

The efficiency of depth discrimination for non-transparent and transparent stereoscopic surfaces

Julian Michael Wallace *, Pascal Mamassian

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

Received 20 August 2003; received in revised form 23 April 2004

Abstract

The perception of transparency in binocular vision presents a challenge for any model of stereopsis. We investigate here how well human observers cope with stereo transparency by comparing their efficiency between transparent and opaque depth judgments. In two experiments, the efficiency measure was computed relative to an ideal observer to take into account the larger correspondence ambiguity in the transparent condition. We found that thresholds for human and ideal observers were consistently higher in the transparent condition than in the opaque condition, across a range of dot densities (Experiment 1) and disparity ratios (Experiment 2). Efficiencies (the ratio of human to ideal performance) were approximately equivalent for the opaque and transparent conditions across all stimulus conditions. Therefore, the cost for stereoscopic transparency can be accounted for by the greater correspondence problem in that condition. Indeed, the fact that efficiencies were very low, around 1%, and decreased with increasing dot density demonstrates that human observers use far less information than is available to perform the task. This account contrasts with previous interpretations for the cost in stereoscopic transparency in terms of inhibitory interactions specific to transparent configurations. We relate our findings to a previous and comparable study of motion efficiency, and discuss our findings in terms of a physiologically plausible model.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Stereopsis; Transparency; Correspondence problem; Ideal observer; Efficiency

1. Introduction

Binocular disparities arise when points at different depths project to the two eyes. Inferring depth back from these binocular disparities is however a very intricate problem. Even if the visual system knew precisely the projection of features in the two eyes, it would still be faced with the combinatorial problem of matching corresponding features between eyes. Indeed, each point in one image could conceivably be matched with any point in the other image. This combinatorial problem that maps a single set of disparities to multiple compatible three-dimensional scenes has been called the *correspondence problem*.

A particularly acute case of the correspondence problem occurs when two disparities are present simul-

taneously in the same visual location. This situation occurs when we view one surface behind another, such as when looking through transparent, reflective surfaces like glass or water. Here we perceive a scene through the transparent surface, yet we also perceive the surface through which we view the scene, due to reflections or specularities. The visual system successfully groups similar disparities and segments dissimilar disparities to recover two surfaces segregated in depth, despite the fact that points from each surface will project to only one point on the retina. Stereo algorithms that employ the uniqueness and continuity constraints of Marr and Poggio (1976, 1979; see also Grimson, 1985) will be unable to recover such scenes, as they do not permit the occurrence of more than one disparity at a given visual location. Indeed, these constraints apply only to smooth opaque surfaces.

Psychophysical studies using variations of random-dot stereograms (Julesz, 1964) have demonstrated that the uniqueness constraint can be violated. Specifically, random dot versions of Panum's limiting case

* Corresponding author. Address: CNRS, INCM, 31 Chemin Joseph Aiguier, 13402 Marseille Cedex 20, France. Tel.: +33-491164523; fax: +33-491774969.

E-mail address: wallace@incm.cnrs-mrs.fr (J.M. Wallace).

(Kaufman, Bacon, & Barraso, 1973) and the double-nail illusion (Weinshall, 1989, 1991) can be perceived as one or more transparent surfaces in depth against an opaque background, although it has been argued that these percepts do not necessarily depend upon non-unique matches (Pollard & Frisby, 1990; Pollard, Mayhew, & Frisby, 1985). The PMF stereo algorithm (Pollard & Frisby, 1990; Pollard et al., 1985) implements the uniqueness constraint by restricting matching to a disparity gradient limit (Burt & Julesz, 1980). This algorithm can recover isolated patches of the different disparities in transparent random dot stereograms (though does not interpolate these patches of disparity to recover two surfaces at different depths). This is in contrast, for example, to the stereo algorithm of Prazdny (1985) that naturally permits the resolution of transparency by using *only* excitatory interactions (and interpolates the computed disparities across the image).

Despite the computational significance of stereoscopic transparency, the psychophysical research is surprisingly sparse. A few studies have assessed the limits of stereoscopic transparency with random dot stereograms that contain two disparities. For such stimuli, there is a continuum of percepts as the difference in disparity is increased (Parker & Yang, 1989; Stevenson, Cormack, & Schor, 1989; Tyler, 1991), from a single plane, through a thickened plane ('pyknostereopsis'), to transparency ('diastereopsis'). The mechanisms underlying stereoscopic transparency were further studied by Akerstrom and Todd (1988). They found that observers were less likely to perceive segregated transparent planes as the overall disparity, and the disparity difference, of the two planes was increased (the disparity differences were above the lower limits previously reported). In contrast, increasing the disparity did not impair the segregation of the opaque surfaces. Akerstrom and Todd (1988) argued that these results demonstrated both facilitatory and inhibitory interactions between different disparity detectors. In the transparent condition, disparity varies sharply across the image, and inhibitory interactions between largely different disparities would limit any facilitatory interactions. They argued that increasing the dot density (and presumably the disparity) would increase the strength of the inhibition, leading to the degraded perception of transparency they found. However, they did not assess the effect of density on a non-transparent display, and so the results cannot be taken as clear evidence of inhibition specific to transparent surfaces. More recently, Gepshtein and Cooperman (1998) also argued for inhibitory interactions between differently tuned disparity detectors. They presented a random dot stereogram of a cylinder behind a transparent plane. Observers were required to report the orientation of the cylinder, horizontal or vertical. They found that, to perform at a particular level, observers required the dot density of the transparent

plane to be lowered as the depth separation between the surfaces was increased. As the effect was weaker for surfaces of different contrast polarities, they argued that this behaviour could be accounted for by inhibitory interactions between disparity detectors for the different surfaces. However, there was no comparison with a non-transparent condition i.e. it is not clear if orientation judgments for a cylinder at different depths depend upon overall density in a non-transparent configuration.

In the present study we use the efficiency measure to quantify the limitations on stereoscopic transparency. Efficiency is an absolute measure of performance computed by comparing human performance with that of the ideal observer that utilises all of the information in a given stimulus to perform a given task optimally (Barlow, 1978; Green & Swets, 1966). Therefore, it is a measure of the amount of visual information actually used by a human observer to perform a task. Here we compute the efficiency for depth discrimination of transparent random dot stereograms and comparable opaque stereograms, in two experiments. Within each experiment, in a transparent condition we presented two populations of dots at different disparities simultaneously, while in the opaque condition we presented each disparity sequentially. The key difference between these conditions is that there is a greater correspondence problem in the transparent case. By comparing performance with an ideal observer that is only limited by correspondence noise (i.e. false dot matches) we could assess whether correspondence noise accounts for the impairment in performance for stereoscopic transparency. In Experiment One, we compared depth discrimination for non-transparent and transparent stereoscopic surfaces as a function of the dot density of the stimuli, and in Experiment Two we made the same comparison but as a function of the disparity ratio between the two surfaces to be discriminated.

2. General methods

Here we describe the basic methods for both experiments. More specific details will be provided for each experiment.

2.1. Human observers

Three experienced psychophysical observers participated, one experimenter (JW), and two paid graduate students (RG & VL). All observers had normal or corrected-to-normal visual acuity and a good stereo acuity.

2.2. Apparatus

Stimuli were presented on a 21" Sony Trinitron Flatscreen monitor via a G4 Power Macintosh running

MATLAB with the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). The monitor refresh rate was set to 75 Hz at a resolution of 1152 by 870. The stimuli were viewed binocularly via a modified Wheatstone mirror stereoscope in a dimly lit room at a viewing distance of 800 mm. Observers used a chin rest to stabilize head position throughout the experiment and fixated on a central white fixation cross, of length 0.30° of visual angle.

2.3. Stimuli

In both experiments, the stimuli were random dot stereograms constructed by randomly placing dots on the left and right images and presenting these images separately to each eye via the Wheatstone stereoscope. Each image consisted of $4.5'$ white squares ('dots') on a black background, 7.5° by 7.5° of visual angle. The remainder of the screen was set to the mean luminance of the stimulus (which varied with the dot density), to maintain a uniform mean luminance across the entire display.

A proportion of the dots were referred to as 'signal dots'; these dots corresponded to the projection of dots on surfaces located either near or far relative to the fixation plane. The remaining dots were referred to as 'noise dots'; these dots were randomly placed independently in each image. Two examples of stimulus are

illustrated in Fig. 1a (without noise) and Fig. 1b (with noise). In Fig. 1a the stereogram contains only two disparities, corresponding to a 'near' transparent surface and a 'far' opaque surface. Fig. 1b contains the same signal disparities, but now a proportion of dots are 'noise dots'. The effect of these added dots is to create more ambiguity in the matching, and results in the perception of dots at many depths. Indeed, some of these matches will be false correspondences resulting from incorrectly matching a signal dot with a non-corresponding noise dot. At the level of noise shown in Fig. 1b it is still possible to perceive the two surfaces, but they are noticeably less clear.

2.4. Procedure

In both experiments, we presented our random dot stereograms in two conditions. In the *transparent* condition, each trial consisted of two disparity signals superimposed in the same interval of 2000 ms duration. One signal (standard or target) was at uncrossed disparity (for a 'far' depth), while the other (target or standard) was at crossed disparity (for a 'near' depth). The depth (near or far) of the target stimulus was randomised across trials. To ensure fusion of the stereograms, each trial was preceded for 500 ms by a fixation cross with nonius lines, centered in the presentation window. The fixation cross was present throughout each




Fig. 1. A transparent stereogram. (a) This stereogram contains two populations of dots, one at crossed disparity and the other at uncrossed disparity. When the stereogram is fused (the two leftward panels are arranged for crossed convergence, and the two rightward panels for uncrossed convergence), a 'near' transparent surface is perceived in front of a 'far' opaque surface. Here the 'far' surface is further from the fixation cross. (b) As in (a), this stereogram contains two populations of dots, one at crossed disparity and the other at uncrossed disparity. However, now a proportion of dots are 'noise', randomly placed in the left and right window. When the stereogram is fused it is still possible to perceive a 'near' transparent surface and a 'far' opaque surface, but they are now embedded in a cloud of dots and are harder to see than before.

trial. In the *opaque* condition, again each trial consisted of two random dot signals, but now presented sequentially in temporal intervals of 2000 ms duration each. In one interval the signal disparity was crossed for a near depth, in the other the signal disparity was uncrossed for a far depth. There was a period of 500 ms between intervals in which only the fixation cross was present, the depth (near or far) of the target stimulus was randomised across trials as was the order of the stimulus presentation.

3. Experiment 1: Effects of dot density

The general aim was to assess whether there is a processing limitation for stereotransparency by comparing efficiencies of depth discriminations for both opaque and transparent conditions. Specifically, in this first experiment we compare the efficiency of depth discrimination in the opaque and transparent conditions across a range of dot densities. As dot density increases, the number of dots and therefore the number of possible correspondences between dots increases. If the mechanisms of stereopsis underlying performance in both the opaque and transparent conditions are sensitive to false correspondences, performance will be similarly impaired as dot density is increased.

3.1. Methods

3.1.1. Stimuli

For *each trial* two sets of signal dots were generated, one for the ‘near’ surface and one for the ‘far’ surface. One of these surfaces could be further from a zero-disparity fixation plane, while the other could be nearer. The surface further from fixation was defined by a *target* disparity, and the surface nearer to fixation by a *standard* disparity. For each surface, we generated a random dot image off-screen, with a total length equal to the desired size of the stereo image plus the disparity for that surface. We then sampled the off-screen image twice to generate the left and right stereo-halves, each stereo-half sampled at a horizontal increment equal to either the standard or target disparity. For example, for a disparity of 6 pixels (far depth) the off-screen random dot image would be sampled at +3 pixels for the left stereo-half and –3 pixels for the right stereo-half. This sampling increment results in corresponding dots to be uniformly displaced in each image at the appropriate disparity (because each stereo-half contained a displaced sample of the image, a small proportion of ‘signal’ dots in each image had no corresponding points). In the transparent condition, the left stereo-halves of each disparity defined surface were superimposed, and similarly for the right stereo-halves. Before presentation of the

stimulus, a proportion of noise dots were randomly placed on the left and right stereo-halves independently.

3.1.2. Procedure

The purpose of the first experiment was to study the effect of dot density on the perception of transparent and opaque random dot stimuli. We used a range of dot densities: 0.005, 0.01, 0.02, 0.04, 0.08, 0.16 and 0.32. These densities correspond to total dot numbers of 50, 100, 200, 400, 800, 1600 and 3200 dots, and to 0.89, 1.78, 3.56, 7.12, 14.24, 28.48 and 56.96 dots per squared degree of visual angle. The dot density refers to the total dot density of the stimulus, such that each interval of the opaque condition had a density of half the total value. For example, a total dot density of 4% corresponds to a 4% dot density for the transparent condition, but a 2% dot density for *each* interval of the opaque condition. Note that we will plot the results with respect to the total dot density. The observer’s task was to decide whether the ‘near’ or ‘far’ surface was further from the fixation plane, a 2-AFC depth discrimination. The two possible alternatives (‘far’ is further from fixation, and ‘near’ is further) are illustrated in Fig. 2 for the case of a transparent stimulus. The standard disparity was fixed at 9′ for all three experiments, while the larger target disparity was fixed at 18′, giving a *disparity ratio* of 2 (18/9). To limit performance, we presented the signals in a number of noise levels by the method of constant stimuli. We tested five noise levels per condition and measured d' for each noise level we tested. In both the transparent and opaque conditions each observer completed 20 practice trials with 0% noise to become familiar with the stimulus before beginning a session for a new condition. There were equal numbers of near-further and far-further trials. Each density condition was blocked, with 40 trials per each noise level (20 near-further, 20 far-further). Within each density condition, trials for different noise levels were randomly interleaved.




Fig. 2. A cartoon illustration of the two stimulus alternatives, viewed from an overhead. The plane of fixation is defined by a fixation cross, and the near and far surfaces by dots at the appropriate disparities. Observers decide which surface is further from this reference plane. On the left of the illustration the far surface is further from the fixation plane, and the example observer makes the correct response. On the right of the illustration the near surface is further from the fixation plane, and the example observer makes the correct response.

3.1.3. Ideal observer

The ideal observer for a given task makes use of all the available information in a given stimulus to perform that task optimally, i.e. maximising the number of correct responses by performing a maximum likelihood estimate (Green & Swets, 1966). For the experiments in this study, the ideal observer is facing the same depth discrimination task as any human observer. The ideal observer needs to represent the disparities displayed in the stimulus, compare these disparities to the disparities of the possible templates, and choose the appropriate template that best matches the disparities in the stimulus (Fig. 3). The disparities of each stimulus are computed by cross-correlating the left and right images of the stimulus. These images are simply binary matrices, in which ‘1’ signals the presence of a dot and ‘0’ is the background. The cross-correlation function describes the quantity of matches at each disparity. It is not a model of the human stereoscopic system, although the cross-correlation function has been used as the basis of a model of human stereoscopic vision (e.g. Cormack, Stevenson, & Schor, 1991), and moreover disparity selective complex cells can be understood as performing a form of cross-correlation (local band-pass filtered and phase insensitive; see Qian & Zhu, 1997). For the

transparent stimulus a single disparity correlation is performed on the left and right images. For the opaque stimulus two disparity correlations are performed, one for each interval. The correlations for both intervals are then summed. At low external noise levels, the peaks of this disparity correlation correspond to the standard and target signals. This can be seen in Fig. 3 (Box A) for a transparent stimulus with 0.70 noise dots (0.30 signal dots), in which the far surface is further. The ideal algorithm computes the likelihood of each possible outcome by comparing the incoming stimulus with a number of ‘templates’. Each template is a representation of the possible alternatives that were illustrated in Fig. 2 (‘far’ is further or ‘near’ is further). These templates are correlations that peak at the expected disparities (Fig. 3, Box B). The exact disparities will correspond to the disparities presented within a given block of trials. In Fig. 3 (Box B) the possible alternatives are given for a disparity ratio of 2. To compute the likelihood of each possible outcome, the ideal algorithm cross-correlates the stimulus correlation with each template (Green & Swets, 1966). The ideal decision rule is then to choose the template that returns the largest cross-correlation value with the stimulus (Fig. 3, Box C), a maximum likelihood decision rule. In the case of low external

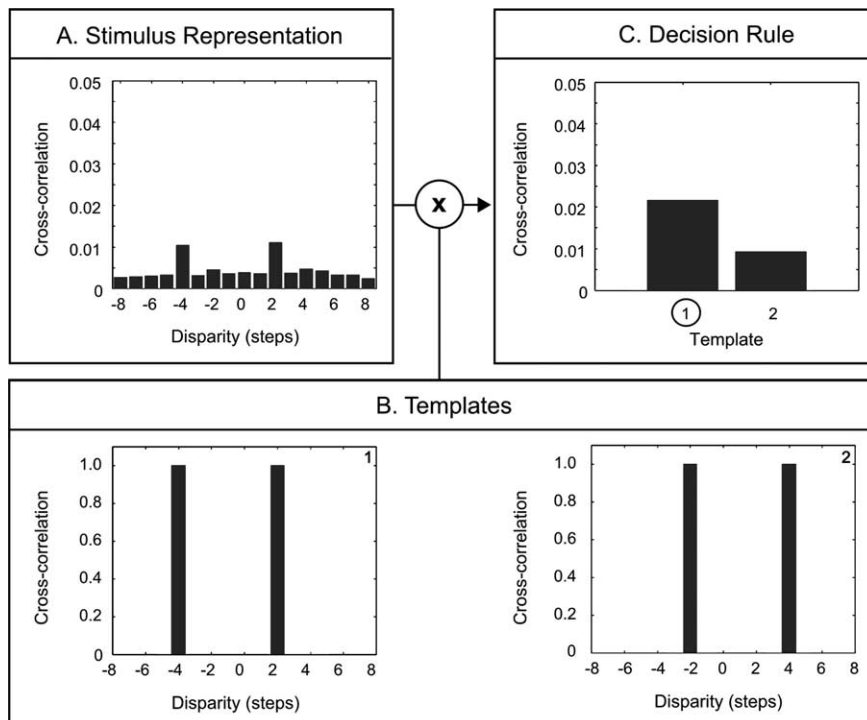


Fig. 3. A schematic illustration of the ideal observer for the depth discrimination task of this study. (1) Stimulus representation: this is the cross-correlation of the left and right images for a transparent stimulus, in which the ‘far’ surface is further from fixation (disparity ratio = 2; dot density = 0.05; proportion of signal dots = 0.30). The correlation peaks at a lag of -4 (a total uncrossed disparity of four dot steps), and +2. (2) Templates: these are the templates for a disparity ratio of 2. Template ‘1’ on the left represents the stimulus in which the ‘far’ surface is further from fixation, and template ‘2’ on the right represents the stimulus in which the ‘near’ surface is further from fixation. (3) Decision rule: the computed correlations for template ‘1’ and template ‘2’ with the stimulus correlation. The correlation is largest for template 1, the correct stimulus, and is selected by the ideal observer.

noise, the template with the highest value will correspond to the actual signal presented, and in Fig. 3 the ideal observer indeed selects the correct template. However, at much lower signal levels the value of the incorrect template can be higher than that of the correct template. Only these occurrences limit the ideal observer performance.

The effects of varying the signal level and the dot density on the stimulus correlation, and therefore the predicted effects on ideal performance, can be seen in Fig. 4. The left columns are correlations for stimuli of 16% density ($d = 0.16$), and the right columns are correlations for stimuli of 32% density ($d = 0.32$). The correlations represented by filled bars are for the opaque condition, and the correlations represented by open bars are for the transparent condition (the open bars are

presented upside-down for better comparison with the filled ones). Each row contains correlations for a particular level of signal, the top row is for 100% signal dots (where the proportion of noise dots is zero, $n' = 0$), the middle row is for 50% signal dots ($n' = 0.50$), and the bottom is for 5% signal dots ($n' = 0.95$). First consider the effects of decreasing the proportion of signal dots (thereby increasing the proportion of noise dots). In the top row two peaks are clearly distinguishable, corresponding to the signal disparities. However, even with 0% noise dots, there are spurious matches at the non-signal disparities, due to matching different signal dots. The ideal observer selects the correct template because the value of the noise at the incorrect signal disparities is much lower than the correct signal disparities. In the middle row the proportion of signal dots has dropped




Fig. 4. Cross-correlations for a number of stimuli, for each correlation ' d ' indicates the dot density and ' n' ' the proportion of noise (so $1 - n'$ is the proportion of signal dots). All the correlations are for a disparity ratio of 2, in which the 'far' surface is further. Dark bars are for the opaque condition, and light bars are for the transparent condition. It can be seen that increasing the noise level decreases the strength of the signal. Increasing the dot density increases both the strength of the signal and of the spurious correlations. Note that the spurious correlations are stronger in the transparent condition than the corresponding opaque condition.

and the corresponding peaks have also dropped, however the value of the noise disparities has not noticeably changed. In the bottom row the proportion of signal dots has been decreased further still. Here the two signal disparities are no longer distinguishable from the background noise in the transparent condition, but are still present in the opaque condition (this is not easily apparent in the 0.16 density correlation, but is clear for the 0.32 density condition). Now the ideal observer is just as likely to select the incorrect template as the correct template in the transparent condition, as the values for the incorrect disparities may be larger than the correct disparities by chance matches. However, in the opaque condition the correct template will be selected. This predicts that the ideal observer thresholds will be higher in the transparent condition. The second aspect of the correlations to consider is the effect of density. As density is increased twofold from the left column to the right column, it is clear that the values of the noise disparities increase. However, the value of the signal disparities also increases. Therefore, dot density will affect ideal performance if the increase in signal and noise amplitudes differs e.g. if the signal amplitude increases proportionally more than the increase in the noise amplitude then ideal performance should improve. We return to these aspects when considering the simulated data. We ran simulations of the ideal observer for both the transparent and opaque conditions. To compute ideal sensitivity, the simulations were performed at five noise levels for each condition, with 400 trials (200 near-further, 200 far-further) per noise level. Efficiency is the ratio of human sensitivity to that of the ideal observer (Barlow, 1978; Tanner & Birdsall, 1958):

$$F = \left(\frac{d'_h}{d'_i} \right)^2 \tag{1}$$

The problem in using this definition is that the ideal observer easily reaches ceiling performance for a suitable range of signal values for the human observer. Thankfully, as we will see in the results section below, d' is proportional to the proportion of signal dots presented. We can therefore (see Appendix A for the derivation) compute efficiency as the squared ratio of the signal thresholds:

$$F = \left(\frac{\theta_i}{\theta_h} \right)^2 \tag{2}$$

3.2. Results

An example of the data obtained is shown in Fig. 5 for both a human observer and a set of simulation of the ideal observer. These data are for the transparent condition, with a dot density of 1%, and a disparity ratio of 2 (standard disparity 0.15°, target disparity 0.30°). It can

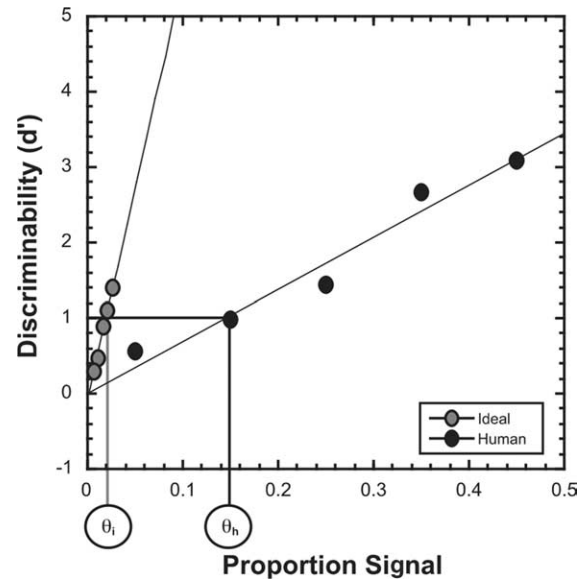
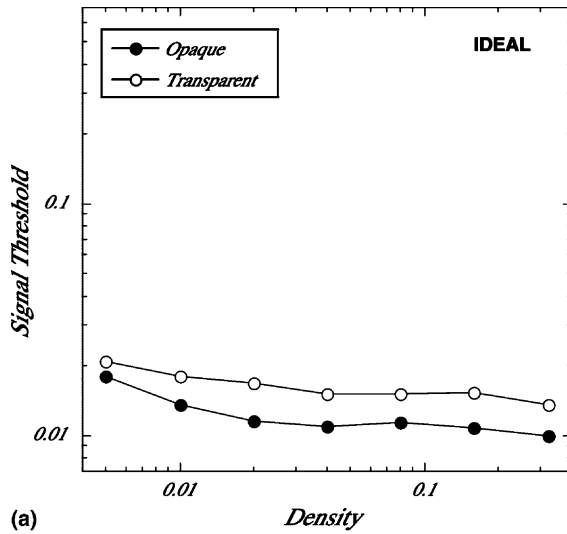



Fig. 5. Sensitivities for a human observer (black symbols) and the simulated ideal observer (grey symbols). A linear function gave very good fits to the data ($r^2 = 0.96$, $r^2 = 0.98$). It is clear that the slope of the line to the ideal observer data is much more steep ($\alpha = 55.9$) than that of the human data ($\alpha = 6.93$). Thresholds (θ_i and θ_h) are taken at $d' = 1$.

be seen that d' increases linearly as the proportion of signal dots is increased (and therefore as the proportion of noise dots is decreased), for both the human and ideal observers. A linear fit constrained to pass through the origin gave an excellent fit ($r^2 = 0.96$ for the human data, $r^2 = 0.98$ for the ideal data). We define the signal threshold (θ_h and θ_i) as the proportion of signal dots required for d' of 1. Note the much higher levels of noise required to limit performance of the ideal observer.

Fig. 6a plots the ideal signal thresholds as a function of the total dot density for both the opaque and transparent conditions. There are two features to these data. The first is that the ideal signal thresholds are consistently higher in the transparent condition than in the opaque condition, across the range of dot densities. This indicates that there is indeed a higher quantity of false matches in the transparent condition (which we saw in Fig. 4). The second feature to these data is that ideal performance tends to improve as dot density is increased. This is somewhat counterintuitive, as increasing dot density increases the number of possible correspondences, which will raise the value of the correlation for the noise disparities. However, increasing dot density will also increase the strength of the signal (see Fig. 4). The improvement in performance indicates that the signal strength initially improves faster than the strength of the correspondence noise (the heights of all the other peaks of the stimulus correlation), but these rates increase similarly from a dot density of around 5%. To assess this, we computed average amplitudes (across 400 trials) for transparent stimuli with a signal proportion



(a)



(b)

Fig. 6. Signal thresholds for Experiment 1 as a function of dot density (disparity ratio equals 2). (a) Ideal observer thresholds. (b) Average thresholds for three human observers. Error bars indicate the standard error of the mean across observers.

equal to 1, and a disparity ratio of 2. We then took the average of the peak amplitudes (that correspond to the two signal disparities that the ideal observer isolates with the correct template), and compared this to the average baseline amplitude (that correspond to the two signal amplitudes that the ideal observer isolates with the incorrect template). We found that the peak amplitudes actually rise faster than the base amplitudes, and this determines the improvement in ideal observer performance.

Fig. 6b plots the signal thresholds for three human observers in the opaque and transparent conditions as a function of the total dot density. The error bars are standard errors of the mean across observers. By comparison with Fig. 6a, it is clear the performance is much worse than ideal performance in both opaque and




Fig. 7. Average efficiencies for three human observers as a function of dot density (disparity ratio = 2). Error bars indicate the standard error of the mean across observers.

transparent conditions. However, similarly to the ideal data, the thresholds for the transparent depth discrimination are consistently higher than those in the opaque condition. Thus, more signal dots are required to perform depth discrimination in the transparency case at an equivalent level of performance to the opaque case. The effect of dot density on human observer performance contrasts with the ideal observer. While there is an initial improvement in performance at the low dot densities, performance declines as dot density is further increased.

Fig. 7 plots the computed efficiencies for the three observers as a function of dot density. Error bars are standard errors of the mean across observers. The efficiencies for the opaque and transparent conditions are approximately equal. The cost in performance (higher signal thresholds) for depth discrimination of transparent surfaces does not translate into a lower efficiency, but is in fact compensated for by comparing human performance with that of the ideal observer. A second aspect of these data is that efficiency decreases similarly for both the opaque and transparent conditions as dot density increases.

3.3. Discussion

In this experiment we compared human performance for depth discrimination of transparent and opaque surfaces as a function of dot density. The first issue that concerned us was whether discrimination performance is impaired for transparent stereograms. We found that the human observer's signal thresholds were higher in the transparent condition than the opaque condition. This could imply that there is an additional limitation on performance in the transparent case, such as inhibi-

tory interactions between different disparities. However, we found that ideal thresholds were also higher in the transparent condition. By computing the efficiency, we could assess the relative cost between human and ideal observer performance. Indeed, we found that efficiencies were approximately equal in the two conditions. Therefore, the limitations on ideal performance account for the limitations on human performance. In other words, false matching accounts for the higher thresholds in the transparent condition, for both the human observers and the ideal observer. The similarity in the opaque and transparent efficiencies implies that there is no additional processing limitation in recovering depth from transparent stereograms (at least at the depths tested here).

The effect of dot density confirms that this mechanism is limited by correspondence noise. Efficiency decreases similarly with increasing dot density in both opaque and transparent conditions, indicating that the human observers are increasingly impaired as the number of potential matches increases. Indeed, the maximum efficiencies are around 1%, indicating human observers use far less information than is available to perform the task, i.e. human observers use only a proportion of the available disparity samples. We can think of this limitation in terms of the stimulus cross-correlations in Fig. 4. From Fig. 4 we saw that decreasing the level of signal decreased the height of the peaks in the correlation. If human observers cannot use all of the available disparity information, these peaks will be lower than the ideal case (and so will indeed require more signal dots than the ideal observer to raise the peaks above the background correspondence noise). This result confirms the finding of Harris and Parker (1992) in which the efficiency of detecting a step-change in depth declined as the number of dots in their stereograms was increased. Similarly, Cormack, Landers, and Ramakrishnan (1997) found that the efficiency for detecting correlation in dynamic random dot stereograms decreased with increasing dot density. Moreover, we find that this effect of density is true also for depth discrimination of transparent stereograms. The similarity in the findings across these different studies is striking given the differences in stimuli, tasks, and the corresponding ideal observers. This encourages the view that all the studies are tapping into the same correspondence noise limited mechanism.

Our maximum efficiency of around 1% does contrast with maximum efficiencies of approximately 20% reported by Harris and Parker (1992) and Cormack et al. (1997). This difference can partially be attributed to the number of stimulus elements used in the studies. The minimum dot density in the previous two studies corresponded to approximately 4 visible dots. Here, the lowest dot density of 0.5% corresponds to 50 dots, and so poses a greater correspondence problem than the

minimum 4 dots of the previous studies. Therefore, it is conceivable that if we had reduced density (and therefore the number of stimulus elements) even further then the efficiencies would have followed the upward trend in that direction and approached a maximum of 20% efficiency. We noticed that at the lower densities used here the perception of a surface was very weak, but was stronger as the dot density was further increased. However, it is difficult to quantify this subjective change. We ran a short experiment in which our observers were required to indicate whether they did perceive surfaces or just noise, a task similar to that of Akerstrom and Todd (1988). We found that, over the same range of dot densities tested here, surface perception thresholds for transparency were generally higher than the opaque condition. However, we found that while this basic effect was qualitatively similar across observers, thresholds varied a lot between observers. These variations are likely to be a direct result of the subjective nature of the task, suggesting that the criterion for surface perception differs between observers (and possibly within observers over trials). Indeed, this inconsistency in the subjective results validates our use of a more objective depth discrimination task to probe the mechanisms underlying stereoscopic transparency.

4. Experiment 2: Effects of disparity

In Experiment 1 we fixed the disparities of the standard and target surfaces, resulting in a constant disparity ratio, and we varied the dot density. Here we aimed to see if equal efficiencies are found in the transparent and opaque conditions across a range of disparity ratios, keeping the dot density constant. Following both Akerstrom and Todd (1988) and Gepshtein and Cooperman (1998), we predict that performance in the transparent condition should be increasingly impaired as the disparity between the two surfaces is increased.

4.1. Methods

4.1.1. Stimuli

The stimuli were random dot stereograms as described in the General Methods section, constructed as described in Experiment 1.

4.1.2. Procedure

We presented transparent and opaque stereograms as described in the General Methods section. Here we presented transparent and opaque random dot stimuli at a range of disparity ratios. We fixed the standard disparity to 0.15° ($9'$), and used five target disparities of 0.30° , 0.45° , 0.60° , 0.75° and 0.90° , giving disparity ratios of 2, 3, 4, 5 and 6. We used a fixed dot density of

0.05. To limit performance, we presented the signals in a number of noise levels by the method of constant stimuli. We tested five noise levels per condition and measured d' for each noise level we tested. In both the transparent and opaque conditions each observer completed 20 practice trials with 0% noise to become familiar with the stimulus before beginning a session for a new condition. There were equal numbers of near further and far further trials. Each condition was blocked, with 40 trials per each noise condition (20 near-further, 20 far-further). Within each condition, trials for different noise levels were randomly interleaved.

4.1.3. Ideal observer

The ideal observer for this task was identical to that described in Experiment 1 in detail. The quantity of matches of a given disparity is given by the cross-correlation of the left and right images. This is then compared with templates, by correlation. The templates used by the ideal observer described the two possible disparity combinations (the location of the peaks in the templates) for a given condition of disparity ratio. The ideal observer then selects the template with the highest correlation, a maximum likelihood decision rule.

4.2. Results

Fig. 8a plots the ideal signal thresholds as a function of disparity ratio for both the transparent (open symbols) and opaque (filled symbols) conditions. There are two features to these data. The first is that the ideal signal thresholds are consistently higher in the transparent condition than in the opaque condition, across the range of disparity ratios. The second feature to these data is that ideal performance is constant across the disparity ratios. Indeed, there is no reason to expect an effect of increasing the difference in disparity between the standard and target surfaces. This simply changes the location of the peaks of the disparity correlations. The only limitation on ideal performance is the disparity noise.

Fig. 8b plots the signal thresholds for three observers as a function of disparity ratio for both the transparent (open symbols) and opaque (filled symbols) conditions. As in Experiment 1, error bars are standard errors of the mean across observers. Again, there are two features to these data. The first is that transparent thresholds are consistently higher than opaque thresholds. The second feature is that there is little effect of disparity ratio, thresholds are more or less constant across the range of disparities tested. There does appear to be a trend for thresholds to increase with disparity under the transparent condition associated with an increase in variability, and indeed this tendency occurs in only one of our three observers.

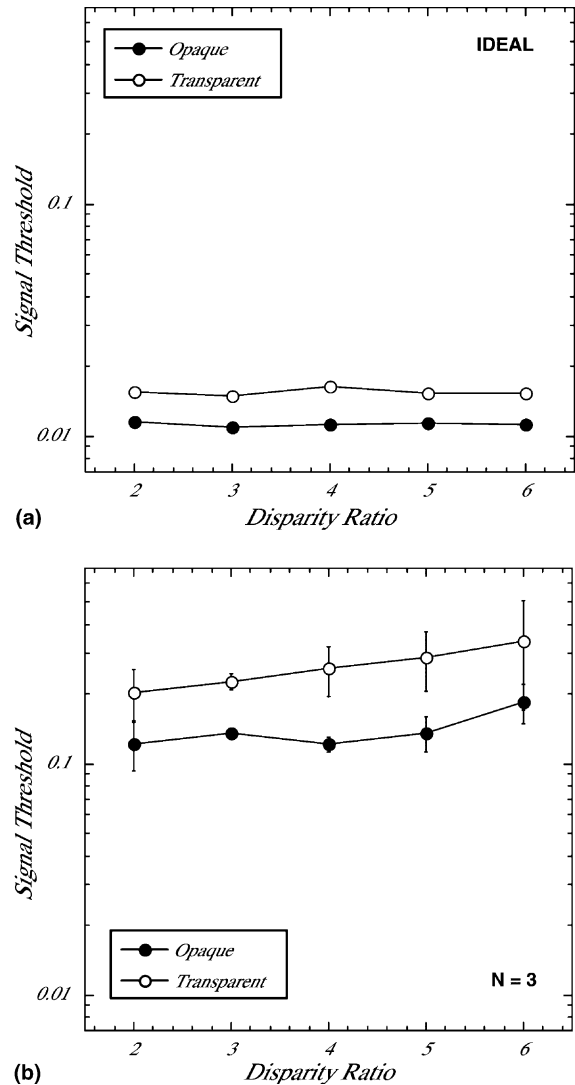


Fig. 8. Signal thresholds for Experiment 2 as a function of disparity ratio (dot density equals 0.05). (a) Ideal observer thresholds. (b) Average thresholds for three human observers. Error bars indicate standard errors of the mean across observers.

Fig. 9 plots the computed efficiencies for the three observers as a function of disparity ratio. Error bars are standard errors of the mean across observers. The efficiencies are similar for the opaque and transparent conditions across the range of disparity ratios. The cost in performance (higher signal thresholds) for depth discrimination of transparent surfaces does not translate into a lower efficiency. Efficiencies are constant across the disparity ratios for both the opaque and transparent conditions, and similar in amplitude across conditions.

4.3. Discussion

In this second experiment we compared human performance for depth discrimination of transparent and opaque surfaces as a function of disparity ratio. This




Fig. 9. Average efficiencies for three human observers as a function of disparity ratio (dot density = 0.05). Error bars indicate standard errors of the mean across observers.

addressed the question of whether discrimination performance is impaired for transparent stereograms across a range of disparity ratios, and whether the disparity ratio has an effect on depth discrimination. We found that signal thresholds are consistently higher in the transparent condition than the opaque condition, for both the human observers and the ideal observer, across a threefold range of disparity ratios. We also found that efficiencies were approximately equal in the two conditions across the range of disparity ratios. This confirms the finding of Experiment 1, false matching accounts for the cost in the transparent condition. However, we found that there is no effect of disparity ratio on depth discrimination of transparent or single opaque surfaces. This is in contrast to the findings of Akerstrom and Todd (1988) who found that increasing the disparity difference between transparent surfaces impaired perceived transparency. Our results are also in contrast to the findings of Gepshtein and Cooperman (1998) who found that the limiting density to discriminate an oriented cylinder behind a transparent plane decreased as the depth between the surfaces was increased, which they termed the ‘farther worse’ effect. The difference between the present study and the Akerstrom and Todd (1988) study may be due to the disparities used. Here we fixed a standard disparity at $\pm 9'$ and increased a target disparity in steps up to $\pm 54'$. In contrast, Akerstrom and Todd (1988) used a minimum difference of $\pm 7'$ and $\pm 21'$ up to a maximum of $\pm 49'$ and $\pm 63'$ (although the exact disparities varied across observers). Therefore the largest absolute disparity in their study was $112'$, while here it is $63'$. It is possible that the effect of disparity on stereo-transparency found by Akerstrom and Todd (1988) is due to a problem in fusing the two-planes simultaneously. Indeed, it was noted by Akerstrom and

Todd (1988) that their observers found they had to make a considerable effort to see the two surfaces in their stereograms, even over long presentation times (up to 35 s), suggesting the need for vergence eye movements. In contrast, here observers were instructed to fixate on a zero disparity cross and could perceive transparency at a relatively short duration.

Effects of disparity on surface perception have been attributed to inhibitory interactions at the level of surface representations. This was suggested by the Gepshtein and Cooperman (1998) study, in which the ‘farther worse’ effect persisted when the two surfaces were defined by opposite polarities, although the overall magnitude of the effect was less than the same polarity condition. Indeed, this parallels Akerstrom and Todd’s (1988) finding that perceived transparency was impaired by increasing the disparity difference between chromatically defined surfaces, but to a lesser extent than a single colour condition. There is evidence for inhibitory interactions in disparity tuning, though not specifically at the level of a surface representation. Specifically, Stevenson, Cormack, and Schor (1991) found that adapting to a particular disparity resulted in a threshold elevation in the disparity sensitivity function, and Cormack, Stevenson, and Schor (1993) found that correlation thresholds for a given disparity were raised by the presence of a different disparity. The disparity tuning functions derived from these studies were very similar (see Cormack et al., 1993), with clear inhibitory regions. Stevenson, Cormack, and Schor (1992) showed that their tuning functions could be modelled by a number of narrowly tuned disparity channels with inhibitory lobes (a centre-surround receptive field), but did not rule out a mutual inhibition between disparity-tuned channels. The lack of an effect of disparity in the present study suggests that the range of disparities we used were beyond the range of any inhibitory interactions, and so favours an account of disparity domain inhibition in terms of narrowly tuned disparity channels with inhibitory lobes, rather than a mutual inhibition between disparity channels, or disparity defined surfaces.

5. General discussion

5.1. Summary of results

In this study we have computed the efficiency for depth discrimination of transparent and similar opaque random dot stereograms. The advantage of our approach was twofold. Our objective method not only gives a more reliable estimate of perceptual performance free of subjective criteria, but the efficiency measure allows the experimenter to normalize that performance to the information available in the stimulus. An efficiency experiment thereby allows us to compare performance

across observers (because it is objective) and across task (because performance is normalized to absolute performance). In Experiment 1 we found that the efficiencies were approximately equal for the transparent and opaque conditions. This demonstrated that the higher thresholds in transparency can be accounted for by a greater incidence of false matches in that condition, and do not necessarily imply inhibitory interactions specific to surface overlap (Akerstrom & Todd, 1988; Gepshtein & Cooperman, 1998). The very low efficiencies we found, of around 1% or less, in both the opaque and transparent conditions suggest there is a problem in using all the available signal information. In support of this we also found that increasing the dot density, thus increasing the number of possible correspondences, decreased the efficiency in both conditions. The effect of density is therefore not, as has been suggested (Akerstrom & Todd, 1988), a behaviour unique to transparent configurations. Furthermore, the similarity in the effect of density implies a common mechanism, rather than for example an inhibitory mechanism specifically sensitive to configurations of overlapping transparent depth planes. In Experiment 2, we found that the efficiencies in the opaque and transparent condition were approximately equal across a range of disparity differences, supporting the finding of Experiment 1. In addition, we found that there was no effect of disparity ratio on stereoscopic transparency. This contrasts with other studies that have found an effect of disparity on transparency, attributed to mutual inhibition between disparity detectors or disparity defined surface representations (Akerstrom & Todd, 1988; Gepshtein & Cooperman, 1998). Task and stimulus differences may underlie this inconsistency. Our findings suggest that if there are inhibitory interactions in the disparity domain, they are probably restricted to a small range around the preferred disparity. We consider further implications of these results in the following sections.

5.2. Correspondence noise limitations

The low efficiencies we find suggest that human observers are unable to use most of the available disparity information to perform depth judgments. This supports previous findings of low efficiencies for other stereo tasks (Cormack et al., 1997, 1994; Harris & Parker, 1992). Both Harris and Parker (1992) and Cormack et al. (1997) found efficiencies of 20% or less and, as we found here, their efficiencies declined as the dot density of their stereograms was increased. We saw that increasing the dot density increased the level of false matches in the stimulus (see Fig. 4), thus creating a greater correspondence problem. Therefore, these results suggest that the mechanisms of stereopsis are limited by correspondence noise, i.e. the greater the correspondence problem the less effective the system is

at solving it. We provide some suggestions for the mechanisms underlying this behaviour in the following sections.

5.3. Similarities with motion mechanisms

The present work uses a methodology comparable to that of another study we have conducted on motion transparency (Wallace & Mamassian, 2003). In that study, we presented random dot kinematograms in which randomly placed dots were displaced to the right or left by a particular amount on subsequent frames. In a transparent condition, both directions of motion were presented simultaneously, while in a coherent condition the directions of motion were presented sequentially (comparable to the opaque condition of the present study). The task was to decide on the direction of the fastest moving surface ('left' or 'right'), in a way analogous to the depth discrimination of the present study. Performance was limited by varying the number of dots allocated to the moving surfaces, while the remaining 'noise' dots were randomly placed on each frame. The ideal observer for this speed discrimination task was therefore similar to the ideal observer for the depth discrimination task, cross-correlating subsequent frames of the motion stimulus rather than between the eyes, and performing a maximum likelihood decision rule by template matching. In an effort to equate the tasks further, we were careful to maintain as similar parameters as possible between the motion and the present stereo experiment to facilitate comparison between the studies. We used a similar projected dot and stimulus sizes, and used an identical range of dot densities and speed/disparity ratios.

Similarly to the results presented here, we found an effect of correspondence noise on coherent and transparent motion, the efficiencies declining as dot density was increased. However, the maximum efficiencies for the motion study were considerably higher, around 10% compared to 1% here. This suggests that motion mechanisms succeed in maintaining a better signal-to-noise ratio than stereo mechanisms. In the following section we suggest that this improvement is due to a suppressive interaction effectively reducing the correspondence noise. Moreover, our motion stimulus was inherently dynamic, consisting of 10 frames of uniform displacements, compared to the one-shot presentation of the binocular images here. The ideal observer is essentially identical in the motion and stereo cases, correlating consecutive pairs of frames, and so the improved efficiency may also indicate that the human motion system is in fact integrating over a longer period than a pair of frames. There was an informational limit in the transparent motion condition (as indicated by ideal observer performance), but in contrast to the present study we found there was a residual cost for processing trans-

parent motion (indicated by higher efficiencies in that condition than the coherent condition). This residual cost was present across the range of dot densities and disparity ratios we tested.

The residual cost for transparent motion is consistent with a range of psychophysical evidence and the modified motion energy model developed by Qian, Andersen, and Adelson (1994b). The modified motion energy model was proposed to account for a series of psychophysical and physiological findings (Qian & Andersen, 1994; Qian, Andersen, & Adelson, 1994a). In these studies, random dot kinematograms were constructed in which pairs of dots moved in opposite directions in close spatial proximity. These stimuli were perceived as ‘flicker’. However, when the dots were unpaired such that they no longer move in close spatial proximity, transparency was perceived. Similarly, when paired dots were presented such that each dot had a different disparity, the previous ‘flicker’ percept was abolished and observers could segregate the two planes of motion. To account for this result, Qian et al. (1994b) introduced disparity selectivity into the filters of the motion energy model, and restricted the opponent motion inhibition within disparity-tuned cells (and within a small spatial region hypothesised to correspond to the size of MT ‘subunits’). There is no inhibitory interaction between the disparity-tuned channels in this model. The findings of our motion and stereo studies are entirely consistent with this model, evidencing inhibition between opposite directions of motion but no inhibition between different disparities (but still consistent with a centre-surround disparity tuning). Moreover, our study demonstrates that only a fraction of the information is used by the human visual system, and it provides quantitative estimates of these fractions for motion and stereo transparency.

5.4. Neural substrate

As described in the introduction, traditional stereo algorithms that employ the uniqueness and continuity constraints of Marr and Poggio (1976, 1979) will be unable to recover stereoscopic transparency, as they do not permit the occurrence of more than one disparity at a given visual location. More recently a range of computational models have been proposed that pass the test of transparency to varying degrees of success (Gray, Pouget, Zemel, Nowlan, & Sejnowski, 1998; Pollard et al., 1985; Prazdny, 1985; Read, 2002; Tsai & Victor, 2003). The later of these models incorporate physiological constraints of the underlying mechanisms (Anzai, Ohzawa, & Freeman, 1999a, 1999b, 1999c; DeAngelis, Ohzawa, & Freeman, 1991; Freeman & Ohzawa, 1990; Ohzawa, DeAngelis, & Freeman, 1990, 1996, 1997), understood to compute a ‘disparity-energy’ (Ohzawa, 1998; Qian, 1994). These models differ from

traditional accounts of disparity processing, in that the mechanisms do not solve the correspondence problem directly for individual dots, but rather through phase and/or position shifts between the spatial frequency tuned receptive fields for the left and right eyes.

The modified motion energy model of Qian et al. (1994a, 1994b) provides an account of motion processing up to the level of area MT. Indeed, Qian and Andersen (1994) found that the modulation of MT activity in response to paired and unpaired random dot displays correlated with a change in the perceived transparency of these displays. Specifically, the responses to paired displays were suppressed compared to paired displays of opposite motions. What is more interesting in terms of the present study, was their finding that MT responses to random noise (of the kind used here) were small, similar to those to paired dot patterns. This suggested that the suppression of opposite motion signals serves to combat the unwanted effects of correspondence noise. The modified motion energy model also includes disparity selectivity, but there is no suppressive interaction between dissimilar disparities. We suggest that this difference can account for the higher efficiencies in the motion study i.e. the suppressive interactions between motion results in a weaker residual noise response than in the disparity case, where there is no suppressive interaction between different disparities. As we previously suggested, our finding of a lack of residual cost for stereotransparency does not necessarily rule out inhibitory interactions in the disparity domain, if the inhibition is restricted to a narrow range around a central disparity-tuned region. Recent physiological evidence finds that disparity and motion selective MT neurons do possess inhibitory surround regions but for the same disparity as the central excitatory region, which could facilitate the segregation of the image into different regions or surfaces as the response is maximal when surround stimulation is different from center selectivity (Bradley & Andersen, 1998). The question then remains, how do disparity selective mechanisms combat correspondence noise, if not by a mechanism of mutual inhibition? One possibility is by spatial pooling. In the same way that MT spatial pooling can serve to combat the motion correspondence problem (Barlow & Tripathy, 1997), it may also serve to combat the stereo correspondence problem. Furthermore, such a pooling operation may account for the low efficiencies, as it would effectively reduce the quantity of disparity samples used to perform the task.

6. Conclusions

The present study provided quantitative estimates of the efficiency of human observers in a stereo transparency task. We found very small efficiencies suggesting

that the stereo transparency mechanism is significantly impaired by the correspondence problem. We found no evidence for disparity inhibition as there were no efficiency difference between transparency and opaque conditions. These results contrasts with motion mechanisms which appear to use inhibitory mechanisms to combat correspondence noise. However, a spatial pooling of disparity information may serve to combat correspondence noise to some extent, at the expense of discarding some segregation information and thus decreasing even further the efficiency of the system.

Acknowledgements

Part of this work was presented at the annual meeting of the Vision Sciences Society in May 2002. This study was supported by an EPSRC quota fellowship (JMW), EPSRC research grant GR/R57157/01 (PM) and the EC research training network HPRN-CT-2002-00226.

Appendix A. Efficiency

We provide here the steps in the derivation of our efficiency measure. Efficiency is defined as the ratio of human sensitivity to that of the ideal observer (Barlow, 1978; Tanner & Birdsall, 1958):

$$F = \left(\frac{d'_h}{d'_i} \right)^2 \quad (\text{A.1})$$

The problem in using this definition is that the ideal observer easily reaches ceiling performance for a suitable range as for the human observer. We find experimentally that d' is proportional to the signal:

$$d'_h(s) = \alpha_h \cdot p(s) \quad (\text{A.2})$$

$$d'_i(s) = \alpha_i \cdot p(s) \quad (\text{A.3})$$

And so, following Harris and Parker (1992):

$$F = \left(\frac{\alpha_h}{\alpha_i} \right)^2 \quad (\text{A.4})$$

In the present experiment, we take our thresholds (θ_h and θ_i) at a d' equal to 1. Thus:

$$\alpha_h = \frac{1}{\theta_h} \quad (\text{A.5})$$

$$\alpha_i = \frac{1}{\theta_i} \quad (\text{A.6})$$

Therefore, substituting Eqs. (A.5) and (A.6) into Eq. (A.4):

$$F = \left(\frac{\theta_i}{\theta_h} \right)^2 \quad (\text{A.7})$$

References

- Akerstrom, R. A., & Todd, J. T. (1988). The Perception of Stereoscopic Transparency. *Perception & Psychophysics*, 44(5), 421–432.
- Anzai, A., Ohzawa, I., & Freeman, R. D. (1999a). Neural mechanisms for encoding binocular disparity: receptive field position versus phase. *Journal of Neurophysiology*, 82, 874–890.
- Anzai, A., Ohzawa, I., & Freeman, R. (1999b). Neural mechanisms for processing binocular information. I. Simple cells. *Journal of Neurophysiology*, 82, 891–908.
- Anzai, A., Ohzawa, I., & Freeman, R. (1999c). Neural mechanisms for processing binocular information. II. Complex cells. *Journal of Neurophysiology*, 82, 909–924.
- Barlow, H. (1978). The efficiency of detecting changes of density in random dot patterns. *Vision Research*, 18, 637–650.
- Barlow, H., & Tripathy, S. P. (1997). Correspondence noise and signal pooling in the detection of coherent visual motion. *Journal of Neuroscience*, 17(20), 7954–7966.
- Bradley, D. C., & Andersen, R. A. (1998). Center-surround antagonism based on disparity in primate area MT. *Journal of Neuroscience*, 18(18), 7552–7565.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433–436.
- Burt, P., & Julesz, B. (1980). A disparity gradient limit for binocular fusion. *Science*, 208, 615–617.
- Cormack, L. K., Landers, D. D., & Ramakrishnan, S. (1997). Element density and the efficiency of binocular matching. *Journal of the Optical Society of America A—Optics Image Science and Vision*, 14(4), 723–730.
- Cormack, L. K., Stevenson, S. B., & Schor, C. M. (1991). Interocular correlation, luminance contrast and cyclopean processing. *Vision Research*, 31(12), 2195–2207.
- Cormack, L. K., Stevenson, S. B., & Schor, C. M. (1993). Disparity-tuned channels of the human visual system. *Visual Neuroscience*, 10, 585–596.
- Cormack, L. K., Stevenson, S. B., & Schor, C. M. (1994). An upper limit to the binocular combination of stimuli. *Vision Research*, 34(19), 2599–2608.
- DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1991). Depth is encoded in the visual cortex by a specialized receptive field structure. *Nature*, 352, 156–159.
- Freeman, R. D., & Ohzawa, I. (1990). On the neurophysiological organization of binocular vision. *Vision Research*, 30, 1661–1676.
- Gepshtein, S., & Cooperman, A. (1998). Stereoscopic transparency: a test for binocular vision's disambiguating power. *Vision Research*, 38(19), 2913–2932.
- Gray, M. S., Pouget, A., Zemel, R. S., Nowlan, S. J., & Sejnowski, T. J. (1998). Reliable disparity information through selective integration. *Visual Neuroscience*, 15, 511–528.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Grimson, W. E. L. (1985). Computational experiments with a feature based stereo algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(1), 17–34.
- Harris, J. M., & Parker, A. J. (1992). Efficiency of stereopsis in random-dot stereograms. *Journal of the Optical Society of America A—Optics Image Science and Vision*, 9(1), 14–24.
- Julesz, B. (1964). Binocular depth perception without familiarity cues. *Science*, 145, 356–362.
- Kaufman, L., Bacon, J., & Barraso, F. (1973). Stereopsis without image segregation. *Vision Research*, 13, 137–147.
- 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.

- Ohzawa, I. (1998). Mechanisms of stereoscopic vision: the disparity energy model. *Current Opinion in Neurobiology*, 8, 509–515.
- Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1990). 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.
- Parker, A. J., & Yang, Y. (1989). Spatial properties of disparity pooling in human stereo vision. *Vision Research*, 29(11), 1525–1538.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Pollard, S. B., & Frisby, J. P. (1990). Transparency and the uniqueness constraint in human and computer stereo vision. *Nature*, 347(6293), 553–556.
- Pollard, S. B., Mayhew, J. E. W., & Frisby, J. P. (1985). Pmf—a stereo correspondence algorithm using a disparity gradient limit. *Perception*, 14(4), 449–470.
- Prazdny, K. (1985). Detection of binocular disparities. *Biological Cybernetics*, 52, 93–99.
- Qian, N. (1994). Computing stereo disparity and motion with known binocular cell properties. *Neural Computation*, 6, 390–404.
- Qian, N., & Andersen, R. A. (1994). Transparent motion perception as detection of unbalanced motion signals. II. Physiology. *The Journal of Neuroscience*, 14(12), 7367–7380.
- Qian, N., Andersen, R. A., & Adelson, E. H. (1994a). Transparent motion perception as detection of unbalanced motion signals. I. Psychophysics. *The 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. *The Journal of Neuroscience*, 14(12), 7381–7392.
- Qian, N., & Zhu, Y. (1997). Physiological computation of binocular disparity. *Vision Research*, 37, 1811–1827.
- Read, J. C. A. (2002). A Bayesian model of stereopsis and motion direction discrimination. *Biological Cybernetics*, 86, 117–136.
- Stevenson, S. B., Cormack, L. K., & Schor, C. M. (1989). Hyperacuity, superresolution and gap resolution in human stereopsis. *Vision Research*, 29(11), 1597–1605.
- Stevenson, S. B., Cormack, L. K., & Schor, C. M. (1991). Depth attraction and repulsion in random dot stereograms. *Vision Research*, 31(5), 805–813.
- Stevenson, S. B., Cormack, L. K., & Schor, C. M. (1992). Disparity tuning in mechanisms of human stereopsis. *Vision Research*, 32, 1685–1694.
- Tanner, W. P., & Birdsall, T. G. (1958). Definitions of d' and h as psychophysical measures. *The Journal of the Acoustical Society of America*, 30(10), 922–928.
- Tsai, J. T., & Victor, J. D. (2003). Reading a population code: a multi-scale neural model for representing binocular disparity. *Vision Research*, 43, 445–466.
- Tyler, C. (1991). Cyclopean vision. In D. Regan (Ed.), *Binocular vision* (pp. 38–74). London: MacMillan.
- Wallace, J. M., & Mamassian, P. (2003). The efficiency of speed discrimination for coherent and transparent motion. *Vision Research*, 43, 2795–2810.
- Weinshall, D. (1989). Perception of multiple transparent planes in stereo vision. *Nature*, 341, 737–739.
- Weinshall, D. (1991). Seeing “ghost” planes in stereo vision. *Vision Research*, 31(10), 1731–1748.