analysing small patches of information in a scene, rather than by considering individual dots. Similarly, other researchers have attempted to model disparity computation through an analysis that begins by processing the image through a set of linear filters (Jones & Malik, 1992) or through a set of model V1 binocular neurones (Prince & Eagle, 2000; Qian & Zhu, 1997;

Read, 2002a, 2002b)—so called 'disparity energy' neurones (Ohzawa, DeAngelis, & Freeman, 1990, 1996, 1997). Such an analysis utilises informationally enriched matching primitives (Anderson & Nakayama, 1994; Jones & Malik, 1992), allowing for an increasingly effective use of similarity constraints. Indeed, it has been shown that—in random-dot stereograms—correspondence ambiguity is almost entirely removed by an analysis based on small image patches (Sanger, 1988). However, the use of spatially extended stimulus patches as matching primitives also allows for the implicit implementation of a nearest disparity constraint (Prince & Eagle, 2000; Read, 2002a, 2002b).

Since V1 binocular neurones signal only absolute disparities (Cumming & Parker, 1999) and not disparity differences or gradients (Nienborg, Bridge, Parker, & Cumming, 2004), only a single disparity value may be signalled by a given image patch. Thus, increasing the spatial extent of the patch to be analysed increases the tolerance of the matching algorithm to small dissimilarities between left and right images that may represent the variation of disparity. That is, constant disparities (i.e. fronto-parallel surfaces), that result in only a horizontal translation of one patch compared to another will be preferred over more complex patterns of disparity that produce complex transformations between local patches. This preference for fronto-parallelism has previously been used as an implicit nearest disparity constraint (Prince & Eagle, 2000; Read, 2002a, 2002b).

An important issue in using extended patches of information as matching primitives is that the size of the patch may limit the disparity information that may be extracted from a stimulus (Banks, Gepshtein, & Landy, 2004). As already noted, if an image patch can have only a single disparity associated with it, the matching algorithm may miss fine variation in disparity. In particular, fine variation in disparity may be missed if the patch size is too large. Thus, in using large image regions to compute disparity, the visual system is confronted with a problem of segmentation: how can different disparities—as are present in the two-surface percept of our stimulus—be detected within a local region? If however, the patch is too small, the matching algorithm will not exhibit the seemingly important nearest disparity bias. That is, the correspondence problem will be reduced to a process that matches features akin to individual dots (cf. Banks et al., 2004; Kanade & Okutomi, 1994). Thus, when too small image regions are used to compute disparity, the visual system is confronted with a problem of integration, requiring a matching algorithm that addresses the classical false-target problem (Marr & Poggio, 1976, 1979). One possible means of overcoming these problems of segmentation and integration is by considering a correspondence matching procedure that utilises multiple patch sizes (Banks et al., 2004; Jones & Malik, 1992; Marr & Poggio, 1979; Prazdny, 1985). Such a procedure could be implemented either through lateral connections between V1 binocular neurones, or through feedback from higher visual areas with larger receptive fields, with information from multiple patch sizes (i.e. spatial scales) combined, perhaps, through coarse-to-fine or fine-to-coarse processing (Menz & Freeman, 2003; Menz & Freeman, 2004a, 2004b; Rohaly & Wilson, 1993; Smallman, 1995). Given these options, a nearest disparity constraint would be implicit in the integration of information across space.

### 5.2. Conclusions

Ambiguous stimuli have been used to study many visual phenomena under the assumption that the tendency to perceive such stimuli in a particular fashion can elucidate underlying neural mechanisms. Although previous research has made use of ambiguous stereograms (e.g. Anderson & Nakayama, 1994; Kontsevich, 1986; Mitchison & McKee, 1987a, 1987b; Papathomas & Julesz, 1989; Petrov, 2002; Ramachandran & Cavanagh, 1985), such stimuli have not been used extensively as a tool to examine the success of correspondence matching models. The stimulus and method presented here is ideally suited to the quantitative investigation of binocular correspondence matching and the testing of such models. In particular, this stimulus may aid our understanding of the factors mitigating segmentation and integration in stereo.

Finding factors that influence the tendency to make the nearest neighbour or nearest disparity match in such a stimulus may help us understand the computational and neural processes underlying stereo correspondence matching. Furthermore, the extent to which a model may predict whether the visual system chooses a nearest neighbour or nearest disparity solution to an ambiguous stereogram provides a stringent test of the adequacy of such a model and may offer a greater computational challenge than traditional random-dot stimuli.

## Acknowledgments

This work was supported by the EPSRC, UK (Grant No. GR/R57157/01). The authors wish to thank the anonymous reviewers for their helpful comments.

## References

Anderson, B. L., & Nakayama, K. (1994). Toward a general theory of stereopsis: binocular matching, occluding contours, and fusion. *Psychological Review, 101*, 414–445.

Arditi, A., Kaufman, L., & Movshon, J. A. (1981). A simple explanation of the induced size effect. *Vision Research, 21*, 755–764.

Banks, M. S., Gepshtein, S., & Landy, M. S. (2004). Why is spatial stereoresolution so low? *Journal of Neuroscience, 24*, 2077–2089.

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

Brewster, D. (1844). On the knowledge of distance given by binocular vision. *Transactions of the Royal Society of Edinburgh, 15*, 663–674.

Cumming, B. G., & Parker, A. J. (1999). Binocular neurons in V1 of awake monkeys are selective for absolute, not relative, disparity. *Journal of Neuroscience, 19*, 5602–5618.

Halpern, D. L., & Blake, R. (1988). How contrast affects stereoacuity. *Perception, 17*, 483–495.

Howard, I. P., & Rogers, B. J. (2002). Seeing in Depth, Vol. 2: Depth Perception. I Porteous: Ontario.

Jones, D. G., & Malik, J. (1992). Computational framework for determining stereo correspondence from a set of linear spatial filters. *Image and Vision Computation, 10*, 699–708.

Julesz, B. (1964). Binocular depth perception without familiarity cues. *Science, 145*, 356–362.

Kanade, T., & Okutomi, M. (1994). A stereo matching algorithm with an adaptive window: theory and experiment. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 16*, 920–932.

Kontsevich, L. L. (1986). An ambiguous random-dot stereogram which permits continuous change of interpretation. *Vision Research, 29*, 517–519.

Legge, G. E., & Gu, Y. (1989). Stereopsis and contrast. *Vision Research, 29*, 989–1004.

Marr, D., & Poggio, T. (1976). Cooperative computation of stereo disparity. *Science, 194*, 283–287.

Marr, D., & Poggio, T. (1979). A computational theory of human stereo vision. *Proceedings of the Royal Society of London. Series B: Biological Sciences, 204*, 301–328.

McKee, S. P., & Mitchison, G. J. (1988). The role of retinal correspondence in stereoscopic matching. *Vision Research, 28*, 1001–1012.

Menz, M. D., & Freeman, R. D. (2003). Stereoscopic depth processing in the visual cortex: a coarse-to-fine mechanism. *Nature Neuroscience, 6*, 59–65.

Menz, M. D., & Freeman, R. D. (2004a). Temporal dynamics of binocular disparity processing in the central visual pathway. *Journal of Neurophysiology, 91*, 1782–1793.

Menz, M. D., & Freeman, R. D. (2004b). Functional connectivity of disparity-tuned neurons in the visual cortex. *Journal of Neurophysiology, 91*, 1794–1807.

Mitchison, G. J., & McKee, S. P. (1987a). The resolution of ambiguous stereoscopic matches by interpolation. *Vision Research, 27*, 285–294.

Mitchison, G. J., & McKee, S. P. (1987b). Interpolation and the detection of fine structure in stereoscopic matching. *Vision Research, 27*, 295–302.

Nienborg, H., Bridge, H., Parker, A. J., & Cumming, B. G. (2004). Receptive field size in v1 neurons limits acuity for perceiving disparity modulation. *Journal of Neuroscience, 24*, 2065–2077.

Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1990). Stereoscopic depth discrimination in the visual cortex: neurons ideally suited as disparity detectors. *Science, 249*, 1037–1041.

Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1996). Encoding of binocular disparity by simple cells in the cat's visual cortex. *Journal of Neurophysiology, 75*, 1779–1805.

Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1997). Encoding of binocular disparity by complex cells in the cat's visual cortex. *Journal of Neurophysiology, 77*, 2879–2909.

Papathomas, T. V., & Julesz, B. (1989). Stereoscopic illusion based on the proximity priniciple. *Perception, 18*, 589–594.

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

Petrov, Y. (2002). Disparity capture by flanking stimuli: a measure for the cooperative mechanism of stereopsis. *Vision Research, 42*, 809–813.

Pollard, S. B., Mayhew, J. E. W., & Frisby, J. P. (1985). PMF: a stereo correspondence algorithm using a disparity gradient limit. *Perception, 14*, 449–470.

Prazdny, K. (1985). Detection of binocular disparities. *Biological Cybernetics, 52*, 93–99.

Prince, S. J. D., & Eagle, R. A. (2000). Weighted directional energy model of human stereo correspondence. *Vision Research, 40*, 1143–1155.

Qian, N., & Zhu, Y. (1997). Physiological computation of binocular disparity. *Vision Research, 37*, 1811–1827.

Ramachandran, V. S., & Cavanagh, P. (1985). Subjective contours capture stereopsis. *Nature, 317*, 527–530.

Read, J. C. A. (2002a). A Bayesian model of stereopsis depth and motion direction discrimination. *Biological Cybernetics, 86*, 117–136.

Read, J. C. A. (2002b). A Bayesian approach to the stereo correspondence problem. *Neural Computation, 14*, 1371–1392.

Rohaly, A. M., & Wilson, H. R. (1993). Nature of coarse-to-fine constraints on binocular fusion. *Journal of the Optical Society of America, A, 10*, 2433–2441.

Sanger, T. D. (1988). Stereo disparity computation using gabor filters. *Biological Cybernetics, 59*, 405–418.

Sato, N., & Yano, M. (2000). A model of binocular stereopsis including a global consistency constraint. *Biological Cybernetics, 82*, 357–371.

Smallman, H. S. (1995). Fine-to-coarse scale disambiguation in stereopsis. *Vision Research, 35*, 1047–1060.

Smallman, H. S., & McKee, S. P. (1995). A contrast ratio constraint on stereo matching. *Proceedings of the Royal Society of London. Series B: Biological Sciences, 260*, 265–271.

Tsai, J. J., & Victor, J. D. (2003). Reading a population code: a multiscale neural model for representing binocular disparity. *Vision Research, 43*, 445–466.

Zhang, Z., Edwards, M., & Schor, C. M. (2001). Spatial interactions minimize relative disparity between adjacent surfaces. *Vision Research, 41*, 2995–3007.