

Stereo Matching Precedes Dichoptic Masking

SUZANNE P. McKEE,* MARY J. BRAVO,* DOUGLAS G. TAYLOR,* GORDON E. LEGGE†

Received 15 April 1993; in revised form 7 September 1993

Stereo matching can intervene to prevent dichoptic masking. In a dichoptic masking paradigm we measured the contrast threshold for a bar target, presented to one eye, as a function of the contrast of an identical masking bar, presented at retinal correspondence in the other eye. Confirming previous studies of dichoptic masking with sinusoidal gratings, the test bar thresholds rose proportionally with increasing masking contrast. This threshold elevation was almost nullified when an extra bar was presented to the eye seeing the test stimulus. Release from masking occurred when the disparity between the masking bar and extra bar was < 20 min arc over a range of contrast levels (8–45%), and for bars containing either broad spatial frequency spectra or bars with only high spatial frequencies (peak = 12 c/deg). The latter result rules out an explanation for the release from masking based on contrast discrimination in low spatial frequency channels. The extra bar was effective in releasing the test bar from masking as long as the extra bar's contrast was greater than about one-fifth the contrast of the mask, a result that suggests that there is a contrast threshold for stereo matching. We interpret our findings to indicate that a stage of stereo matching occurs prior to the neural site limiting dichoptic contrast discrimination.

Stereopsis Binocular vision Contrast masking Retinal correspondence

INTRODUCTION

Contrast increment thresholds are commonly measured by superimposing a low contrast test stimulus on a base or pedestal contrast stimulus of identical spatial and temporal characteristics. The threshold is the smallest test contrast that can be reliably detected in the presence of the contrast pedestal. The function relating the test threshold to the pedestal contrast is dipper-shaped. When the pedestal is zero, the "absolute" contrast threshold is obtained. For near-threshold pedestal contrasts, the test threshold "dips" to values less than the absolute threshold. As the pedestal contrast rises to suprathreshold levels, the test threshold rises, revealing the nonlinear nature of contrast processing; the contrast pedestal *masks* the presence of the low contrast increment.

In monocular contrast masking, the masking contrast pedestal and the test stimulus are presented to the same eye, while the other eye is shown a uniform field. In binocular contrast masking, identical masking pedestals with superimposed test stimuli are presented to both eyes, usually at corresponding retinal loci. In dichoptic masking, the masking contrast is presented to one eye, while the test contrast is presented to the other eye at retinal correspondence with the mask. Superficially, the monocular and dichoptic procedures appear the same.

However, dichoptic thresholds are substantially higher than monocular (or binocular) thresholds measured at the same masking contrast. Using grating targets, Legge (1979, 1984b) and Levi, Harwerth and Smith (1979, 1980) found that, as the masking contrast was increased beyond three times the contrast threshold, dichoptic thresholds rose with a slope close to 1.0 on logarithmic coordinates. In comparable measurements of monocular contrast masking, thresholds rose with a slope between 0.5 and 0.7 (Legge, 1979, 1981, 1984a; Legge & Kersten, 1983). Legge (1984a, b) explained the difference between monocular and dichoptic masking by proposing that the monocular signals were squared prior to binocular combination. Clearly, squaring the test increment with the pedestal would produce a larger, more detectable, contrast change than squaring the increment separately and adding it to the squared pedestal. This quadratic summation rule was the essential feature of Legge's "binocular energy" model, a model that accounted for much of the data on binocular contrast summation.

The use of quadratic summation to account for binocular interaction has been confined to in-phase sinusoidal gratings of the same spatial frequency. Its validity has not been studied for non-sinusoidal targets, e.g. bars, and its relation to stereopsis is unknown. The passive summation of local contrast signals at retinal correspondence will not account for all aspects of binocular contrast summation. For one thing, the amount of dichoptic masking depends on the similarity of the test and masking stimulus; not surprisingly, the greatest

*Smith-Kettlewell Eye Research Institute, 2232 Webster Street, San Francisco, CA 94115, U.S.A.

†Department of Psychology, University of Minnesota, Minneapolis, MN 55455, U.S.A.

masking is found when the spatial frequency of test and mask are identical (Legge, 1979). Disparity also affects binocular masking. Several studies have found a reduction in binocular masking when the test signal was presented at a different horizontal or vertical disparity from the masking noise (Henning & Hertz, 1973, 1977; Moraglia & Schneider, 1990, 1992). These results suggest that monocular contrast signals are combined by binocular mechanisms which are tuned for spatial frequency and disparity. Certainly, there is abundant physiological evidence for cortical units with these properties (De Valois & De Valois, 1988; Poggio & Fischer, 1977; Freeman & Ohzawa, 1990; Ohzawa, DeAngelis & Freeman, 1990).

The main function of the human binocular system is to combine the monocular images, assigning a unique position in depth to each visible feature. This problem of identifying which features in the monocular half-images belong together is often called the binocular correspondence problem. The existence of cortical units tuned for disparity or other specific stimulus attributes does not guarantee a unique depth map. Because of the considerable overlap in their tuning for position, spatial frequency and disparity, many different units are stimulated by any target presented at any position and disparity. How does the brain identify the best stereo match, given many alternatives? In the last two decades, a number of computational models have been proposed to explain how the human brain, or a computer, might perform stereo matching (Marr, 1982; Grimson, 1981; Frisby & Pollard, 1991). Generally, these matching models consist of two stages—a preliminary stage in which all potential matches between compatible elements are identified, and a second stage in which the best possible match for each element is selected and alternative matches are eliminated. Plausible constraints, based on psychological observations as well as the physical characteristics of natural surfaces, are imposed on both stages to determine the best match and resolve matching ambiguities.

Psychophysical data on binocular contrast summation have not been incorporated into computational models of stereo matching. This neglect of contrast information is curious, because there is a strong empirical relationship between stereopsis and contrast. The direction of disparity (crossed or uncrossed) can be correctly identified near contrast threshold (Mansfield & Simmons, 1989; Smallman & MacLeod, 1994), and stereoacuity improves with increasing contrast (Halpern & Blake, 1988; Heckmann & Schor, 1989; Legge & Gu, 1989). However, contrast may limit the precision of stereo signals at some early stage of processing, prior to the matching operations. Binocular contrast summation could occur at a neural stage where any pair of monocular contrast signals from the two half-images, falling within a specified orientation, spatial frequency bandwidth and disparity range are combined passively—a stage that conforms roughly to the first stage of the stereo matching models. If this were true, then a low contrast target in one eye would be inescapably masked

by a similar high contrast target at the corresponding locus in the other eye, no matter what processing occurred after this first stage. However, it is possible that binocular contrast summation is contingent on the outcome of the stereo matching operations—that the contrasts of the matching elements are summed only *after* the best match is selected. In this case, whether a low contrast target would be masked by high contrast target at correspondence would depend on the alternative matches lying within the vicinity of the pair.

To study this question, we measured dichoptic masking using the bar targets sketched in Fig. 1. The upper two boxes show the stereo half-images used for a straightforward measurement of dichoptic masking; the bright masking bar (C) in the left eye is presented at retinal correspondence with a dim test bar (ΔC) of identical dimensions in the right eye. The incremental contrast of the test bar is increased until the observer can

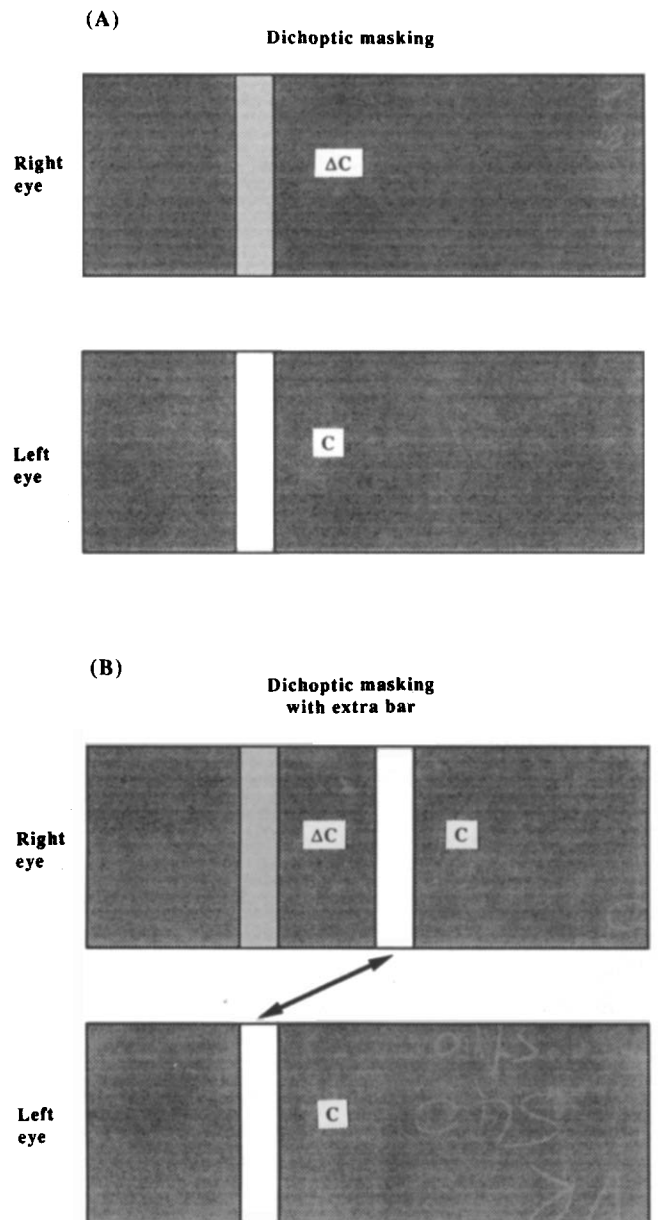


FIGURE 1. Bar targets used to measure dichoptic masking alone (A), and dichoptic masking with extra bar presented adjacent to test bar (B).




FIGURE 2. Diagram of stimulus configuration used for this study. Targets were bright bars superimposed on a uniform field, 1.4×0.76 deg. Nonius lines were continuously visible when uniform fields were illuminated. Bar targets were presented for 200 msec when subject signaled that nonius lines appeared aligned. Disparity measured between centers of targets.

reliably detect its presence. The lower two boxes in Fig. 1 show the stereo half-images for a slightly different masking configuration; an extra bar, equal in contrast to the masking bar, is now presented adjacent to the test bar in the right eye. Since the extra bar is, in all respects, identical to the masking bar in the left eye, it seems likely that at small disparities these two bars will be matched, and that the observer will see a single bright bar at a depth behind the fixation plane. If this match between the mask and the extra bar occurs, does the contrast threshold for detecting the test bar change? Despite the match with the extra bar, the masking bar may still elevate the contrast threshold of the test bar at retinal correspondence. Alternatively, by providing a suitable match, the extra bar may release the test bar from the dichoptic masking effects on the bar in the left eye, so the test threshold would be lower in the presence of the extra bar. In short, does dichoptic masking precede or follow stereo matching?

METHODS

The stimulus configuration

For most of these experiments, the targets were bright vertical bars, 6 min arc wide and 46 min arc long, presented on rectangular uniform fields, 1.4 deg wide by 46 min arc in height. Figure 2 shows a typical experimental configuration. The test bar was presented to the right eye in the center of the uniform field, the dichoptic masking bar was presented to the left eye at retinal correspondence with the test bar, and the extra bar,

when presented, was shown to the right eye, on the right side* of the test bar, at an uncrossed disparity with respect to the masking bar. In one experiment, the disparity between the mask and extra bars was varied, and to accommodate the most extreme disparity (45 min arc), the test and mask bars were displaced together 20 min arc to the left of center. For all experiments, short black nonius lines were superimposed on the upper and lower edges of the uniform fields, 10 min arc to the left of center. The nonius lines were actually pieces of black chart tape, about 6 min arc in width, that had been aligned visually with the edge of a bright guide line, and applied to the face of each monitor; thus, they were visible whenever the uniform background was illuminated. The location of each monocularly-presented feature was defined with respect to the physical location of each nonius line. Therefore, if the nonius lines were perfectly aligned (specifying the same horizontal direction in the two eyes), a monocular feature, presented a small distance from the nonius line in one eye, would occupy the corresponding retinal location of a monocular feature presented to other eye at the same distance from the nonius line viewed by the other eye. Equivalently, these two monocular features would have zero horizontal disparity with respect to the fixation plane.

We also used luminance-balanced triphasic targets to confine target energy to a high spatial frequency bandwidth; these triphasic targets consisted of bright vertical bars, 2.8 min arc wide, flanked by two thin dark bars, each 1.4 min arc wide. For physically-balanced stimuli, the integrated luminance of center and surround is equal to the uniform background, but there is no guarantee that physically-balanced stimuli are balanced for the human visual system. For example, the response to a luminance decrement may not equal the response to a luminance increment due to early non-linearities in

*Some thresholds were also measured with the extra bar on the left of the test bar at an uncrossed disparity with respect to the masking bar. The thresholds were identical to those measured with the extra bar on the right.

luminance processing. A non-linearity of this type would introduce a small contrast signal at low spatial frequencies, i.e. a d.c. component (Burton, 1973). To insure that the center and flanks were psychophysically-balanced, we measured detectability (d') for a fixed contrast difference as we systematically varied the luminance ratio of the center to the flanks. Since contrast sensitivity falls monotonically with increasing spatial frequency beyond the peak at 2–3 c/deg, we assumed that minimum detectability would correspond to the ratio that best isolated a high spatial frequency bandwidth.

We found that the luminance ratio that produced minimum detectability for the triphasic target was roughly constant over the contrast range which we could test (2.8–7%). In agreement with the results from Bravo and Blake (1992), the dark flanks were lower in luminance at the ratio that produced minimum detectability than at the ratio that produced physical balance. When the triphasic stimuli were physically balanced for luminance, the ratio of the luminance increment to the decrement was about 60%, but in the case of the psychophysically-balanced stimuli the ratio was closer to 50%. The psychophysically-balanced triphasic targets were invisible at long viewing distances (>4 m), even at contrasts as high as the 31% used in our experiments—further evidence of their high spatial-frequency bandwidth. We measured the luminance distribution of the psychophysically-balanced triphasic targets using the 1 min arc narrow slit of a Pritchard photometer, and performed a Fourier transform of the distribution to estimate the amplitude of the signal as a function of spatial frequency. Figure 3 shows the amplitude spectrum of the psychophysically-balanced triphasic target; the d.c. component is due to the low luminance of the flanks required to produce the

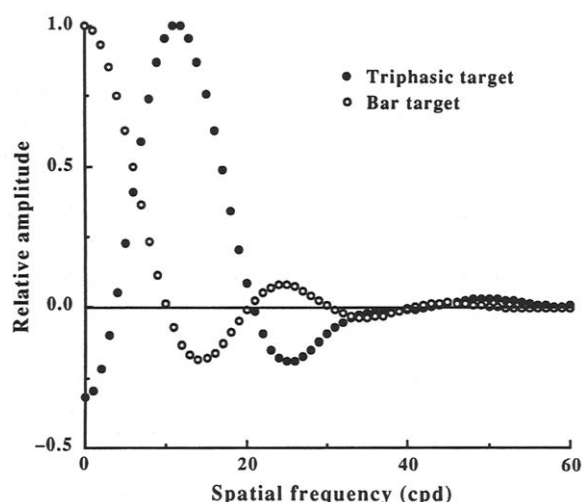


FIGURE 3. Amplitude spectrum of bars and triphasic high spatial frequency targets used in this study. The triphasic target consisted of a bright central bar, 2.8 min arc wide, flanked by dark bars 1.4 min arc wide. Target luminance was balanced psychophysically by varying the ratio of the luminance of central and flanking bars to produce minimum detectability at a given contrast. Flanks were lower in luminance at the ratio that produced minimum detectability than at the ratio that produced physical balance, as shown by the negative d.c. component in the spectrum.

minimum detectability. The peak spatial frequency of these targets is about 12 c/deg. For comparison, the amplitude spectrum of the bar targets is also shown in Fig. 3.

The targets were drawn by computer-generated signals on the screens of two Hewlett-Packard 1332A monitors, each equipped with a P4 phosphor. A custom-designed function generator was used to convert the point-addressable X - Y display monitors into a raster-scan mode. A high-speed (125 kHz) triangle wave was fed into the Y -axis to create a vertical raster which was moved horizontally by a slower sawtooth ramp (244 Hz) fed into the X -axis. There were 1024 lines drawn within a small region in the center of each monitor. At the 1.5 m viewing distance, the center-to-center distance between each raster line was 7 sec arc. The luminance of each raster line could be set independently by varying the Z -axis poke values. The Z -axis had 8 bits of resolution, and the gain could be changed to achieve a larger or smaller luminance range as needed for greater flexibility in measuring contrast thresholds. For most experimental conditions, the stimulus bars were filled with raster lines of the same poke value, creating a physically-uniform bar, but for some test conditions, the poke values did not allow sufficient luminance resolution. In these cases, smaller steps in luminance were achieved by regularly alternating between two poke values for adjacent raster lines to create an apparently uniform field a half-step between the two luminance values. The images on the two monitors were superimposed by a beam-splitting pellicle; orthogonally-oriented polarizers were placed in front of the monitors and the observer's eyes to guarantee that the features on only one screen, including the uniform field and the appropriate nonius line, were visible to each eye. The polarizers in front of the subject's eyes were mounted in a fixed frame attached to a table, so that changes in head orientation did not affect the degree of isolation of left and right images. When observers viewed the stimuli presented on one screen through crossed polarizers, no light from that screen was visible. "Ringing" and other non-uniformities at the edge of the background fields were obscured by placing a black mask with a rectangular aperture in front of each screen; the rectangular apertures were aligned visually by viewing them through the pellicle.

The luminances for each Z -axis poke value and half-step value were measured with a Pritchard photometer, using a 6 min arc probe centered on a vertical bar, 12 min arc wide, that was presented in the center of the uniform background. The luminance of the background field on each screen was 6.6 cd/m² when viewed through the pellicle and polaroids. The luminance measurements for the bar targets were converted into Michelson contrast according to the formula:

$$\text{contrast} = \frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}}$$

Overhead fluorescent lighting located about 2 m from the CRT monitors supplied indirect illumination of the experimental setting at a moderate photopic level. Room

furniture and experimental equipment were clearly visible. All luminance measurements were made with the room lights on, under the same conditions used for the experiments.

Procedure

In the context of this study, it was essential that the masking bar and the test bar were presented at retinal correspondence, so at the beginning of each trial, the observer was shown only the uniform fields with the superimposed nonius lines. When the nonius lines appeared to be aligned, the observer pressed a button to signal the presentation of the test stimulus. The test, mask, and extra bars were then presented for 200 msec, a duration somewhat longer than convergence latency, but certainly too short to permit the completion of a voluntary convergence movement during a test trial. In a pilot study, we tried using a duration of 150 msec; the results for this shorter duration showed the same pattern as the results presented here, but absolute thresholds were elevated constraining the range of possible contrasts. After 200 msec, the bars disappeared, but the uniform field remained visible until the observer responded. After the response of the observer, or after a 2 sec wait period, the uniform field was briefly extinguished for 100 msec, and then reappeared. The brief extinction signalled the beginning of a new test trial. Following this extinction, observers were required to wait 800 msec before initiating a new trial, to preserve a stable state of contrast and/or luminance adaptation. If the observer detected a large shift in the nonius lines during a trial, they did not respond to the presentation. Usually, observers noted only small random movements, less than half the width of the nonius lines, as would be expected from convergence instability.

For all experiments, we made the same measurements—the contrast threshold for the test bar presented to the right eye in the center of the uniform field. We used a “yes-no” procedure to determine contrast threshold, rather than a two-alternative force-choice procedure with two temporal intervals, because we thought observers might initiate a convergence movement prior to or during the second temporal interval. On each trial, we presented either a test bar at some faint test contrast or a blank. Subjects indicated by pressing one of two buttons whether they thought the test bar was present. Feedback was provided. Each test contrast level was presented in a separate block of 80 trials with equal probability of test bar being present or absent on each trial. The hits and false alarms from the 80 trials were used to estimate d' (Elliott, 1964). A minimum of three contrast levels (240 trials total) were used to determine the contrast corresponding to a threshold criterion of $d = 1.0$. This threshold criterion was estimated by fitting a straight line (least squares criterion) to the data points plotted in $\log d'$ vs \log contrast coordinates (see Foley & Legge, 1981). For some conditions, we made repeated measurements, each based on at least 240 trials, to evaluate threshold variability; the error shown in the tables is the standard error from three or more measure-

ments, or the sample standard deviation of two measurements.

For most of our measurements, we found that a straight line produced an excellent fit to the data ($r^2 \geq 0.9$). The exceptions were the dichoptic masking thresholds (mask and test bar alone without the extra bar). Some of the dichoptic functions were non-linear with plateaus between rising segments. Three of the four observers (MB, SM, DGT) reported distinct cues which may correlate with the different segments of the functions. We will analyze the fine structure of the dichoptic psychometric functions in future research. For the time being, we collected additional data so that straight-line fits were moderately good, enabling us to obtain conservative estimates of the contrast yielding $d' = 1$. Some observers reported a slight change in the position or depth of the masking bar when a very low contrast test bar was presented. To obscure this position cue, we randomly varied the lateral position of the dichoptic mask bar (when presented alone), or of the dichoptic mask and the test bar together (when both were presented), over a small range (± 1 min arc).

Observers

Four experienced psychophysical observers participated in these experiments. All had normal, or corrected to normal, visual acuity (20/20) for the viewing distance of 1.5 m, and normal stereopsis. All four observers were aware of the purpose of the experiments, and they had practiced in many of the experimental conditions before collecting the data presented here.

RESULTS

Dichoptic and monocular masking

Our first experiment measured monocular and dichoptic masking using the bar targets described above. For the monocular contrast thresholds, the masking and test bars were superimposed in one eye, and a uniform background field was presented to the other eye. The dichoptic thresholds were measured by presenting the test bar to the right eye at retinal correspondence with the masking bar presented to the left eye. The results from two subjects are shown in Fig. 4; the thresholds and the masking contrasts have been divided by the absolute contrast threshold of each subject. The absolute monocular thresholds for the four subjects ranged from 1.1 to 1.6%. These thresholds are considerably higher than the thresholds typically measured with either sinusoidal or bar targets. We attribute these high monocular thresholds to the low luminance of the background (6.6 cd/m^2) and the short vertical extent (46 min arc) of our targets. In a study using narrow bars presented at low luminance (4.6 cd/m^2), Legge (1978) measured absolute thresholds that were only slightly lower than ours, and his targets were much longer. Since we were chiefly interested in the contrast range that produced substantial threshold elevation, we measured dichoptic thresholds only at contrasts greater than five times

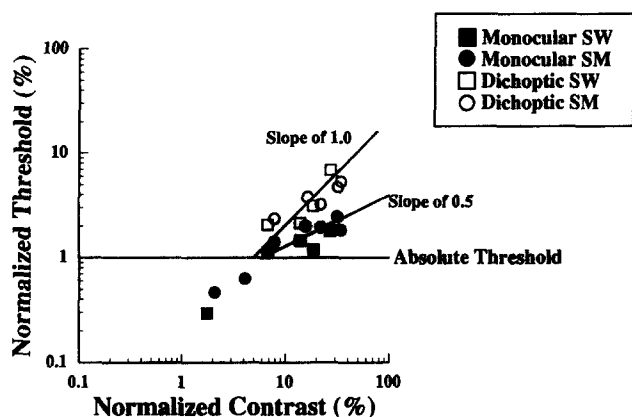


FIGURE 4. Contrast thresholds as a function of pedestal contrast for bar targets used in this study. Solid symbols show monocular thresholds for a contrast increment superimposed on pedestal contrast presented to the right eye. Open symbols show dichoptic thresholds for test bar presented to right eye with a masking bar presented at retinal correspondence in left eye. Thresholds and pedestal contrast have been divided by absolute monocular contrast threshold for each subject ("normalized"). Straight lines with slopes of 1.0 and 0.5 are shown as guides; they were not fitted to the data. Stimulus duration 200 msec.

threshold. A few monocular thresholds were measured at lower contrasts to establish that we could obtain the "dip" in the discrimination function that is characteristic of monocular contrast functions. The increment thresholds for both subjects "dipped" below absolute threshold at pedestal contrasts 2–3 times the absolute threshold.

Our findings show that the basic relationship between monocular and dichoptic masking, previously observed for gratings, also holds for bars. The dichoptic thresholds are higher than the monocular thresholds at the same masking contrast, and they appear to follow a steeper slope. The straight lines in Fig. 4 showing slopes of 0.5 and 1.0 are presented only as visual guides; they are *not* fitted to the data, but the data do fall close to these lines. Nevertheless, there is one striking difference between our results and previous findings. The normalized dichoptic thresholds are considerably lower than those observed with sinusoidal grating targets, by as much as half a log unit. The normalized monocular thresholds are also lower, but even when our monocular data are used to predict the dichoptic function, our dichoptic data are below predicted values [see Appendix, Fig. A1(C)]. For example, the dichoptic thresholds predicted from straightforward application of the quadratic summation rule to our monocular data should be roughly 12% at a contrast of 30%; instead the thresholds are <8%. It is tempting to ascribe this difference to lateral shifts of the mask produced by convergence movements. Legge (1979), however, also used nonius alignment for his measurements of dichoptic masking at 16 c/deg, where convergence jitter could easily have shifted the mask by as much as one-quarter of a period from trial-to-trial, since the precision of

nonius alignment is not much better than 1 min arc (McKee & Levi, 1987). He found that the normalized thresholds at 16 c/deg were in good agreement with the normalized thresholds at 0.25 c/deg where small shifts in the masking grating would have had a negligible effect. Moreover, Georgeson (1988) found the same amount of dichoptic masking for in-phase 1 c/deg gratings as for gratings with a 90 deg spatial phase shift between mask and test. We conclude that dichoptic masking for bar targets is qualitatively similar to dichoptic masking for gratings, but differs quantitatively; threshold elevation is less and undershoots the value expected from quadratic summation.

Dichoptic masking vs stereo matching

We next compared the dichoptic thresholds measured with the mask alone to the same thresholds measured when an extra bar was placed adjacent to the test bar—the comparison sketched in Fig. 1. For this experiment, the contrast of the adjacent extra bar was identical to the contrast of the mask and the center-to-center disparity between the mask and the extra bar was 10 min arc for subjects SM and DGT, and 15 min arc for subject MB. In Fig. 5, we have plotted test thresholds for both conditions as a function of the masking contrast. The addition of the extra bar produced a consistent improvement in the thresholds. At the highest masking contrast we could produce, the difference between the two conditions ranged from 0.3 to 0.9 log units.

Still, the test threshold for the extra bar condition hovered above absolute threshold for much of the tested range, so the release from masking was substantial, but not complete. The extra bar by itself might have elevated the threshold of the test bar. To test this, we measured the effect of the extra bar presented alone on the monocular test threshold. The extra bar by itself had very small effects on the test threshold that do not account for the incomplete release from dichoptic masking. For subject MB, the threshold with an extra bar of 43% contrast was 1.1%, essentially equal to her absolute monocular threshold (1.25%), but for subject SM, the threshold was 1.9%, only slightly below her threshold when both the dichoptic mask and the extra bar were present (2.1%).* However, an extra bar of 21% contrast presented alone had no effect the test threshold of subject SM, and her test threshold was still above her absolute threshold at this contrast (21%) when both mask and extra bar were presented together (see solid circles for top graph, Fig. 5). Thus, lateral masking from the extra bar alone does not explain the slight elevation above absolute threshold.

To return to the main result, what accounts for the substantial release from masking? Perhaps the improvement in threshold was produced by a change in convergence that shifted the mask away from correspondence with the test bar—a change in convergence induced by the extra bar. The stimulus duration was too brief to permit a voluntary change in convergence during a test trial, but the observers, anticipating the appearance of a target behind the fixation plane, might have initiated a

*The extra bar was located at a distance of 15 min arc for subject MB, but only 10 min arc for subject SM.

change in convergence when they pressed the button to begin a trial. Because the nonius lines were always visible, the observers would have noticed a large change in convergence, but they may not have detected small shifts. Observers reported that perceived shifts in the alignment of the nonius lines were generally much less than half the width of the nonius lines, or <3 min arc.

We therefore asked whether small lateral shifts of the masking bar away from retinal correspondence with the test bar could account for the observed release from masking. Dichoptic masking thresholds were measured according to our standard procedure, except that the

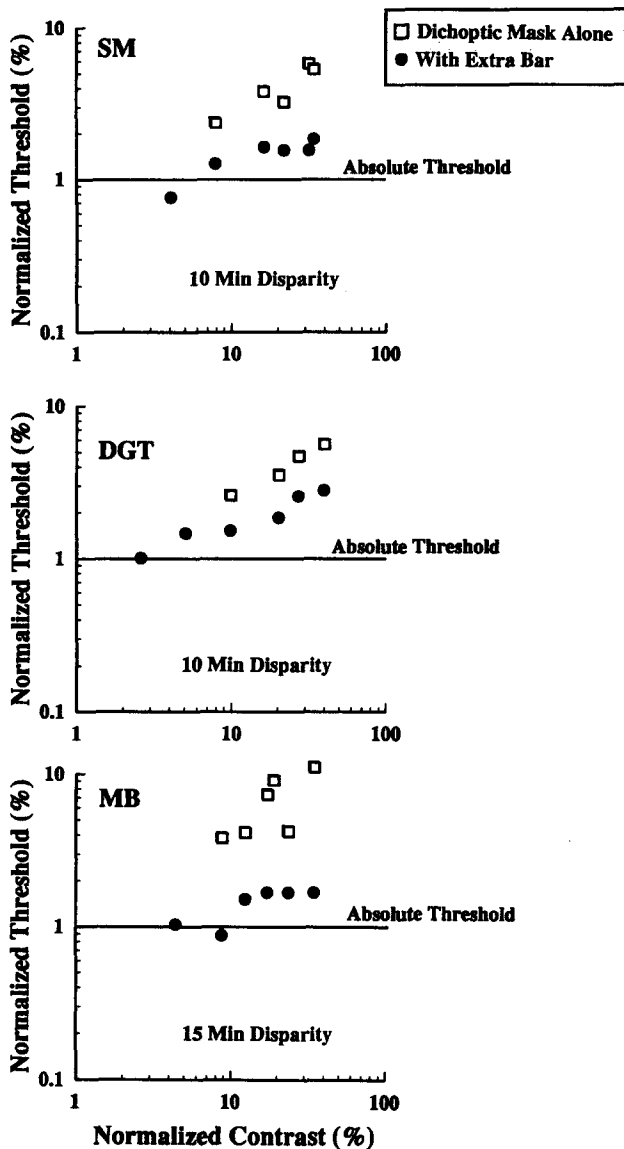


FIGURE 5. Dichoptic thresholds plotted as a function of masking contrast for the two conditions sketched in Fig. 1. Open symbols show dichoptic thresholds for test bar presented alone in right eye with masking bar in left eye. Solid symbols show dichoptic thresholds with an extra bar, matched in contrast to masking bar, presented adjacent to test bar in right eye. Center-to-center disparity between masking and extra bar 10 min arc for observers SM and DGT, and 15 min arc for observer MB. Stimulus duration 200 msec. Thresholds and masking contrast divided by absolute monocular contrast threshold for each observer.

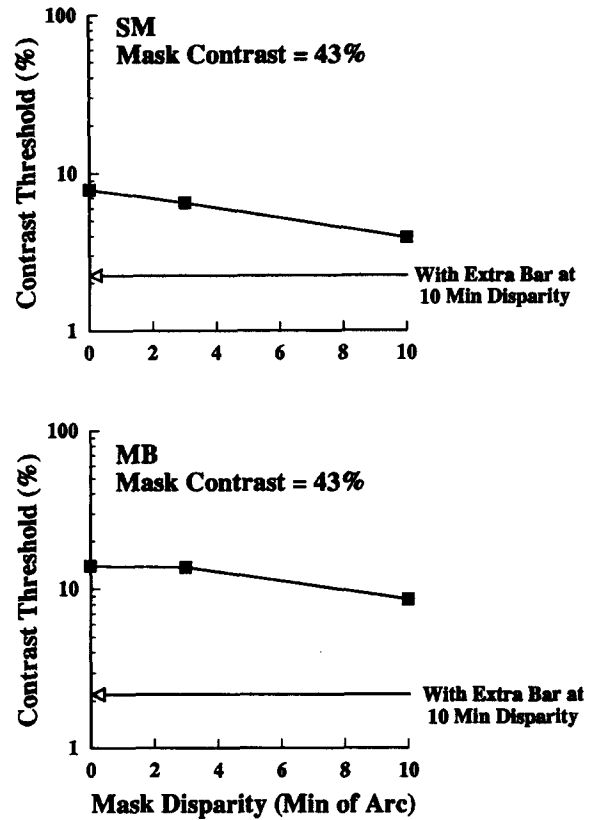


FIGURE 6. Dichoptic thresholds (not normalized) for test bar presented alone in right eye as a function of the disparity of the masking bar presented alone in left eye. Even with a 10 min arc offset, dichoptic thresholds for test bar alone are higher than thresholds measured with extra bar presented adjacent to test bar (shown by horizontal arrows). Stimulus duration 200 msec. Masking contrast 43% for all measurements. Disparity measured from center-to-center of bars. Thresholds remain elevated when mask is moved away from retinal correspondence with test bar.

mask was presented with a crossed disparity with respect to the test bar, i.e. to the right of retinal correspondence. There was no extra bar in the right eye's stimulus for these measurements. The results of changing the disparity of the mask bar are shown in Fig. 6 for a mask contrast of 43%. There is a decline in the effectiveness of the mask with increasing disparity, but even at 10 min arc, the thresholds are higher than the thresholds measured with the extra bar present (indicated by horizontal arrows). It is thus unlikely that shifts in convergence can account for our results.

Before we conclude that our results demonstrate the operations of stereo matching, we must first consider an explanation based on the output of a low spatial-frequency filter. Our bar targets have broad spatial-frequency spectra (Fig. 3) and could stimulate low as well as high-frequency mechanisms. Mechanisms tuned to low frequency (e.g. 1 c/deg) have relatively coarse spatial resolution and would not resolve the test bar and the extra bar when they are separated by 10 min arc. Imagine that such a mechanism is stimulated by the masking bar in one eye and the extra bar in the other eye, both having the same contrast. This is equivalent to presenting this low-frequency mechanism with a normal

binocular stimulus. The addition of the test stimulus to one eye simply increases the monocular signal strength in that eye. This means that the discrimination task at the output of the low-frequency filter is no longer a dichoptic discrimination; instead, it is a variant of binocular masking (equivalent to a masking contrast C in both eyes, and an incremental test contrast, ΔC , in one eye). We expect that the dependence of monocular test contrast ΔC on binocular masking contrast C would resemble findings reviewed in the Introduction. If so, threshold elevation would be less with the extra bar than for dichoptic masking. The reduction in threshold elevation associated with the addition of the extra bar could be due to a conversion from dichoptic masking to a variant of binocular masking at the output of a low-frequency filter. Legge's quadratic summation rule leads to the same prediction. In the Appendix, we show the calculations that lead to this prediction, and functions for binocular (or monocular) masking and for dichoptic masking [see Fig. A1(C)]. The predicted function for a monocular increment presented with a binocular mask falls between these two functions, rising with a slope that is parallel to the binocular masking function. At high contrasts, this function is about 0.5 log units below the dichoptic function—an improvement in threshold comparable to the results shown in Fig. 5.

The low-frequency filter explanation also predicts that the threshold elevation for the test bar should be about the same whether the extra bar is located 10 min arc to its right or superimposed on top of it. In Table 1, we show the actual (*not* normalized) test thresholds measured when the extra bar was presented at either 0 or 10 min arc disparity. The thresholds for DGT are the same in these two conditions, but, the prediction fails for the other two observers. Their thresholds are actually better at a disparity of 10 min arc than at 0 disparity, suggesting that matching operations account for the release from masking observed in their data. We next measured the disparity range of these putative matching operations for these two observers (see Fig. 7). As predicted by the quadratic summation rule, their thresholds at 0 disparity were below their dichoptic masking thresholds, but in agreement with the data in

Table 1, their thresholds were lower still at disparities between 10–20 min arc. At yet larger disparities, these two observers were unable to fuse the mask with the extra bar, and their test thresholds rose to the dichoptic masking level. Apparently, the disparity range for matching is limited to about 20–30 min arc for bar targets presented to the fovea.

Our findings are consistent with the existence of two separate mechanisms that account for the release from masking when the extra bar is added. One is mediated by a low-frequency filter and appears to use quadratic summation as a binocular combination rule. It accounts for the data of observer DGT and for the case in which the extra bar is superimposed on the test bar. The second mechanism, perhaps functioning at moderate to high spatial frequencies, seems to depend on matching operations within the human stereo system. Most likely, both mechanisms operate in all three observers. If the observers are using an optimum strategy, they will choose the mechanism that is most sensitive in the tested condition. For subject DGT, the low frequency mechanism may always be the most sensitive for detecting the contrast of bar targets. The other two subjects may alternate between the two mechanisms depending on the contrast level of the mask and the disparity of the extra bar. This alternation between mechanisms would explain why their thresholds in the extra bar condition remain higher than the absolute threshold at mask contrasts above 10% (see solid circles in Fig. 5). For example, the absolute monocular threshold for the test bar could depend on the low frequency mechanism, but the release from masking could be mediated by matching operations in a higher frequency mechanism. Even if the release from masking were complete, the absolute contrast threshold of this higher frequency mechanism would probably be higher than the absolute threshold of the low frequency mechanism. Thus, the threshold in the extra bar condition would always be slightly above the measured absolute threshold of the test bar; it would correspond to the absolute contrast threshold of the less sensitive, higher frequency mechanism.

Stereo matching at high spatial frequencies

To establish that matching could indeed contribute to the observed release from masking for all three subjects, we repeated the measurements with high spatial frequency triphasic targets—targets that would be essentially invisible to a low-frequency filter (see Methods). The stimulus arrangements were identical to those used for the bar experiments, except that the center-to-center disparity between the mask and the extra triphasic target was reduced to 7 min arc. At disparities greater than about 8 min arc, the mask and extra triphasic targets did not fuse, producing instead a confusing diplopic image with no compelling sense of depth. This disparity limitation is predictable from the abundant psychophysical evidence that the disparity range scales with spatial frequency (Felton, Richards & Smith, 1972; Tyler, 1975; Schor, Wood & Ogawa, 1984; Smallman & MacLeod, 1994). In Table 2, we show three

TABLE 1. Threshold for test bar when contrast of mask and extra bar = 43%

	Extra bar at 0 disparity	Extra bar at 10 min arc disparity
SM	3.6 ± 0.3	2.1 ± 0.3
DGT	3.7 ± 0.8	3.3 ± 0.3
MB	6.6 ± 0.1	2.2 ± 0.5

Contrast thresholds for test bar presented to right eye with masking bar presented at retinal correspondence in left eye. Left column: thresholds when test bar presented superimposed on extra bar in right eye at retinal correspondence with masking bar in left eye (0 disparity). Right column: thresholds when extra bar presented to right eye adjacent to test bar at a disparity of 10 min arc (center-to-center) with respect to masking bar. Contrast to mask and extra bars = 43%.

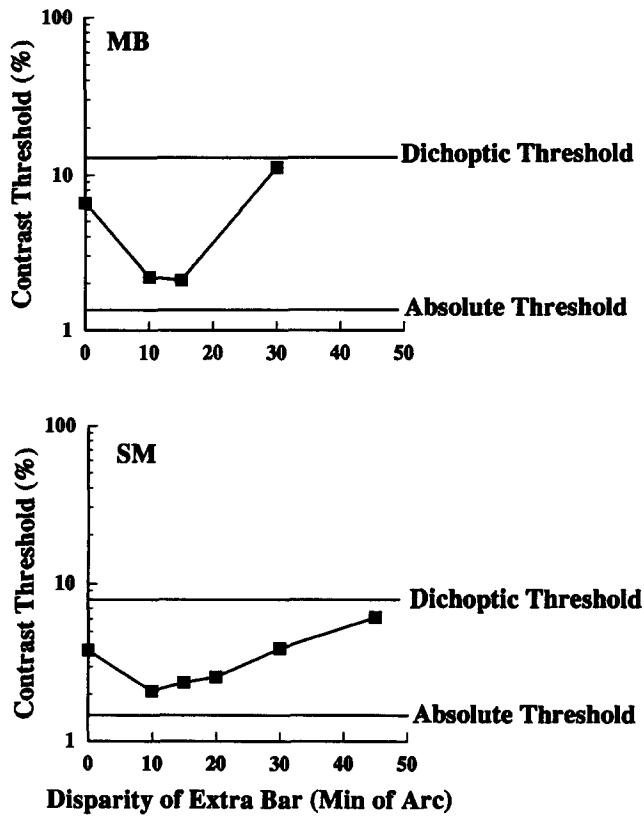


FIGURE 7. Dichoptic thresholds (not normalized) for test bar as a function of the center-to-center disparity between mask and extra bars. Test bar presented to right eye at retinal correspondence with masking bar presented to left eye. Extra bar presented to right eye on the right side of the test bar. Stimulus duration 200 msec. Contrast of mask and extra bars is 43%.

different thresholds measured with these triphasic targets: the absolute monocular contrast threshold, the dichoptic masking threshold with a 31% contrast mask, and the threshold with the extra triphasic target presented adjacent to the test target. The results for all three subjects were similar, although masking was substantially less for subject DGT than for the other two subjects. As we found with the bars, the presence of the extra target released the test target from masking.* Thus, the measurements with the high spatial frequency targets confirm our previous conclusion—matching does contribute to the release from masking.

Matching with unequal contrasts

Many computational models of stereopsis impose a “uniqueness” constraint on stereo matching—each element in one half-image can be matched with only one

*The test thresholds in the presence of the extra bar were not significantly different from the absolute thresholds for these triphasic targets, showing a complete release from masking. This outcome could be fortuitous. Note that at masking contrasts of 3–5 times threshold, the release from masking with bar targets was also complete. Unfortunately, the contrast range for these triphasic targets was limited by the low luminance available for the present study, so we could not explore the release from masking at higher contrasts.

TABLE 2. High spatial frequency targets

	Absolute test threshold	Dichoptic masking threshold 31% contrast	Threshold with extra bar at 7 min arc disparity
SM	6.2 ± 0.6	19.6 ± 3.2	6.3 ± 0.8
DGT	4.8 ± 0.1	9.2 ± 1.3	5.6 ± 0.6
MB	5.8 ± 0.4	24.0 ± 0.8	7.0 ± 1.8

Contrast thresholds for high spatial frequency triphasic targets (see Fig. 3 for amplitude spectrum) presented to right eye. Left column: absolute monocular contrast thresholds. Middle column: dichoptic thresholds in presence of masking target (31% contrast) presented to left eye at retinal correspondence with test target. Right column: dichoptic thresholds when an extra triphasic target (31% contrast) is presented to right eye adjacent to test target at a disparity of 7 min arc (center-to-center) with respect to masking target.

element in the other half-image (Marr & Poggio, 1976, 1979; Marr, 1982; Frisby & Pollard, 1991). Once the best match has been chosen, alternative matches are eliminated via some sort of “inhibition”. For our austere stimulus configuration, the competition among potential matches is fairly simple; the high contrast masking bar can match with the faint contrast test bar at retinal correspondence, or the high contrast extra bar at a disparate position, or both (see depth map in Fig. 8). Making two matches would violate the uniqueness constraint, so, if “uniqueness” constrains human stereo matching, the operations that eliminate one of the two matches may account for our results. Let us assume that the best match is between the two high contrast bars—the mask bar and the extra bar (bright circle in Fig. 8 where the lines of sight cross). In the Marr and Poggio models (1976, 1979), inhibition from the best match is implemented along the monocular lines of sight. This line-of-sight inhibition (diagrammed by arrow in Fig. 8) could turn off the monocular masking signal from the




FIGURE 8. Depth map showing projected matches between targets presented to the two eyes. High contrast masking bar in left eye can be matched with the high contrast extra bar in right eye, or with faint test bar at retinal correspondence, or with both bars. Line-of-sight inhibition (shown by diagonal arrow) can prevent the mask from combining with the test bar, thereby releasing it from masking.

left eye at the site of the test bar, thereby permitting the monocular test bar signal from the right eye to “sneak through” unmasked.

What makes the match between the two high contrast bars *better* than the match between the mask and the test bars? In our simple configuration, only disparity and relative contrast can determine the best match. We have already shown that the release from masking depends on the disparity between the mask and extra bar. If no suitable alternative match falls within a limited disparity range, the masking bar will combine with the test bar, producing an elevation in threshold. For targets lying within this range, our results so far indicate that contrast equivalence is a more powerful determinant of matching than target disparity; the match between the masking bar and the extra bar is better than the alternative match at zero disparity presumably because the contrast of the test bar is so different from the contrast of the mask. To examine the influence of contrast on stereo matching in greater detail, we systematically varied the contrast of the extra bar, and measured the test bar threshold for a fixed masking bar contrast. In Fig. 9, test thresholds are plotted as a function of the contrast of the extra bar. Fig. 9(A,B) shows the data from observer SM for mask contrasts of 46 and 22% respectively; Fig. 9(C) shows data from observer MB at a mask contrast of 24%.

Contrast similarity is not a rigorous constraint on stereo matching. The graphs show that there is a large range of extra bar contrasts that produce a release from masking. The dark vertical arrow in each graph indicates the contrast of the extra bar that equals the contrast of the mask, but the test threshold does not reach a minimum at this point. In fact, test thresholds reach a shallow minimum when the contrast of the extra bar is roughly a half log unit below the contrast of the mask. This result probably rules out any major contribution from a low spatial frequency mechanism that sums together the test and extra bar contrasts, at least in these two subjects. The thresholds predicted by the quadratic summation model reach a shallow minimum when the contrasts of the mask and the extra bar are equal; the addition of a weak pedestal contrast, i.e. the dim extra bar, to the test bar should have almost no effect on the dichoptic threshold, contrary to our findings [see Appendix, Fig. A1(D)].

As the contrast of the extra bar was reduced below one-fifth of the masking bar contrast, the test thresholds rose precipitously. Note, that at this contrast level, the extra bar was close to the dichoptic masking threshold of the test bar measured with the mask alone. We observed some peculiar changes in the appearance of the targets in this low contrast condition. In our previous experiments with equal contrasts and small disparities (< 20 min arc), a single bright bar, corresponding to the fused percept of the mask and the extra bar, was seen on every trial at a depth behind the fixation plane. The test bar, when visible, appeared as a faint shimmering line slightly to the left and in front of this bright bar. The observers usually made their judgements on the basis of the number of bars (one or two) they saw in the stimulus.

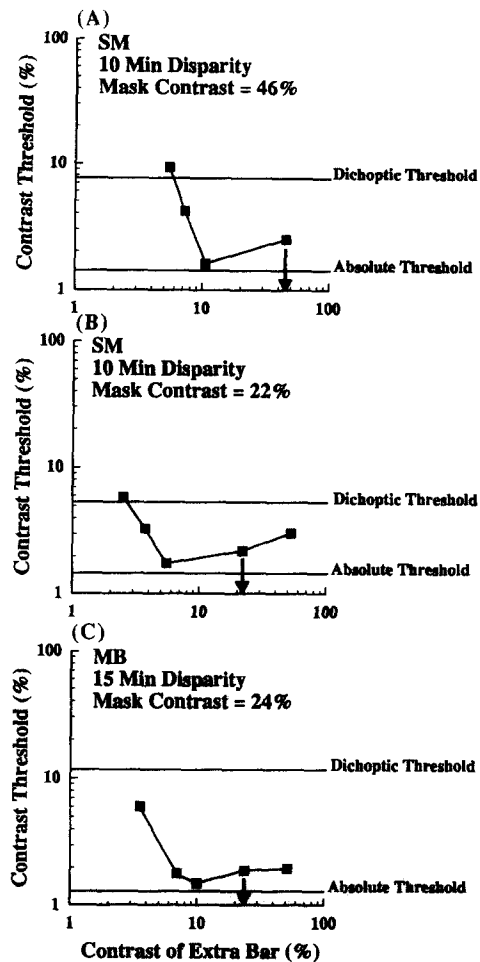


FIGURE 9. Dichoptic thresholds (not normalized) for test bar as a function of the contrast of the extra bar presented in the right eye adjacent to test bar. (A) Contrast of masking bar, presented in the left eye at retinal correspondence with test bar, = 46%, disparity between mask and extra bar = 10 min arc (center-to-center), observer SM. (B) Mask contrast = 22%; disparity = 10 min arc, observer SM. (C) Mask contrast = 24%; disparity = 15 min arc, observer MB. The upper horizontal lines show the dichoptic thresholds measured with the mask alone; the lower horizontal lines show the absolute monocular thresholds for the test bar. The vertical dark arrows show the extra bar contrast which is equal to the mask contrast.

In the present experiment, the same percept was seen as long as the contrast of the extra bar was greater than one-fifth of the mask. When the contrast of the extra bar fell much below this value, the appearance changed radically. Sometimes the bright masking bar appeared behind the plane, but sometimes it appeared in the fixation plane at the locus of the test bar. Initially, both observers thought that the shift implied that the test bar had been presented on that trial—that the mask was now making a match with the test bar in preference to the dim extra bar. No such luck! The visible shift of the mask was *not* correlated with the presence of the test bar so the thresholds rose. In the absence of a solid match at the disparate position of the extra bar, the mask moved to the position corresponding to its monocular location in the fixation plane, whether or not there was a corresponding stimulus in the other eye at this location.

We conclude that stereo matching has a contrast threshold. A high contrast feature in one eye can be combined with a low contrast feature in the other eye, but the lower contrast must be greater than about one-fifth of the higher contrast. The fact that the mask combined with the test bar in our dichoptic masking condition (no extra bar) would seem to contradict this conclusion, because the test bar was generally below the requisite contrast threshold. However, our observations on the shifting location of the mask, described above, suggest that the mask and the test bar are not "actively" matched. Consider what happens if we present the masking bar to one eye and a uniform field to the other. The masking bar is easily visible, and it occupies the same visual location it would occupy if a bar of identical contrast were presented at retinal correspondence in the other eye. If all retinal signals, even monocular signals, stimulate binocularly-driven cortical units, then the monocular bar stimulates the same population of units that respond to a binocular bar presented in the fixation plane. Similar reasoning applies to the test bar. The contrast threshold for the test bar may be elevated in dichoptic masking because the mask and test bars happen to share the same set of binocular units. In this view, retinal correspondence is special because it is the default position of unmatched features. Note that if stereo matching has a threshold, then the putative inhibition invoked in Fig. 8 to enforce "uniqueness" may be unnecessary. There is no need to inhibit a match that is never made.

As Fig. 9 shows, there is also a gentle rise in threshold as the contrast of the extra bar increases. Given the broad spectrum of the bar targets, the shape of this function could reflect shifts from one spatial mechanism to another with changes in relative sensitivity. But there is another explanation. Recently, a number of papers have demonstrated a non-specific mutual inhibition among neighboring striate units in the cat, a process called contrast normalization (Robson, 1988; Bonds, 1989; Heeger, 1992). Some psychophysical models have also incorporated contrast normalization to explain local contrast masking phenomena (Grossberg & Mingolla, 1985; Grossberg, 1987; Lubin & Nachmias, 1990; Graham, 1991; Bergen & Landy, 1991). In our stimulus, the test threshold may rise because of non-specific inhibition from the high contrast match between the mask and the extra bar. This non-specific inhibition depends on contrast and proximity. We noted that the extra bar, by itself, had a slight effect on the test thresholds, but presumably, the test threshold would also rise if the contrast of the extra bar alone were high enough. However, the binocular-summed contrast of the match is always greater than the contrast of the extra bar alone, and the binocular location of the match is also closer than the extra bar alone, if we assume that the match is located somewhere between the monocular half-images (Sheedy & Fry, 1979; Mansfield, Akutsu & Legge, 1992), so its influence is greater than the extra bar alone.

DISCUSSION

Dichoptic masking is contingent on stereo matching. Whether a high contrast target in one eye will mask a low contrast test target in the other eye depends on the other potential matches in the vicinity of the test target. We found that a high contrast bar in one eye strongly masked the contrast threshold of a test bar presented in the other eye at retinal correspondence. When we presented another high contrast monocular bar adjacent to the test bar, the masking effect nearly disappeared. The superior match between the two high contrast features apparently intervened to prevent masking at retinal correspondence.

Many years ago, Hirsch (1948) observed that a monaural signal presented with monaural noise to one ear became more detectable if the same noise was added in phase to the other ear, an effect known as masking level difference (MLD). Binaural signals were also easier to detect if the binaural phase difference of the signal was different from the binaural phase difference of the noise—binaural masking level difference (BMLD). A number of studies have attempted to identify a visual analog of this effect (Henning & Hertz, 1973, 1977; Moraglia & Schneider, 1990, 1992; Schneider & Moraglia, 1992). Generally, these studies have shown that if the disparity of the visual noise is different from the disparity of the signal, the signal is more detectable than if signal and noise have the same disparity. Our results are undoubtedly related to these "binocular unmasking" effects, but there is an important difference. In our paradigm, it is difficult to specify the disparity of the signal, i.e. the test bar. For example, if the test bar is paired stereoscopically with the masking bar at retinal correspondence (zero disparity), threshold elevation should be the same with or without the extra bar. Schneider and Moraglia (1992) proposed separate monocular and binocular channels, arguing that detection is based on the channel with the higher signal-to-noise ratio. However, if there are separate monocular channels, there should never be any dichoptic masking because a monocular channel is not affected by a mask in the other eye. We argue instead that, when the stereo system matches the mask bar with the extra bar, the test bar becomes an unmatched feature that, by default, is detected by a *binocular* mechanism tuned to zero disparity. We think our results are better considered in the context of human stereo matching.

In most studies in human stereo matching, the targets are random-dot stereograms or other patterns that contain a large number of similar elements. Thus, in combining the monocular half-images, the stereo system is challenged with many equally valid matches. The early stereo matching models (Julesz, 1971; Nelson, 1975; Dev, 1975; Marr & Poggio, 1976) concentrated on resolving the ambiguities posed by these complex targets, and paid little attention to the specific characteristics of the matching elements—they were just the pixels that made up the stereo half-images. These early models imposed the simplest "compatibility" constraint on the matching elements, requiring only that black dots were

matched to black dots and white dots to white dots. In subsequent models (Marr & Poggio, 1979; Grimson, 1981; Frisby & Pollard, 1991), the matching elements became more refined tokens, e.g. zero-crossings, that were derived from pre-processing of the retinal image by scaled spatial band-pass filters or other operators. The compatibility constraint was made more specific—that the zero-crossings have the same sign and/or orientation, and perhaps arise from the same spatial frequency bandwidth. Nevertheless, like their predecessors, these models also operated on monocular place-markers; graded information about contrast was discarded prior to the postulated matching operations.

Our results show that contrast does affect stereo matching, as others have also demonstrated (Jordan, Geisler & Bovik, 1990; Foley, 1976; Krol & van de Grind, 1979). We also found that the contrasts of the matching features need not be identical, in support of Marr's contention that "gray level" information is not an important constraint on stereo matching. However, this conclusion misses the point. *If binocular contrast summation occurs after the matching operations, then graded information about contrast must be preserved beyond the neural level where matching occurs.* The human stereo system will readily match features of different contrasts, but this result does not imply that graded information about contrast has been discarded prior to the matching operations.

Recently, Jones and Malik (1991) proposed a computational model of stereo matching that utilizes linear spatial filters—the same type of filters that have been used to model human spatial vision. These filters are tuned for spatial frequency, position, and orientation, and are driven by stimulus contrast. Stereo matching is based on identifying the image patch in one stereo half-image that is most similar to the image patch in the other stereo half-image after convolution with a family of scaled spatial filters. The minimum difference is determined by comparing the outputs of spatial filters, convolved with a patch of one half-image, to the filter outputs associated with each point within a defined search area in the other half-image. After the initial matches are chosen on the basis of minimum differences, additional constraints based on surface continuity and occlusion are then used to improve depth reconstruction. Jones and Malik showed that this filter-based matching algorithm can do a reasonable job of reconstructing a depth map from random dot stereograms and natural images. Clearly, their approach, if modified for human binocular vision, could explain our results. The filtered signals generated by the extra bar would be more similar to the filtered signals generated by the mask bar than to those generated by the test bar, since contrast affects the magnitude of the filtered signals. There is still a question about how the brain could implement this choice physiologically. Are less optimal matches inhibited by the strongest pairing?

We suggested above that stereo matching has a contrast threshold. Similar features that differ in contrast by more than a factor of five are not matched. Interestingly,

Legge and Gu (1989) found that they could not measure a stereo threshold when the contrast ratio of the stereo half-images fell below about one-fifth. Mansfield, Akutsu and Legge (1992) measured the perceived visual direction of band-limited disparate targets as a function of the contrast ratio of the stereo half-images. They observed that the visual direction shifted systematically toward the monocular location of the higher contrast half-image, suggesting again that a stereo match requires some minimum contrast ratio between the two half-images. Since, in physiological terms, differences in spatial frequency, orientation or position can be mimicked by differences in contrast (e.g. Kontsevich & Tyler, 1994), this matching threshold could be generalized to the multiple dimensions of the Jones–Malik filters. A threshold of this kind could eliminate most alternative matches. However, if stereo matching is limited only by a threshold and a range of permissible disparities (the "search area"), then it is certainly possible that double matches could occur under some circumstances. One such circumstance is the well-known Panum's limiting case, where a single feature in one eye is matched with two identical features in the other eye, in obvious violation of the "uniqueness constraint".

Contrast measurements are a new tool for exploring the rules that govern human stereo matching. They also offer the prospect that two different approaches to binocular combination—stereo matching models and binocular contrast summation models—can be merged into a more unified view of human binocular processing. In future work, we intend to use contrast thresholds to determine if the constraints commonly assumed in computational stereo matching apply to human stereo matching, and how these constraints influence binocular contrast summation.

REFERENCES

- Bergen, J. R. & Landy, M. S. (1991). Computational modeling of visual texture segregation. In Landy, M. S. & Movshon, J. A. (Eds), *Computational models of visual processing*. Cambridge, Mass.: MIT Press.
- Bravo, M. J. & Blake, R. (1992). The contributions of figure and ground textures to segmentation. *Vision Research*, *32*, 1793–1800.
- Bonds, A. B. (1989). Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience*, *2*, 41–55.
- Burton, G. J. (1973). Evidence for non-linear response processes in the visual system from measurements on the thresholds of spatial beat frequencies. *Vision Research*, *13*, 1211–1225.
- Dev, P. (1975). Perception of depth surfaces in random-dot stereograms: A neural model. *International Journal of Man–Machine Studies*, *7*, 511–528.
- De Valois, R. L. & De Valois, K. K. (1988). *Spatial vision*. New York: Oxford University Press.
- Elliott, P. B. (1964). Tables of d' . In Swets, J. A. (Ed.), *Signal detection and recognition by human observers* (pp. 651–684). New York: Wiley.
- Felton, T. B., Richards, W. & Smith, R. A. (1972). Disparity processing of spatial frequencies in man. *Journal of Physiology, London*, *234*, 349–362.
- Foley, J. M. (1976). Binocular depth mixture. *Vision Research*, *16*, 1263–1267.

- Foley, J. M. & Legge, G. E. (1981). Contrast detection and near-threshold discrimination in human vision. *Vision Research*, *21*, 1041–1053.
- Freeman, R. D. & Ohzawa, I. (1990). On the neurophysiological organization of binocular vision. *Vision Research*, *30*, 1661–1676.
- Frisby, J. P. & Pollard, S. B. (1991). Computational issues in solving the stereo correspondence problem. In Landy, M. S. & Movshon, J. A. (Eds), *Computational models of visual processing*. Cambridge, Mass.: MIT Press.
- Georgeson, M. A. (1988). Spatial phase dependence and the role of motion detection in monocular and dichoptic forward masking. *Vision Research*, *28*, 1193–1206.
- Graham, N. (1991). Complex channels, early local nonlinearities, and normalization in texture segregation. In Landy, M. S. & Movshon, J. A. (Eds), *Computational models of visual processing*. Cambridge, Mass.: MIT Press.
- Grimson, W. E. L. (1981). *From images to surfaces: A Computational study of the human early visual system*. Cambridge, Mass.: MIT Press.
- Grossberg, S. (1987). Cortical dynamics of three-dimensional form, color, and brightness perception: Monocular theory. *Perception & Psychophysics*, *41*, 87–116.
- Grossberg, S. & Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent features. *Perception & Psychophysics*, *38*, 141–171.
- Halpern, D. L. & Blake, R. (1988). How contrast affects stereoacuity. *Perception*, *17*, 483–495.
- Heckmann, T. & Schor, C. M. (1989). Is edge information for stereoacuity spatially channeled? *Vision Research*, *29*, 593–608.
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, *9*, 181–197.
- Henning, G. B. & Hertz, B. G. (1973). Binocular masking level differences in sinusoidal grating detection. *Vision Research*, *13*, 2455–2463.
- Henning, G. B. & Hertz, B. G. (1977). The influence of bandwidth and temporal properties of spatial noise on binocular masking-level differences. *Vision Research*, *17*, 399–402.
- Hirsch, I. J. (1948). The influence of interaural phase on interaural summation and inhibition. *Journal of the Acoustical Society of America*, *20*, 536–544.
- Jones, D. G. & Malik, J. (1991). A computational framework for determining stereo correspondence from a set of linear spatial filters. Report No. UCB/CSD 91/655. Computer Science Division (EECS), University of California, Berkeley, Calif.
- Jordan, J. R., Geisler, W. S. & Bovik, A. C. (1990). Color as a source of information in the stereo correspondence process. *Vision Research*, *30*, 1955–1970.
- Julesz, B. (1971). *Foundations of Cyclopean perception*. Chicago, Ill.: University of Chicago Press.
- Kontsevich, L. L. & Tyler, C. W. (1994). Analysis of stereothresholds for stimuli below 2.5 cyl/deg. *Vision Research*. In press.
- Krol, J. D. & van de Grind, W. (1980). The double-nail illusion: Experiments on binocular vision with nails, needles and pins. *Perception*, *9*, 651–669.
- Legge, G. E. (1978). Space domain properties of a spatial frequency channel in human vision. *Vision Research*, *18*, 959–970.
- Legge, G. E. (1979). Spatial frequency masking in human vision: Binocular interactions. *Journal of the Optical Society of America*, *69*, 838–847.
- Legge, G. E. (1981). A power law for contrast discrimination. *Vision Research*, *21*, 457–467.
- Legge, G. E. (1984a). Binocular contrast summation—I. Detection and discrimination. *Vision Research*, *24*, 373–383.
- Legge, G. E. (1984b). Binocular contrast summation—II. Quadratic summation. *Vision Research*, *24*, 385–394.
- Legge, G. E. & Gu, Y. (1989). Stereopsis and contrast. *Vision Research*, *29*, 989–1004.
- Legge, G. E. & Kersten, D. (1983). Light and dark bars: Contrast discrimination. *Vision Research*, *23*, 473–483.
- Levi, D. M., Harwerth, R. S. & Smith, E. L. (1979). Humans deprived of normal binocular vision have binocular interactions tuned to size and orientation. *Science*, *206*, 852–854.
- Levi, D. M., Harwerth, R. S. & Smith, E. L. (1980). Binocular interactions in normal and anomalous binocular vision. *Documenta Ophthalmologica*, *49*, 303–324.
- Lubin, J. & Nachmias, J. (1990). Discrimination contours in an F/3F stimulus space. *Investigative Ophthalmology and Visual Science (Suppl.)*, *31*, 409.
- Mansfield, J. S. & Simmons, D. R. (1989). Contrast threshold for the identification of depth in bandpass-filtered stereograms. *Investigative Ophthalmology and Visual Science (Suppl.)*, *30*, 251.
- Mansfield, J. S., Akutsu, H. & Legge, G. E. (1992). Interocular contrast differences introduce lateral shifts in the perceived location of binocular depth targets. *Investigative Ophthalmology and Visual Science (Suppl.)*, *33*, 1332.
- Marr, D. (1982). *Vision*. San Francisco, Calif.: (W.H.) Freeman.
- Marr, D. & Poggio, T. (1976). Cooperative computation of stereo disparity. *Science*, *194*, 283–287.
- Marr, D. & Poggio, T. (1979). A theory of human stereopsis. *Proceedings of the Royal Society of London B*, *204*, 301–328.
- McKee, S. P. & Levi, D. M. (1987). Dichoptic hyperacuity: The precision of nonius alignment. *Journal of the Optical Society of America A*, *4*, 1104–1108.
- Moraglia, G. & Schneider, B. (1990). Effects of direction and magnitude of horizontal disparities on binocular unmasking. *Perception*, *19*, 581–593.
- Moraglia, G. & Schneider, B. (1992). On binocular unmasking of signals in noise: Further tests of the summation hypothesis. *Vision Research*, *32*, 375–385.
- Nelson, J. I. (1975). Globality and stereoscopic fusion in binocular vision. *Journal of Theoretical Biology*, *49*, 1–88.
- 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.
- Poggio, G. F. & Fischer, B. (1977). Binocular interaction and depth sensitivity in striate and prestriate cortex of behaving rhesus monkey. *Journal of Neurophysiology*, *40*, 1392–1405.
- Robson, J. G. (1988). Linear and nonlinear operations in the visual system. *Investigative Ophthalmology and Visual Science (Suppl.)*, *29*, 117.
- Schneider, B. & Moraglia, G. (1992). Binocular unmasking with unequal interocular contrast: The case for multiple cyclopean eyes. *Perception & Psychophysics*, *52*, 639–660.
- Schor, C. M., Wood, I. & Ogawa, J. (1984b). Binocular sensory fusion is limited by spatial resolution. *Vision Research*, *24*, 661–665.
- Sheedy, J. E. & Fry, G. (1979). The perceived direction of the binocular image. *Vision Research*, *19*, 201–211.
- Smallman, H. S. & MacLeod, D. I. A. (1994). A size disparity correlation at masking threshold. *Journal of the Optical Society of America*. In press.
- Tyler, C. W. (1975). Spatial organization of binocular disparity sensitivity. *Vision Research*, *15*, 583–590.

Acknowledgements—This research was supported by NEI grant RO1-EY06644 and NEI core grant 5P-30-EY01186. Gordon Legge's participation in this research was supported by NEI grant RO1-EY02857. We wish to thank Dr Scott Watamaniuk for serving as an observer in these experiments.

APPENDIX

In 1984, Legge proposed that the binocular signal strength b is related to left- and right-eye signal strength l and r by the quadratic-summation formula:

$$b = \sqrt{(l^2 + r^2)}.$$

As illustrated by Fig. A1(A), a binocular signal of fixed size b can be represented as the points on a circle. Different points on the circle represent different pairs of left-eye and right-eye signals whose quadratic sum is b . Given binocular signal b , there is a binocular signal

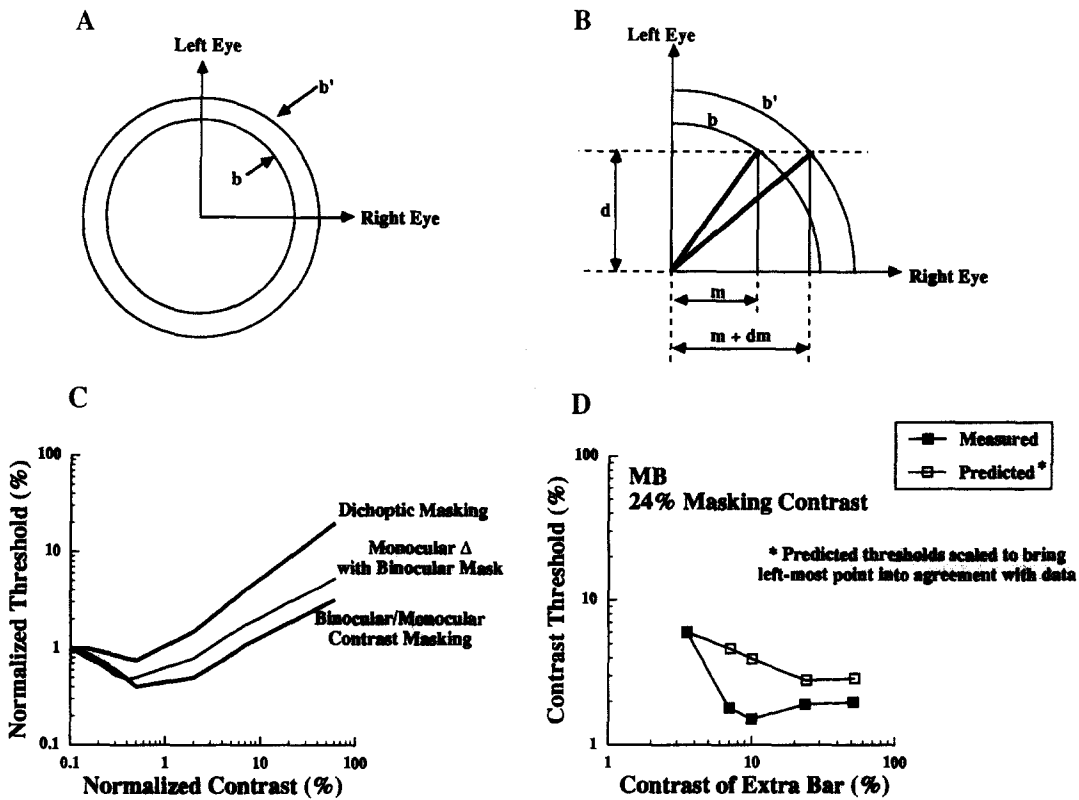


FIGURE A1. Geometric model of binocular signals b and b' . (B) Monocular discrimination threshold dm in the presence of a pedestal m in the same eye and a pedestal d in the other eye. See text. (C) Lower curve: contrast discrimination function based on rough fit to monocular data in Fig. 4. Middle curve: monocular discrimination threshold (dm) predicted from quadratic summation model for equal contrast pedestals ($m = d$) in two eyes. Upper curve: dichoptic masking function predicted from quadratic summation model from lower monocular function. (D) Lower curve from bottom graph in Fig. 9; upper curve shows predictions of quadratic summation model for monocular threshold dm in the presence of unequal pedestals ($m \neq d$) in two eyes. Predicted function scaled so that left-most point agrees with measured threshold.

b' which is a jnd larger. It is represented by a larger circle. Any combination of left-eye and right-eye signals that produces b' should be discriminably different from the pair producing signal b .

The quadratic summation rule does not specify how b' relates to b , i.e. it does not specify the shape of the discrimination function. Let db be the jnd in the binocular signal:

$$db = b' - b.$$

Designate the function relating db to b by:

$$db = f(b).$$

For purposes of this analysis, we assume that the discrimination function $f(b)$ is a dipper type function that approximates the monocular data shown in Fig. 4.

Suppose we present a pedestal of strength m to one eye and a pedestal of strength d to the other. We increase m to find the discrimination threshold dm . When $d = 0$, this is the same as purely monocular discrimination [the lower curve shown in Fig. A1(C)]. When $m = 0$, this is the same as the dichoptic masking experiment. The dichoptic masking function predicted from the monocular data is shown by the upper curve in Fig. A1(C). When $m = d$, dm is a monocular increment of a binocular stimulus.

The binocular signal strength b that occurs with d in one eye and m in the other is [see Fig. A1(B)]:

$$b = \sqrt{(m^2 + d^2)}.$$

For discrimination to occur, m must be incremented by dm so that the resulting binocular signal b' is just discriminable from b . As shown in

Fig. A1(B), dm is just the difference in length between the bases of the two triangles. The base of the large triangle has length $\sqrt{(b'^2 - d^2)}$ and the base of the small triangle is just m . Therefore,

$$dm = \sqrt{(b'^2 - d^2)} - m. \tag{3}$$

Recall that $b' = b + db$ and $db = f(b)$:

$$dm = \sqrt{[(b + db)^2 - d^2]} - m. \tag{4}$$

Equation (4) permits a prediction of the molecular increment threshold for any (m, d) pair of pedestals presented to the two eyes.

In the context of our experiments, dm is the test bar, m is the extra bar, and d is the masking bar. Let us assume that the test and the extra bars are summed together. The central curve in Fig. A1(C), labeled "monocular Δ with binocular mask", shows the predicted contrast discrimination function when the mask and extra bars are the same contrast ($m = d$). In Fig. A1(D), we show the predicted function for a condition when the extra bar is not equal to the masking bar ($m \neq d$). The masking contrast (d) is fixed at 24% and the contrast (m) of the extra bar is varied systematically. The predicted function has been multiplied to bring the left-most point into agreement with the measured threshold; recall that our measured dichoptic thresholds were lower than predicted by the dichoptic masking function shown in Fig. A1(C). Despite this scaling, our measured thresholds are below the predicted values, and the minimum of the data occurs at different value than the predicted minimum.