Available online at www.sciencedirect.com
SCIENCE DIRECT
Vision Research 46 (2006) 57–71
Vision Research
www.elsevier.com/locate/visres

Effect of decorrelation on 3-D grating detection with static and dynamic random-dot stereograms

Stephen Palmisano [a,*], Robert S. Allison [b,c], Ian P. Howard [c]

[a] Department of Psychology, University of Wollongong, NSW 2522, Australia
[b] Department of Computer Science, York University, Ont., Canada M3J 1P3
[c] Centre for Vision Research, York University, Toronto, Ont., Canada M3J 1P3

Received 12 July 2004; received in revised form 4 October 2005

Abstract

Three experiments examined the effects of image decorrelation on the stereoscopic detection of sinusoidal depth gratings in static and dynamic random-dot stereograms (RDS). Detection was found to tolerate greater levels of image decorrelation as: (i) density increased from 23 to 676 dots/deg^2; (ii) spatial frequency decreased from 0.88 to 0.22 cpd; (iii) amplitude increased above 0.5 arcmin; and (iv) dot lifetime decreased from 1.6 s (static RDS) to 80 ms (dynamic RDS). In each case, the specific pattern of tolerance to decorrelation could be explained by its consequences for image sampling, filtering, and the influence of depth noise.

Keywords: Stereo vision; Correspondence; Stereogram; Surface detection

1. Introduction

In both static and dynamic random-dot stereograms (RDS), 3-D surface structure is visible only after the two monocular images are combined by the visual system (Julesz, 1960, 1964, 1971). In viewing such displays, the stereoscopic depth percept is based solely on the positional disparities of corresponding dots in the two eyes' images. However, there is also a potentially complex correspondence problem to be solved (for a review—see Howard & Rogers, 1995). Since the dots in these RDS are identical in contrast polarity, shape, and size, any dot in the left eye's image could be matched with numerous dots in the right eye's image. While this correspondence problem may often be eased by the presence of clusters of dots that are recognizably the same in the two eyes' images, these dot clusters are not essential for binocular matching. Julesz (1960, 1964, 1971) showed that stereoscopic depth could still be seen when these 'micropatterns' are obscured by large numbers of uncorrelated dots in one or both eyes' images. Using static RDS, which represented a central square lying either in front or behind a surround, he noted that as image decorrelation increases:

> "first the corners of the cyclopean square disappear, but a rounded off area in the centre is still perceived in depth. Loss of stereopsis gradually increases with increasing noise. More and more dots appear at other depth planes than that of the square or its surround. Finally it is impossible to detect an area in the centre as being different to the surround" (Julesz, 1971, pp. 275).

In Julesz's original demonstrations, observers had to detect the 3-D structure of surfaces represented by static RDS with various amounts of image decorrelation. However, this image decorrelation would not only have made binocular matching more challenging, but it should also have influenced stereoscopic surface detection—which requires judgments based on perceived depth and surface structure (Harris & Parker, 1994; Palmisano, Allison, & Howard, 2001). More recent research in this area has attempted to isolate the processes responsible for binocular matching by: (1) using dynamic RDS in which the locations of correlated and uncorrelated dots change continually; and (2) having observ-

* Corresponding author. Tel.: +612 4221 3640; fax: +612 4221 4163.
E-mail address: Stephenp@uow.edu.au (S. Palmisano).

0042-6989/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2005.10.005

ers detect the presence of interocular correlation rather than changes in depth (Cormack, Stevenson, & Schor, 1991, 1994; Cormack, Landers, & Ramakrishnan, 1997; Livingstone, 1996; Stevenson, Cormack, Schor, & Tyler, 1992; Tyler & Julesz, 1976, 1978). Observers were instructed to indicate which of two stimuli had the greater interocular correlation in a two-interval-forced-choice task. Since displays typically represented a frontal plane surface, stereoscopic surface detection was assumed to play only a minor role in this task. In general, these studies found that sensitivity to interocular correlation depends on a number of stimulus factors, including display duration (Tyler & Julesz, 1976, 1978), contrast (Cormack et al., 1991), dot density (Cormack et al., 1997), and distance of the surface from the plane of fixation (Stevenson et al., 1992).

There is a sizable literature on the effect of image decorrelation on binocular matching. However, the effect of image decorrelation on stereoscopic vision, which involves both binocular matching and disparity-based surface detection, has received far less attention. Julesz's original demonstrations suggest that coarse depth perception is fairly robust to this type of noise. However, it appears that image decorrelation has marked detrimental effects on fine stereopsis (stereoacuity and latency to resolve complex RDS). For example, Christophers, Rogers, and Bradshaw (1993; also cited in Bradshaw, Rogers, & De Brun, 1995) found that the latency to detect a complex spiral shape in depth almost doubled when they decorrelated their static RDS by 30%. Similarly, Cormack and colleagues (1991) found that the smallest horizontal step change in disparity which could be detected in their dynamic RDS increased by approximately a factor of 3–4 as image decorrelation increased from 10% to 70%.

In the current study, we expanded on these previous investigations: examining the effects of dot density, corrugation spatial frequency and corrugation amplitude on the detection of disparity-defined 3-D surfaces in the presence of image decorrelation. In our main experiments, RDS depicted surfaces with sinusoidal modulations in depth and we increased image decorrelation by replacing correlated dots with uncorrelated dots. This study also appears to be the first to explicitly compare the effects of image decorrelation on 3-D surface detection with static and dynamic RDS. Lankheet and Lennie (1996) describe the following differences in the experience of viewing static and dynamic RDS containing Gaussian-distributed additive disparity noise:[1]

> "It should be noted that detecting correlation in (static random-dot patterns) is quite different from detecting it in dynamic random-dot patterns. In (static random-dot patterns) the depth of individual pixels is clearly seen eventually. In dynamic random-dot patterns on the other hand, the short dot life of individual pixels makes their depth very difficult to resolve. As a result, in noisy dynamic random dot stereograms the depth of the noise itself is not perceived: rather than a cloud of points in three dimensions one perceives an uncorrelated image with little or no depth" (pp. 530).

This observation suggests that the detection of 3-D surfaces might be less affected by decorrelation noise with dynamic RDS than with static RDS. Below we outline three possible reasons why detection performance with dynamic RDS might be expected to exceed that found with static displays. The first possibility is that averaging disparity information over time acts to increase the signal-to-noise ratio for a dynamic RDS, since any spurious dot matches occurring when viewing a dynamic RDS would be uncorrelated over time (Allison & Howard, 2000). However, averaging disparity information over time would have little affect on the signal-to-noise ratio for a static RDS, because both the spurious and correct matches would be stable and correlated over time. The second possibility is based on the fact that image decorrelation will only produce stable depth noise when the RDS is static (in the case of dynamic RDS, the short dot lifetimes would make it more difficult to resolve the depth of individual dots). According to this notion, spurious matches in static RDS might be more disruptive to surface detection than spurious matches in dynamic RDS, as the stable depth noise generated by the former would be inconsistent with the perception of a smooth surface (Lankheet & Lennie, 1996). Potentially, any such advantage for dynamic RDS might be nullified by increased difficulties resolving the depths of individual signal dots. However, there is one important difference between the signal and noise dots in dynamic RDS—unlike the transient localized depths represented by noise dots, the global surface structure represented by the signal dots is stable and supported over time. Thus, it is possible that the short dot-lifetimes in dynamic RDS might minimise the effects of local depth noise, but leave the extraction of the global surface structure relatively unimpaired. Finally, the third possibility is that detection performance might be more tolerant to image decorrelation with dynamic RDS, because these displays should have a higher effective density than a static RDS with the same instantaneous dot density—assuming that the dynamic RDS is viewed for a sufficiently long period and the dot lifetime is shorter than the visual integration time (e.g., 26 ms—Lankheet & Lennie, 1996). If true, one might expect differences between static and dynamic RDS to be maximal for sparse, high spatial frequency corrugation displays—as the multiple surface

[1] In the case discussed by Lankheet and Lennie (1996), all of the dot pairs in their RDS were correlated and originally represented a smooth sinusoidal surface in depth. When Gaussian distributed disparity noise was added to these correlated dots, the result was that the stereo-defined surface appeared jagged—at least when static RDS were used—with the amount of jaggedness depending on the amplitude of this depth noise. Conversely, in the current study, our displays consisted of a mixture of correlated dots (whose disparities represented a smooth sinusoidal surface) and uncorrelated dots. Spurious matches of non-corresponding dots could, however, have indirectly generated depth noise, which would have been very similar to the effects of this additive disparity noise.

features would each be represented by far fewer disparity samples in the static case and image decorrelation would be more likely to result in undersampling and depth noise.

2. Experiment 1: Effects of image density and decorrelation on 3-D surface detection with static RDS

One main goal of this study was to examine the effects of image decorrelation on 3-D surface detection with comparable static and dynamic RDS. However, static and dynamic RDS with the same instantaneous dot densities can have very different effective, or perceived, dot densities. So before comparisons could be made between static and dynamic RDS, we had to determine the effects of physical dot density on the tolerance to image decorrelation in static RDS. This experiment examined four physical display dot densities (23, 89, 178, and 676 dots/deg^2). For each density condition, the total number of dots in each half image remained constant (at either 1831, 7188, 14,412, or 54,746) as the image decorrelation increased from 0% to 100%.

2.1. Method

2.1.1. Observers

Three observers participated in Experiments 1 and 2; the first author (SAP), and two observers DH and MH (who were naive to the experimental hypotheses). Two additional participants, XF and HJ also participated in the control experiments reported in this paper. Experiments were undertaken with the understanding and written consent of each observer. All observers (aged between 30 and 41 years) had participated in many previous experiments on stereoscopic surface detection. They had normal or corrected-to-normal vision and a stereoacuity of at least 20 s of arc (''Randot'' stereovision test). Each observer was given several 100 test trials before their experimental data were collected.

2.1.2. Apparatus

Static RDS were generated on a Macintosh G4 and presented in a Wheatstone stereoscope. Computer monitors (Apple Multiple Scan M2978, 1024×768 pixels, 75 Hz) were placed one to the left and one to the right of the observer and viewed through mirrors mounted at $\pm 45°$ to the frontal plane. The observer's head was restrained by a head–chin rest. The viewing distance was 84 cm. The convergence angle of the stereoscope and the monitor orientations were appropriate for this distance. Each screen subtended 14° in height by 18° in width. The stimuli were presented in a dark room and all surfaces were covered with matte black cloth, cardboard or paint. Each eye saw only one screen and black cardboard apertures blocked the view of the monitor's frame.

2.1.3. Stimuli

Static RDS consisted of two antialiased stereo half images produced by symmetrical oversampling and decimation (see Fig. 1). Each half image subtended a square $9° \times 9°$ area and had one of four different dot densities: 23 (1831 dots), 89 (7188 dots), 178 (14,412 dots) or 676 dots/deg^2 (54,746 dots). Due to random dot overlaps, the average dot density for each condition corresponded to 1%, 3%, 6% or 24% of the display, respectively. All of the dots were blue and subtended an area of 4 arcmin^2. Peak luminance at the center of each antialiased dot was 52 cd/m^2 and the average luminance of the dark background was 0.2 cd/m^2. As dot density increased from 23 to 676 dots/deg^2, the average luminance of a 1° area of the display increased from approximately 0.6 to 12 cd/m^2.

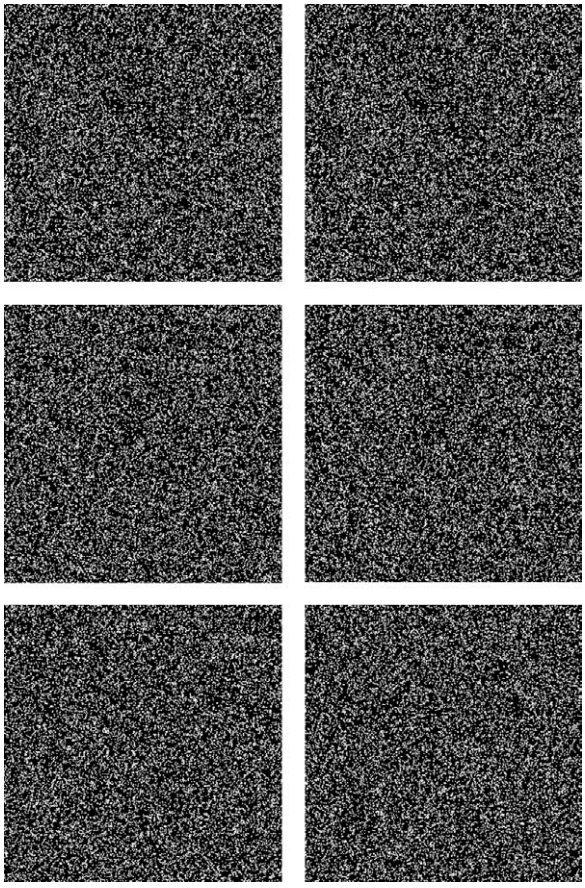

Fig. 1. Random-dot stereogram (RDS) pairs representing examples of the 0.22 cpd stimuli used in Experiment 1. When cross-fused, they portray sinusoidal depth gratings with various levels of image decorrelation. Top, 0% image decorrelation (i.e., pure signal); middle, 30% image decorrelation; bottom, 100% image decorrelation.

The RDS used in these experiments were of two kinds.

1. Each 'signal + noise' display represented a surface with sinusoidal modulations in depth (horizontally oriented ridges), which occurred at one of three spatial frequencies (0.22, 0.44, or 0.88 cpd). The sinusoid's phase varied randomly from trial to trial. In the case of a pure signal display (0% image decorrelation), each dot in the left eye's image had a dot in the corresponding location in the right eye's image. Horizontal disparities were applied

to these corresponding dots by shifting them in opposite directions in the left and right stereo half-images (disparity ranged from +2 to −2 arcmin). For the remainder of the 'signal + noise' displays, 10–90% of the dots in each half-image were uncorrelated. Image decorrelation was increased by replacing randomly selected correlated dot pairs with pairs of uncorrelated dots (one dot to each eye), rather than by adding additional uncorrelated dots to the display.

2. 'Noise' displays were identical to the 'signal + noise' displays, except that 100% of the dots in the left and right eyes images were uncorrelated.

2.1.4. Procedure

Observers were informed that they would be shown a series of displays depicting a surface with sinusoidal modulations in depth (similar to a 'corrugated tin roof' with either 2, 4 or 8 troughs and peaks) and distracter stimuli appearing as either a plane or a 3-D volume. They were instructed to fixate a 12 arcmin cross for each display and indicate whether or not they saw a surface with modulations in depth. Following these instructions and the presentation of sample stimuli, observers commenced the experiment by pressing the space bar on the keyboard. As soon as they had clearly fused each stereogram, observers indicated whether or not the display appeared corrugated by pressing one of two buttons (''yes'' or ''no''). The stereogram was displayed until a response was recorded and then the monitor turned black for 2s. This intertrial interval reduced the likelihood of afterimages and disparity aftereffects. After completing several practice blocks, observers ran twenty-four experimental blocks (six replications of each of the four dot density conditions). Within each block (660 trials), equal numbers of the following conditions were presented in random order: (i) 'signal + noise' and 'noise' displays; and (ii) 0.22, 0.44 and 0.88 cpd corrugation displays. Each noise-level by spatial-frequency condition was presented 10 times per block.

2.1.5. Analyses

Each observer's ''yes'' responses in the presence or absence of a stereoscopically defined depth grating were converted into hit rates (H) and false alarm rates (F), respectively. These rates, expressed as probabilities ranging between 0.0 and 1.0, were then converted into z-scores and used to calculate d prime $\{d' = z(H) - z(F)\}$. Ninety-five percent confidence intervals were then used to determine whether d' values varied significantly across the different experimental conditions. These confidence intervals $\{\mathrm{CI}(d')\}$ were calculated as follows:

$$\mathrm{var}(d') = H(1-H)/N_H\left[\phi(H)\right]^2 + F(1-F)/N_F\left[\phi(F)\right]^2,$$
$$\mathrm{CI}(d') = 1.95 \times \left[\mathrm{var}(d')\right]^{1/2} \quad (1)$$

where N_H = number of signal + noise displays, N_F = an equivalent number of noise displays, $\phi(\mathrm{H}) = 2\pi^{-1/2}\exp[-0.5z(H)^2]$, and $\phi(F) = 2\pi^{-1/2}\exp[-0.5z(F)^2]$ (MacMillan & Creelman, 1991).

2.2. Results and discussion

Stereoscopic surface detection was found to be remarkably robust in the presence of substantial image decorrelation (see Fig. 2)—performance only fell to chance when 70–90% of the dots in the two half-images were uncorrelated.[2] The remarkable tolerance found for image decorrelation was consistent with the findings of Cormack and colleagues (1991) on the ability to detect a step edge in depth from a decorrelated RDS. Extrapolating from their data, the minimum disparity required to correctly detect a step edge in depth at 90% image decorrelation should lie between 1.7 and 2.5 arcmin (for observers' SBS and LKC, respectively). In the current experiment, sinusoid detection performance fell to chance at 90% decorrelation for our densest, 2 arcmin amplitude displays. The similarity of our findings to those of Cormack et al's was quite surprising since our displays had the following stimulus characteristics which were expected to impair stereoscopic detection: (i) they were static as opposed to dynamic RDS; and (ii) they were sparser (ranging from 1% to 24% dot density) than those used by Cormack et al. (50% dot density). However, our static RDS displays provided multiple disparity defined surface features, as opposed to a single step in disparity, which might have compensated for our displays being static and comparatively sparse.

Detection performance in the presence of 10–80% image decorrelation was found to consistently improve as the display density increased from 23 to 676 dots/deg^2 (see Fig. 2). All three observers demonstrated significantly greater tolerance to 10–80% image decorrelation for 676 dots/deg^2 displays compared to 23 dots/deg^2 displays [with d' differences and confidence intervals of 0.8 ± 0.25 (SAP), 0.9 ± 0.25 (MH), and 0.8 ± 0.25 (DH)] (see Figs. 2 and 4). In principle, the current improvements in observer tolerance to image decorrelation with increasing density could have arisen because: (i) the surface troughs and peaks were defined by more disparity samples in high density displays, which in turn would have produced a more compelling percept of the 3-D surface; (ii) the effective signal-to-noise ratio was greater in high density displays (while the absolute proportions of signal dots to noise dots would have

[2] It was conceivable that observers could theoretically have performed the task by detecting the correlation of dots in the two eyes' images rather than by detecting a coherent surface. However, we also ran control experiments where the dots in the distractor displays had the same correlation as the target. Under these conditions, simple detection of correlation was insufficient to perform the task. We found that when the distractor stimulus was either a frontal plane with disparity noise or a volume of dots with the same depth range as the sinusoidal target, the pattern of results was similar to that found with the fully decorrelated distractor.

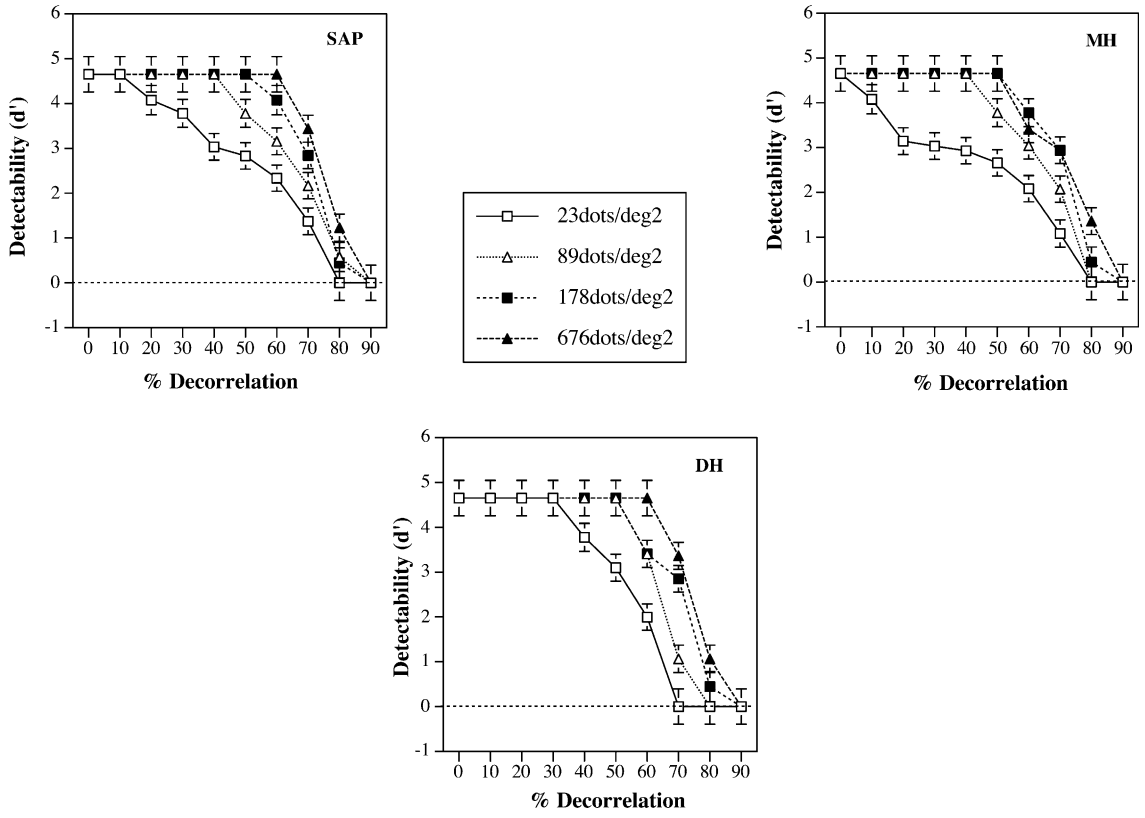

Fig. 2. Effect of display density (23, 89, 178, and 676 dots/deg^2) on sinusoid detection from static RDS with 0–90% image decorrelation for three observers (SAP, MH and DH). The d' values for each density condition shown in this graph were produced by pooling the data from the three different spatial frequency conditions. Within each density condition, correlated and decorrelated displays had identical numbers of dots in each half-image. Error bars represent standard errors of the mean [Experiment 1].

remained constant, the effectiveness of signal filtering or pooling would have improved as the numbers of signal dots increased); and/or (iii) the average amplitude of the stable depth noise produced by spurious matches would have been less for high density displays (while the number of spurious matches would actually increase with the display density, it would become progressively more common to find a nearby match).

It is interesting to compare the above effects of density on stereoscopic surface detection with previous reports of the effects of density on interocular correlation detection. While we found that detection of decorrelated sinusoids consistently improved as dot density increased from 1% up to 24%, Cormack et al. (1997) found that interocular correlation detection gradually improved as display density increased from 0.1% to 2%, but showed little change at higher dot densities (up to 50%). According to the sampling explanation,[3] the different effects of density in two studies could be taken as evidence that fewer correlated dots were required to detect interocular correlation than were required to detect multiple corrugations in depth. Observers in the earlier Cormack et al. study might have undersampled their briefly viewed displays (180 ms), matching sufficient numbers of dots to detect the presence of interocular correlation, but not enough to reconstruct a coherent surface (in this case a frontal plane). Conversely, due to the longer (participant-defined) display durations used in the current experiment, our observers might have been able to match progressively more correlated dots as the display density increased, which in turn would have produced a more compelling percept of the 3-D surface and a greater tolerance to image decorrelation.

Support for the notion that stereoscopic surface detection and interocular correlation detection have different sampling requirements was provided by the results of a control experiment. This control examined the detection performance (of SAP and two naïve observers XF and HJ) for decorrelated sinudoidal surfaces and decorrelated frontal planes (see Fig. 3). For all three observers, the ability to detect sinusoidal disparity gratings was less affected by 10–50% image decorrelation than their ability to detect frontal plane surfaces. Conversely, we found that their ability to detect frontal plane surfaces was less affected by 70–

[3] There were a number of other differences between these two experiments that could also have accounted for this discrepancy. Our displays were static, remained visible until the observer responded (which took 0.5–0.9 s on average), subtended a visual angle of $9° \times 9°$, and consisted of between 1831 and 54,746 dots. Conversely, Cormack et al.'s (1997) displays were dynamic (12 frames at 67 frames/s), lasted only 180 ms, subtended a visual angle of $2° \times 2°$, and consisted of between 1 and 200 dots.

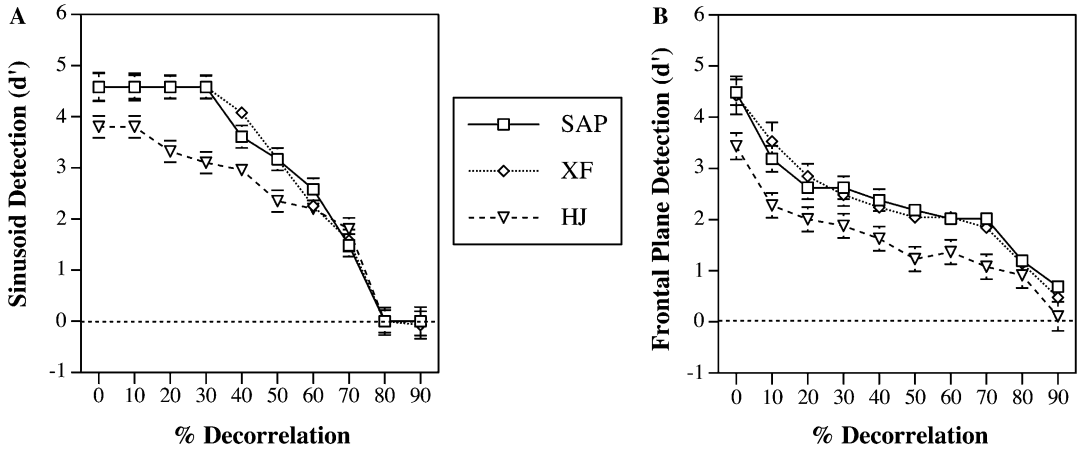

Fig. 3. A comparison of the detection performance of three observers (SAP, XF and HJ) for decorrelated sinusoidal (A) and frontal plane (B) surfaces with equivalent dot densities. In the case of the sinusoidal surfaces—the data from the 0.22 cpd corrugation spatial frequency condition is presented [control experiment].

90% image decorrelation than their ability to detect disparity gratings. Our interpretation of these results is as follows: At low levels of image decorrelation, detection judgments were aided and dominated by the presence of surface structure. Since the sinusoid's locally smooth depth modulations were more salient than the zero-depth structure of the frontal plane surfaces, a detection advantage was found for sinusoidal displays {support for this claim comes from a study by Palmisano, Allison, and Howard (2000) which found that human detection performance was more efficient for disparity-defined sinusoids than for square waves}. Conversely, at high levels of image decorrelation, detection judgments were based predominately on the presence of interocular correlation (as it became increasingly difficult to extract surface structure from the RDS). Since the detection of interocular correlation should have been more straightforward for frontal plane stimuli than for sinusoidal stimuli,[4] a performance advantage was found for frontal planes at high levels of decorrelation. Thus, while the detection of a frontal plane might be regarded as a reasonably pure measure of binocular matching, we argue that additional post-matching processing was required to perceive a sinusoidal surface (consistent with results reported by Harris & Parker, 1994). The above findings confirm that our sinusoidal depth gratings were the appropriate stimuli to investigate the processes involved in stereoscopic surface detection (i.e., not just a subset of these processes).

[4] In principle, one could build a correlation detector that examined matching statistics and signalled if the correlation observed was greater than expected from random variation (i.e., no surface structure need be inferred). In the frontal plane situation, accidental matches should occur at a variety of lags, but true matches should occur only at zero lag. Hence one could pool zero lag correlators over the entire image. By contrast, correlations in sinusoidal stimuli would occur over a range of lags (disparities) and without analysing their spatial coherence they would be difficult to distinguish from accidental matches. Pooling correlators over space without taking into account their spatial relations would be much less effective than in the frontal plane case.

In our main experiment, tolerance to image decorrelation was also found to vary modestly with the spatial frequency of the depth grating (see Fig. 4). All three observers (SAP, MH, DH) had significantly higher sensitivities to 0.22 cpd displays than to 0.88 cpd displays in the presence of 10–80% image decorrelation [d' differences of 0.63 ± 0.25 (SAP), 0.73 ± 0.25 (MH) and 0.4 ± 0.25 (DH)]. However, sensitivities to 0.22 cpd displays were not significantly different from those for 0.44 cpd displays in the presence of 10–80% image decorrelation [d' differences of 0.2 ± 0.25 (SAP), 0.18 ± 0.25 (MH) and 0.16 ± 0.25 (DH)]. This effect of spatial frequency on the detection of partially decorrelated depth gratings appeared to be quite similar to its effect on the minimum disparity required to detect fully correlated depth gratings (Bradshaw & Rogers, 1999; Rogers & Graham, 1982; Tyler, 1974). For example, Rogers and Graham (1982) found that sensitivity to depth gratings peaked with corrugation frequencies between 0.2–0.4 cpd and fell off at lower and higher frequencies. Thus, it is possible that the present spatial frequency effect was simply a reflection of the disparity sensitivity function found previously with displays at the disparity threshold.

However, it is also possible that the detection of decorrelated high spatial frequency sinusoids was more easily disrupted because their troughs and peaks were defined by fewer dot pairs than those in low spatial frequency corrugations. This was due to the fact that within each experimental block, density was held constant for all displays, irrespective of the number of surface features represented. A classical result in signal theory known as the Nyquist–Shannon sampling theory (Shannon, 1949) specifies that, for unambiguous reconstruction, a signal must be sampled at a rate of at least twice its highest frequency component. Recently, it has been demonstrated that humans can resolve stereoscopic gratings with corrugation frequencies approaching the Nyquist limit in sparse random dot stereograms (Banks, Gepshtein, & Landy, 2004). In the present conditions, we deliberately chose modest spatial frequencies and sufficiently dense RDS so that sampling was

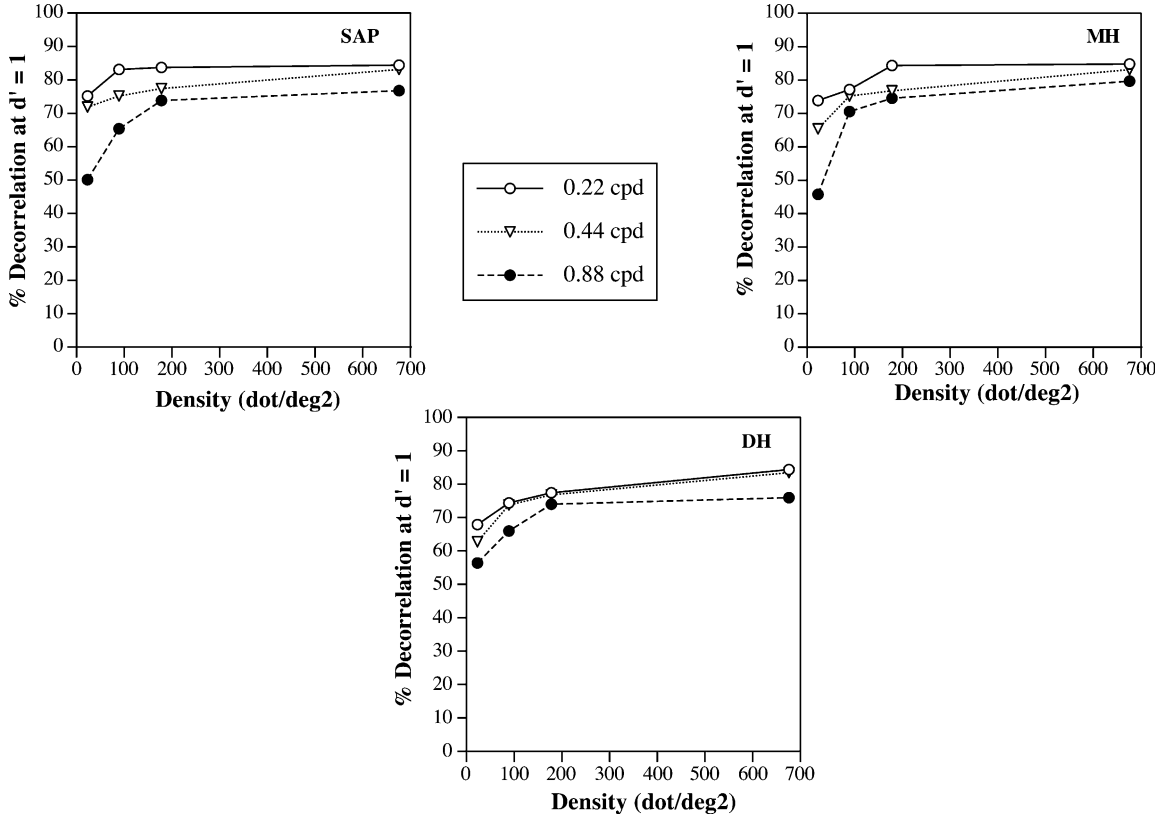

Fig. 4. Effects of dot density and corrugation spatial frequency on sinusoid detection from decorrelated static RDS (SAP, MH, and DH). The figure shows the level of decorrelation which produced d' values of 1 for each of the dot density and spatial frequency conditions examined [Experiment 1].

always several times higher than the Nyquist rate in correlated conditions. Of course, as the dots defining the stereoscopic surface were replaced by uncorrelated dots the effective sampling rate would have declined (to zero when the dots were completely decorrelated). There was only one condition—with the highest spatial frequency (0.88 cpd), lowest density (26 dots per deg^2) and maximum 90% decorrelation—where the signal would have been strictly undersampled (see Fig. 5). However, this analysis depends on the observer correctly matching all of the correlated signal dots and ignoring the uncorrelated noise dots. Matches involving uncorrelated dots would have introduced depth noise in these static RDS and the failure to match correlated pairs would have reduced the effective sampling density.

In summary, the effects of display density and corrugation spatial frequency on the tolerance of human stereoscopic surface detection to image decorrelation could be explained by either disparity/depth noise or sampling considerations. Increasing the density of RDS displays could have improved tolerance to decorrelation by: (i) reducing the likelihood of (effective) undersampling; (ii) increasing filtering effectiveness; and/or (iii) reducing the amplitude of depth noise. Since the effects of depth noise should be modulated by the spatial frequency selectivity of stereopsis and (effective) undersampling should be apparent at high spatial frequencies before low spatial frequencies, both of these factors could also have made it more difficult to perceive smooth peaks or troughs from high spatial frequency displays as image decorrelation increased.

3. Experiment 2: Effects of corrugation amplitude and decorrelation on 3-D surface detection with static RDS

Experiment 1 revealed that the tolerance of stereoscopic surface detection to image decorrelation in static RDS increased as dot density increased and stimulus spatial frequency decreased. However, the depth corrugations in these static displays always had an amplitude of 2 arcmin. Experiment 2 re-examined the tolerance of stereoscopic surface detection to image decorrelation using static displays with four different corrugation amplitudes: 0.5, 1, 2 or 3 arcmin. In both of the density conditions examined (23 and 676 dots/deg^2), the total number of dots in each half image remained constant at either 1831 or 54,746 as image decorrelation increased from 0% to 100%.

3.1. Method

The observers, apparatus, stimuli and procedure were identical to those of Experiment 1, except that four corrugation amplitudes were examined (as opposed to only one). Experimental blocks were run at least six times in a random order—each examined detection performance for a specific

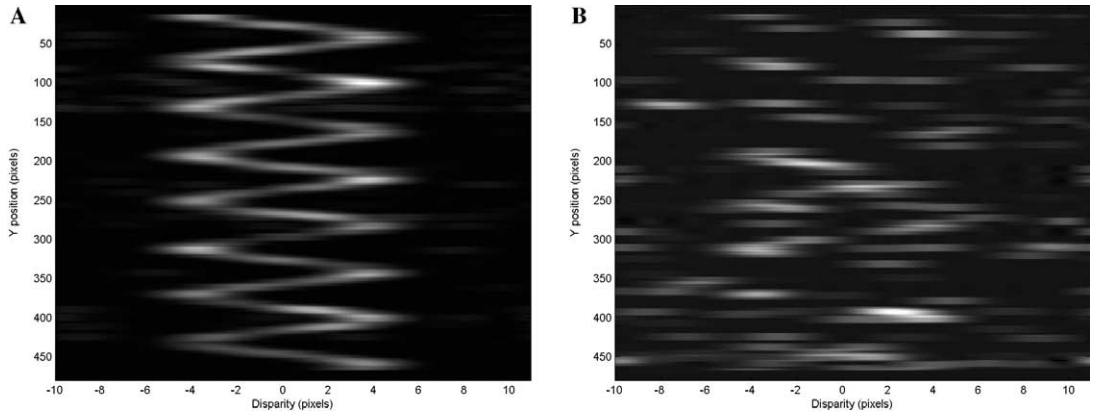

Fig. 5. Effect of image decorrelation on binocular matching by interocular correlation using the square 480×480 pixel images from Experiment 1 (in terms of visual angles these images corresponded to $9° \text{ H} \times 9° \text{ V}$ at the viewing distance used in our experiments). Each panel displays correlation as a function of the position of an 8×8 pixel window in the left and right eyes. The model was run at every location in the image (except the borders), with each eye's window being centred on the same line. Summed correlation at each location was represented by intensity, with brighter values indicating higher levels of correlation. (A) 0% decorrelation, the window correlator algorithm produced a clear sinusoidal ridge of high correlation corresponding to the disparity signal (which in this case was a 0.88 cpd RDS with an amplitude of 2′ and a density of 23 dots/deg^2). However, in (B) a noisy version of this signal condition (90% image decorrelation) was poorly represented by the model.

corrugation amplitude (0.5, 1, 2 or 3 arcmin), display density (either 23 or 676 dots/deg^2) and corrugation spatial frequency (0.22 and 0.88 cpd). Within each of block, equal numbers of the 'signal + noise' and 'noise' displays conditions were presented in a random order (each condition was presented 10 times per block).

3.2. Results and discussion

Stereoscopic surface detection was perfect or near perfect for pure signal displays at the smallest corrugation amplitude examined (0.5 arcmin)—indicating that these disparity-defined corrugations were well above detection threshold for all observers. As the level of image decorrelation increased (i.e., 10–80%), all three stimulus manipulations (density, spatial frequency and amplitude) were found to produce significant differences in detection performance. Detection performance was significantly less tolerant to 10–80% image decorrelation for the 0.5 arcmin corrugation amplitude compared to the average detection performance for the three larger corrugation amplitudes (1, 2 and 3 arcmin), indicated by d' differences of -0.4 ± 0.2 (SAP), -0.5 ± 0.2 (MH) and -0.25 ± 0.2 (DH) (see Fig. 6). On face value, this amplitude finding appears inconsistent with the notion that the visual system detects disparity by measuring correlation over a finite, frontoparallel area (e.g., Banks et al., 2004)—which would predict that disparity detection based on correlation signals should decline as the disparity gradients increased. However, the relatively low spatial frequencies and modest amplitudes of our target displays ensured that the peak disparity gradient was modest and much lower than classical disparity gradient limits. Further, the use of sinusoidal gratings ensured that much of the disparity change in the stimulus was well below even these modest disparity gradients. We argue below that the observed effect of corrugation amplitude could have been produced if stereoscopic surface detection was more susceptible to disparity noise in the case of the smaller amplitude display.

As in the previous experiment, all three observers demonstrated significantly higher sensitivities to 0.22 cpd displays than to 0.88 cpd displays in the presence of 10–80% image decorrelation [d' differences of 0.6 ± 0.17 (SAP), 0.65 ± 0.17 (MH) and 0.5 ± 0.17 (DH)]. Similarly, all three observers demonstrated significantly greater tolerance to 10–80% image decorrelation for 676 dots/deg^2 displays compared to 23 dots/deg^2 displays [d' differences of 0.5 ± 0.17 (SAP), 0.6 ± 0.17 (MH) and 0.4 ± 0.17 (DH)]. However, the effects of display density on detection also appeared to interact with those of corrugation amplitude. Detection performance for *dense 0.5 arcmin* displays was substantially more tolerant to 10–80% image decorrelation than that for *sparse 0.5 arcmin* displays [d' differences of 0.9 ± 0.35 (SAP), 0.83 ± 0.35 (MH) and 0.54 ± 0.35 (DH)]. By comparison, detection performance for *dense 1–3 arcmin* displays was only modestly more tolerant to image decorrelation than that for *sparse 1–3 arcmin* displays [d' differences of 0.4 ± 0.3 (SAP), 0.51 ± 0.3 (MH) and 0.4 ± 0.3 (DH)].

It is possible that the above interaction between static display density and corrugation amplitude occurred because the *sparse 0.5 arcmin RDS* were more susceptible to the stable depth noise produced by spurious dot matches than RDS with larger corrugation amplitudes. For example, if this stable depth noise often approached or exceeded the amplitude of the 0.5 arcmin corrugation, then it is possible that more correlated dots would have been required to produce the percept of a smooth, continuous surface. The improvements found in noise tolerance for these 0.5 arcmin RDS as density increased, could have been due to either an

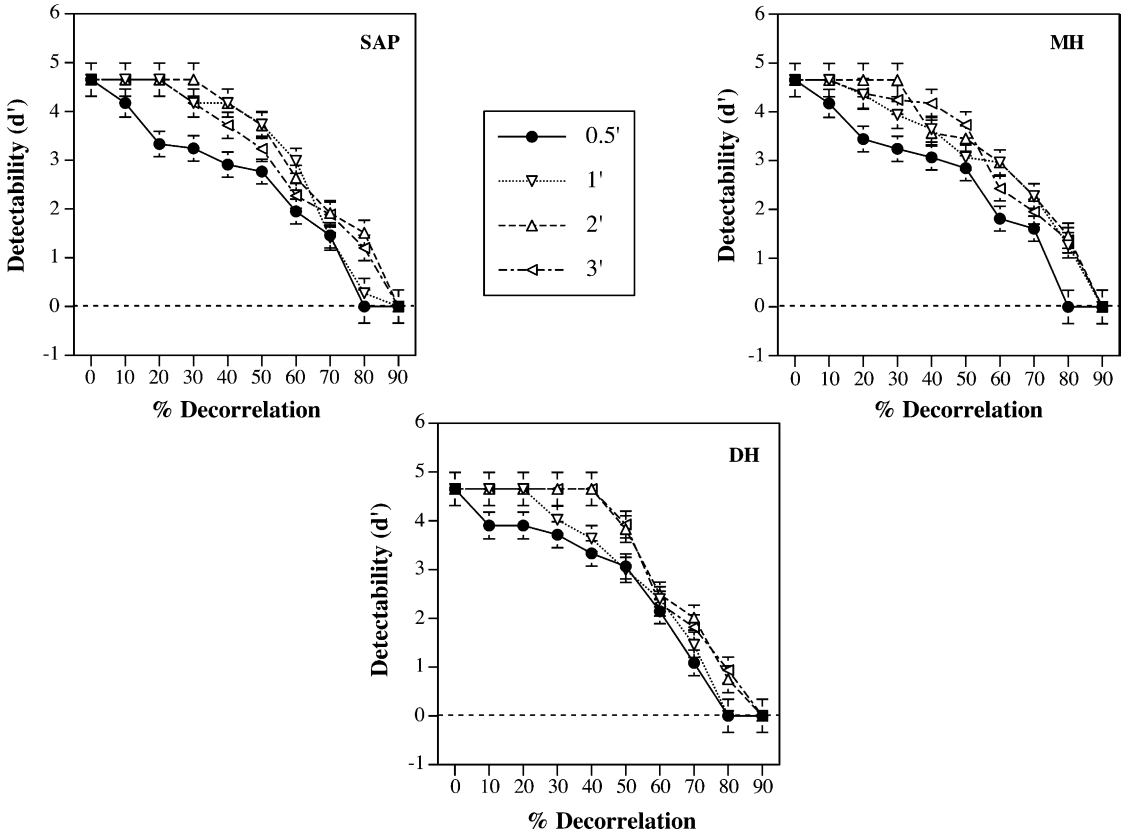

Fig. 6. Effect of corrugation amplitude (0.5, 1, 2, or 3 arcmin) on sinusoid detection from static RDS with 0–90% image decorrelation for three observers (SAP, MH and DH). The d' values for each corrugation amplitude condition shown in this graph were produced by pooling the data from the different dot density and corrugation frequency conditions. Error bars represent standard errors of the mean [Experiment 2].

increase in the effective signal-to-noise ratio (due to improved filtering precision) and/or a reduction in the amplitude of the disparity noise (due to the increased likelihood of a nearby spurious match). Since sparse-static displays with larger corrugation amplitudes should have been less susceptible to the effects of large amplitude depth noise, the improvements provided by increasing the display density would have been more modest for these displays.

4. Experiment 3: 3-D surface detection with decorrelated static and dynamic RDS

In Experiment 3, we examined whether stereoscopic surface detection was more robust in the presence of image decorrelation using dynamic, rather than static, RDS. We manipulated the effects of dot lifetime, display density, corrugation spatial frequency, and corrugation amplitude. As proposed in the introduction, if dynamic RDS are viewed for a sufficiently long period, detection performance might be expected to exceed that found with static displays due to one or more of the following reasons: (i) averaging over time during a dynamic RDS should increase the signal-to-noise ratio of the display; (ii) the stable depth noise produced by spurious matches in static RDS might be more disruptive to surface detection than the transient noise effects produced by dynamic RDS; and (iii) if dot lifetimes were shorter than the visual integration time, the effective display density should be higher than that of a static display with the same instantaneous dot density.

4.1. Method

4.1.1. Observers

Observer DH was replaced by a naïve observer MEL (30 years of age), who met the observer requirements mentioned previously.

4.1.2. Stimuli

Static RDS and the individual 'frames' of dynamic RDS were identical to the displays used in the previous experiment: they had one of two physical dot densities (23 or 676 dots/deg^2) and one of two corrugation amplitudes (0.5 or 2 arcmin). In the case of dynamic RDS,[5] all dot positions and their binocular disparities (in the case of correlated dots) were revised every 26 or 80 ms. Both

[5] Reducing the dot lifetime of a dynamic RDS reduced the effective/perceived display contrast relative to static RDS. However, despite this potential confound, a detection advantage was still found for dynamic RDS over static RDS.

static and dynamic RDS displays were presented for a fixed display duration of 1.6 s. To achieve this constant duration, the number of 'frames' presented increased from 20 to 60, as the dot-lifetimes of dynamic RDS decreased from 80 ms down to 26 ms. Experimental blocks were run at least 25 times in a random order—each examined detection performance for a specific dot lifetime (1.6 s, 80 ms, or 26 ms), corrugation amplitude (0.5, 1, 2 or 3 arcmin), display density (either 23 or 676 dots/deg^2) and corrugation spatial frequency (0.22 and 0.88 cpd). Within each block, equal numbers of 'signal + noise' and 'noise' displays conditions were presented in a random order (each condition was presented four times per block).

4.2. Results and discussion

Overall, all three observers were significantly more tolerant to 10–80% image decorrelation in dynamic RDS than in static RDS [d' differences of 0.7 ± 0.14 (SAP), 0.5 ± 0.14 (MH) and 0.52 ± 0.12 (MEL)] (see Figs. 7 and 8). Contrary to the prediction that reducing dot lifetime from 80 to 26 ms would increase the signal-to-noise ratio of dynamic displays, this manipulation was found to produce no further improvement in noise tolerance. Specifically, detection performance with 26 ms dynamic RDS was not found to be significantly different to that found with 80ms dynamic RDS in the presence of 10–80% decorrelation [d' differences of 0.02 ± 0.1 (SAP), 0.07 ± 0.1 (MH) and 0.05 ± 0.1 (MEL)]. However, as performance with 80ms dynamic RDS was near to that found for the highest density displays, this null finding might indicate a ceiling effect in the sampling of stereoscopic surfaces.

Consistent with the findings of the earlier experiments, detection performance with static RDS containing 10–80% image decorrelation improved as: (i) the physical density increased from 23 to 676 dots/deg^2 [dense versus sparse d' differences were 0.6 ± 0.17 (SAP), 0.7 ± 0.19 (MH) and 0.7 ± 0.16 (MEL)—see Fig. 8]; (ii) the spatial frequency decreased from 0.88 to 0.22 cpd [low versus high spatial frequency d' differences were 0.6 ± 0.2 (SAP), 0.6 ± 0.2 (MH) and 0.53 ± 0.16 (MEL)—see Fig. 8]; and (iii) as the corrugation amplitude increased above 0.5 arcmin [2 arcmin versus 0.5 arcmin d' differences were 0.6 ± 0.2 (SAP), 0.7 ± 0.2 (MH) and 0.3 ± 0.16 (MEL)].

As expected, increasing the density of static RDS improved tolerance to image decorrelation to a greater extent than increasing the density of dynamic RDS. Detec-

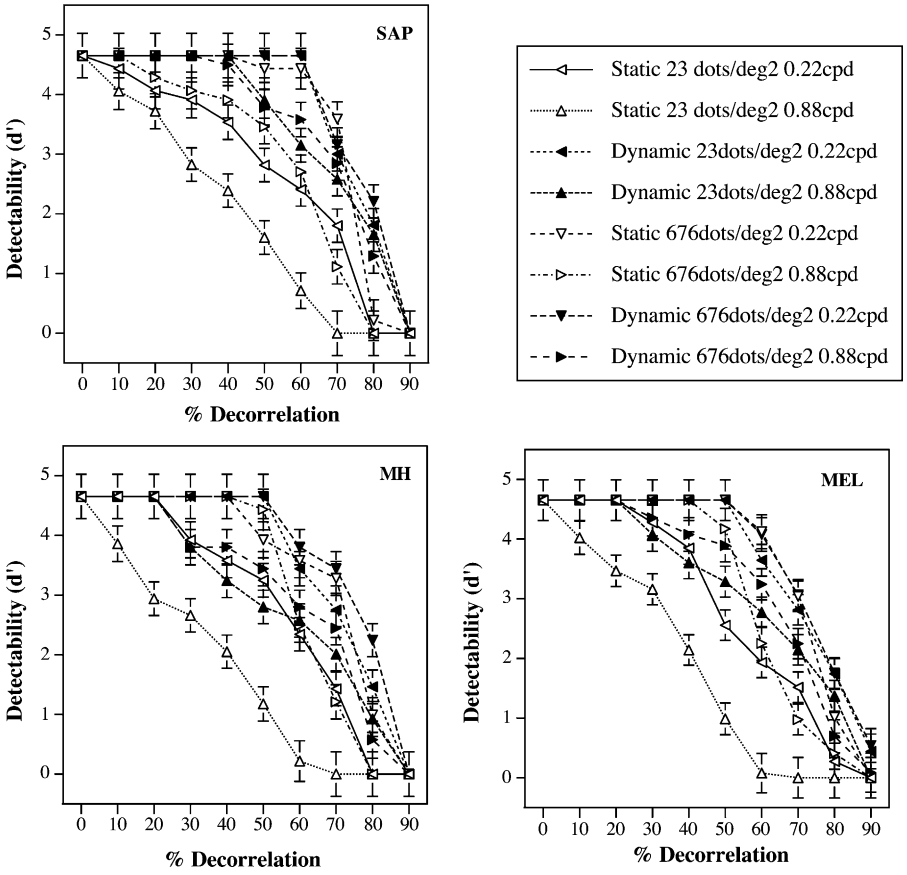

Fig. 7. Effect of display density (23 or 676 dots/deg^2) corrugation spatial frequency (0.22 or 0.88 cpd) on sinusoid detection from static and dynamic RDS with 0–90% image decorrelation for three observers (SAP, MH and MEL). Error bars represent standard errors of the mean [Experiment 3].

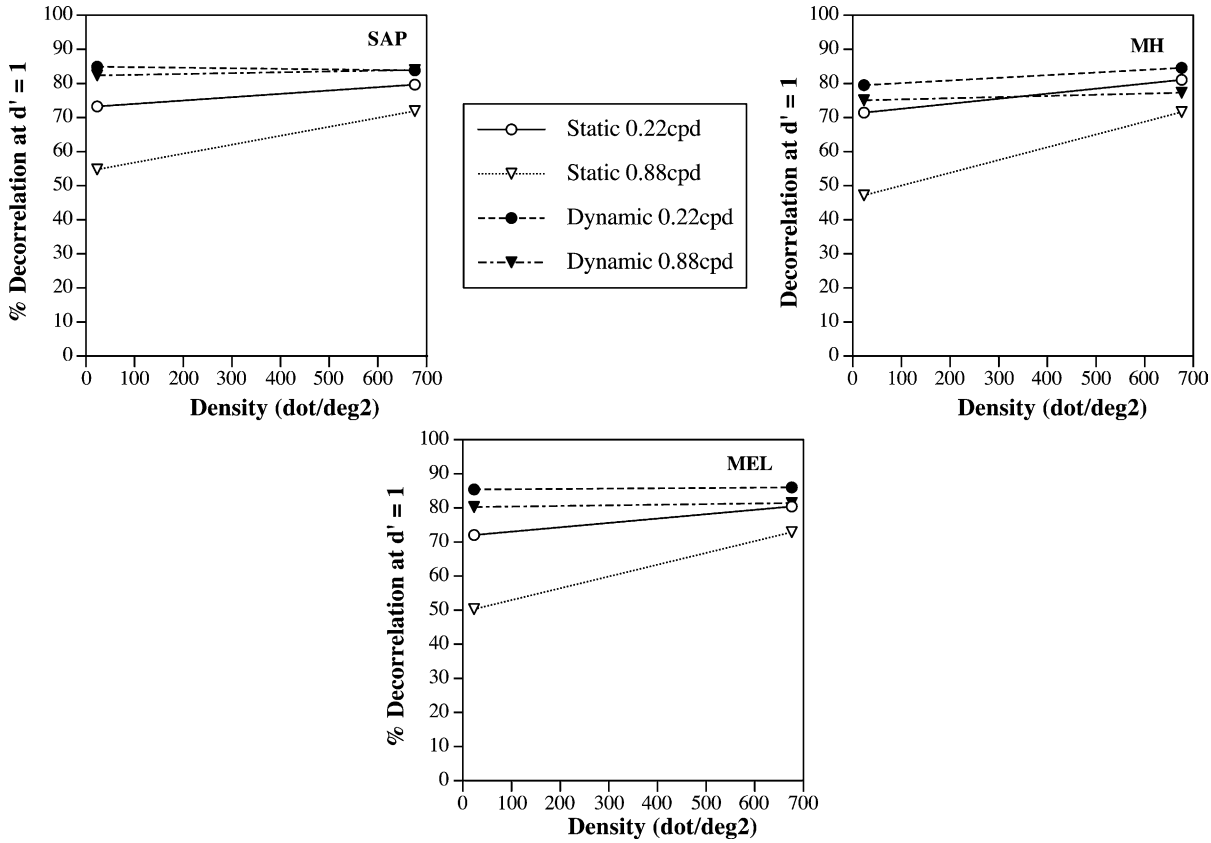

Fig. 8. Effects of dot density (23 or 676 dots/deg^2), display type (static or dynamic RDS) and corrugation spatial frequency (0.22 or 0.88 cpd) on sinusoid detection from decorrelated RDS (SAP, MH, and MEL). The figure shows the level of decorrelation which produced d' values of 1 for each of the dot density, display type and spatial frequency conditions examined [Experiment 3].

tion was substantially more tolerant to 10–80% image decorrelation with *sparse-dynamic* RDS than with *sparse-static* RDS. However, only two of the observers were significantly more tolerant to 10–80% decorrelation for dense-dynamic RDS than for dense-static RDS. In Experiment 2, we argued that tolerance to image decorrelation was reduced for *sparse-small-amplitude* displays because observers were particularly susceptible to the stable depth noise produced by these static displays (which would have more often approached or exceeded the corrugation amplitude). Consistent with the notion that the transient noise effects produced by dynamic RDS were less disruptive than the stable depth noise produced by static RDS, no such interaction between corrugation amplitude and display density was found for dynamic RDS.

Across all three observers, tolerance to 10–80% image decorrelation was found to be quite similar for *sparse-dynamic* and *dense-dynamic* displays [d' differences of 0.07 ± 0.2 (SAP), 0.2 ± 0.2 (MH) and 0.09 ± 0.16 (MEL)—see Fig. 8]. This suggests that: (i) the increase in effective density produced by reducing the dot life time of *sparse-dynamic* RDS was sufficient to improve performance to near ceiling levels; and (ii) *dense-dynamic* displays might have posed a more serious correspondence problem than *sparse-dynamic* displays.

In Experiments 1 and 2, we found the following spatial frequency by physical density interaction: (i) *sparse-high-spatial-frequency* RDS were more susceptible to decorrelation noise than *dense-high-spatial-frequency* RDS; (ii) however, *low-spatial-frequency* RDS were quite tolerant to decorrelation noise across the range of display densities. Based on this interaction, we might also expect that observers would become more tolerant to decorrelation noise for *high-spatial-frequency* displays when the dot lifetime was reduced (from 1.6 s down to 80 or 20 ms), as this would increase the effective density of the RDS. Consistent with this proposal, we found that reducing the dot lifetime substantially improved detection for *sparse 0.88 cpd* RDS, but had less marked effect on detection for *dense 0.88 cpd* RDS (see Fig. 8).

5. General discussion

Consistent with previous research (Cormack et al., 1997; Julesz, 1960, 1964, 1971), stereoscopic surface detection was found to tolerate substantial image decorrelation of both static and dynamic RDS. However, the extent of this tolerance to image decorrelation was shown to depend on a number of stimulus characteristics. Specifically, we found that stereoscopic surface detection improved as: (i) the density of static RDS increased from 23 to 676 dots/deg^2

(Experiments 1–3); (ii) the corrugation spatial frequency decreased from 0.22 to 0.88 cpd (Experiments 1–3); (iii) the amplitude of the depth corrugation of static (but not dynamic) RDS increased above 0.5 arcmin (Experiments 2–3); and (iv) the dot lifetime decreased from 1.6 s to 80 ms (holding display duration constant at 1.6 s—Experiment 3).[6] Also, interactions were found between these different stimulus factors. Most importantly, the greatest improvements in detection performance were found when dot lifetime was reduced from 1.6 s to 80 ms if the display was sparse (23 dots/deg^2) and the corrugation spatial frequency was high (0.88 cpd).

We argue that the above patterns of tolerance to image decorrelation can be explained by the following two factors. First, spurious matches between correlated and uncorrelated dots would have interfered with stereoscopic surface detection by reducing the number of disparity samples available to the observer (a loss of effective sampling). Second, spurious matches, between either correlated and uncorrelated dots or pairs of uncorrelated dots, would have generated depth noise, which would have been inconsistent with the surface represented by the disparity signal. Both 'sampling' and 'depth noise' accounts predict that surface detection difficulties produced by image decorrelation should be mitigated by increasing the physical or effective density of RDS (compared to sparse-static displays). In principle, these improvements could have been produced because dense-static and sparse-dynamic displays: (i) provided additional disparity samples which aided in stereoscopic surface reconstruction and increased the effective signal-to-noise ratio; and (ii) generated depth noise which had either a smaller amplitude or appeared less stable.

According to the sampling account, tolerance to decorrelation improved as physical or effective density increased, because observers were able to extract larger numbers of disparity samples in dense-static and sparse-dynamic RDS (compared to sparse-static RDS). As the numbers of correct (and spurious) matches increased with the density, this eventually allowed the observer to adequately reconstruct the surface from decorrelated RDS. However, as can be seen by the outputs of our finite-sized window correlator model in Fig. 9, for most of the decorrelated stimulus conditions examined, each surface feature could still be represented by multiple disparity samples {in this case, the panels in the right column represent the outputs for a sparse (23 dots/deg^2), low spatial frequency (0.22 cpd) RDS with 80% decorrelation}. In these conditions, the model outputs suggest that the detection difficulties produced by image decorrelation were primarily due to the effects of disparity/depth noise interfering with surface reconstruction.

In principle, increasing the physical or effective density well above the Nyquist limit could actually have improved tolerance to decorrelation, because it increased the effective signal-to-noise ratio of the RDS. In the case of dynamic RDS, averaging disparity information over time would have acted to increase the signal-to-noise ratio, since any spurious dot matches occurring when viewing a dynamic RDS would be uncorrelated over time. Even in the case of static RDS, filtering precision should have improved with the number of available signal dots, leading to an increase in the effective signal-to-noise ratio. Consistent with this notion, Fig. 10 shows that for our window correlator model, identical disparity gratings were noticeably more visible in the presence of 80% image decorrelation with 676 dot/deg^2, as opposed to 23 dots/deg^2, RDS (especially for displays with the higher corrugation frequency). This suggests that noise tolerance improved as density increased because spatial filtering and pooling became more effective, which in turn, caused the sinusoidal signal to become more salient and more easily distinguished from the disparity noise.[7]

Finally, it is also possible that observers were more tolerant to image decorrelation in dense-static displays because they provided more noise dots than sparse-static displays with the same level of image decorrelation. Increasing the number of uncorrelated dots in a display would have increased the number of binocularly fused dots with stable perceived depths that were inconsistent with the smooth 3-D surface represented by the signal. However, as the population mean disparity of the depth noise was zero, the greater the number of noise dots, the more likely the observed local mean disparity of the depth noise would approximate to zero. As a result, the disparity of the depth noise might have been more likely to approximate zero in dense-static displays than in sparse-static displays. Consistent with the depth noise account, noise tolerance was always superior for dynamic RDS for two of our observers (SAP and MEL)—even when compared to the noise tolerance for the densest static RDS examined (676 dots/deg^2). However, MH's tolerance to image decorrelation with dense-static RDS sometimes rivaled that found for both sparse-dynamic and dense-dynamic RDS—suggesting that she was less

[6] We ran a series of control experiments to determine whether the pattern of results found for dot density, corrugation amplitude, spatial frequency and dot lifetime would persist when other aspects of the stimulus were manipulated. As average luminance varied with dot density in the main experiment, we ran a static RDS control experiment in which the average luminance was varied for each density. This control confirmed that the effect of density on decorrelated surface detection persisted across the range of average luminances tested (0.6–12 cd/m^2). We also replicated many of the effects in this paper with a larger RDS dot size (6 arcmin^2). Finally, we found that manipulations of density, amplitude, spatial frequency and dot lifetime produced roughly similar patterns of decorrelation tolerance for sinusoidal and square wave corrugations. However, marked differences in patterns of tolerance to decorrelation were found when detecting frontal plane surfaces.

[7] Support for this notion is also provided by the finding that the surface detection performance of an ideal observer became significantly more tolerant to additive disparity noise as the display density increased (see Fig. 11).

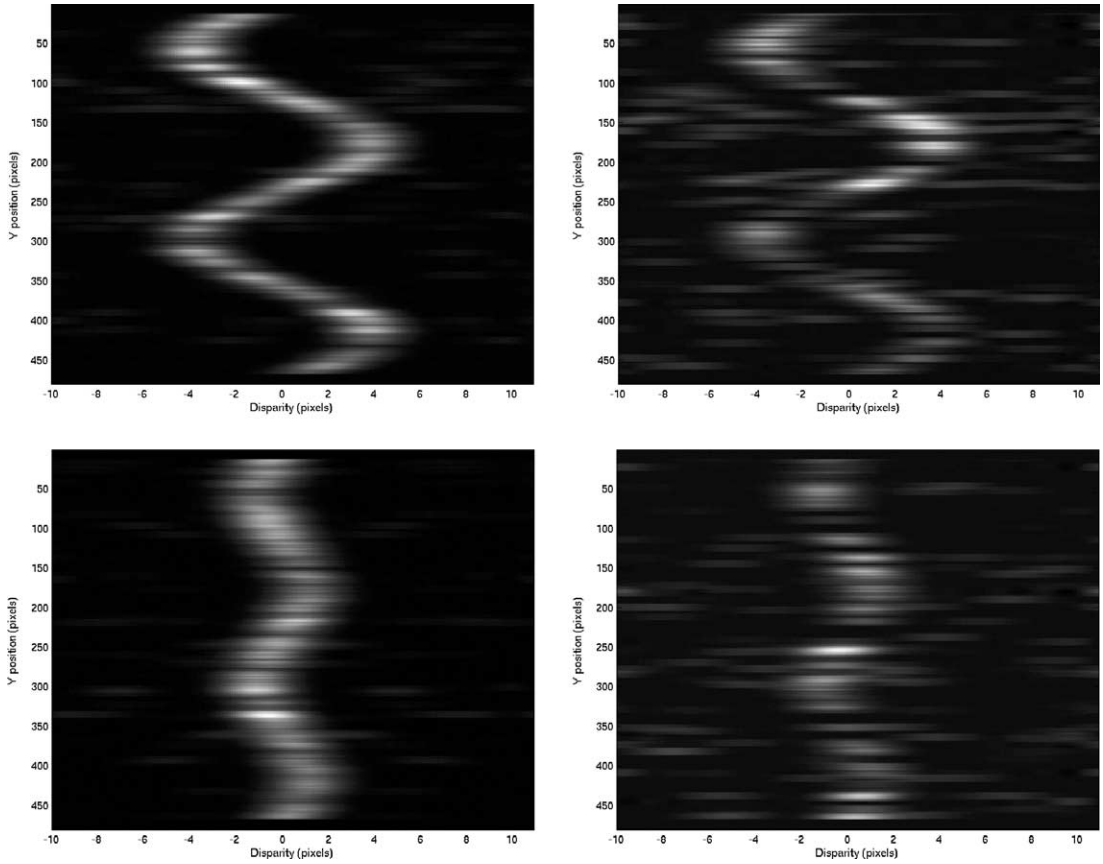

Fig. 9. Effect of 0% (left panels) or 80% (right panels) image decorrelation on the output of the same windowed correlator used in Fig. 5. All of the RDS conditions represented have a density of 23 dots/deg^2 and a spatial frequency of 0.22 cpd. The two top panels were produced by a disparity signal with a corrugation amplitude of 2 arcmin, whereas the bottom two panels were produced by a disparity signal with a corrugation amplitude of 0.5 arcmin.

susceptible to the stable depth noise in static displays relative to the other observers.

Importantly, manipulations of RDS amplitude and spatial frequency were predicted to produce slightly different results according to the 'sampling' and 'depth noise' explanations of image decorrelation. For example, it would be difficult to account for the finding in Experiment 2 that sparse static 0.5 arcmin RDS were more susceptible to image decorrelation than sparse larger amplitude displays, if one assumed that decorrelation only interfered with surface detection by reducing the number of effective samples below the Nyquist rate.[8] Based on this sampling account, the effects of decorrelation noise should have been quite similar for both the sparse 0.5 arcmin and the sparse 2 arcmin amplitude RDS—since the signal is clearly visible for both decorrelated amplitude conditions when modeled by our windowed correlator (see Fig. 9). Rather, we have argued that sparse static 0.5 arcmin RDS were more susceptible to the stable depth noise (produced by spurious dot matches) than sparse static 1–3 arcmin RDS displays, because the depth noise would have been larger with respect to the signal in the former case. This amplitude effect was found to disappear when display density increased—which could have been due to the resulting decrease in the average amplitude of the disparity noise or to an increase in the effective signal-to-noise ratio. Further support for this depth noise account of the static corrugation amplitude findings was also provided when we failed to find a similar amplitude effect for dynamic RDS in Experiment 3 (where spurious matches were less likely to result in stable depth noise).

In principle, increasing the spatial frequency of our sinusoidal displays from 0.22 to 0.88 cpd while holding dot density constant, would be expected to reduce the tolerance to image decorrelation, as each surface feature would be defined by progressively fewer disparity samples (eventually reducing effective sampling to the Nyquist rate). However, research has also shown that sensitivity to sinusoidal depth gratings peaks with corrugation frequencies between 0.2 and 0.4 cpd (e.g., Rogers & Graham, 1982). Thus, it was possible that at suprathreshold levels of disparity, the detection of surfaces with higher spatial frequency corrugations would be more susceptible to the effects of disparity noise. The notions that sampling

[8] Such an account might predict that static displays with larger corrugation amplitudes should be more (not less) susceptible to image decorrelation than displays with smaller amplitudes (e.g., if the visual system detects disparity by measuring correlation over a finite, fronto-parallel area—Banks et al., 2004).

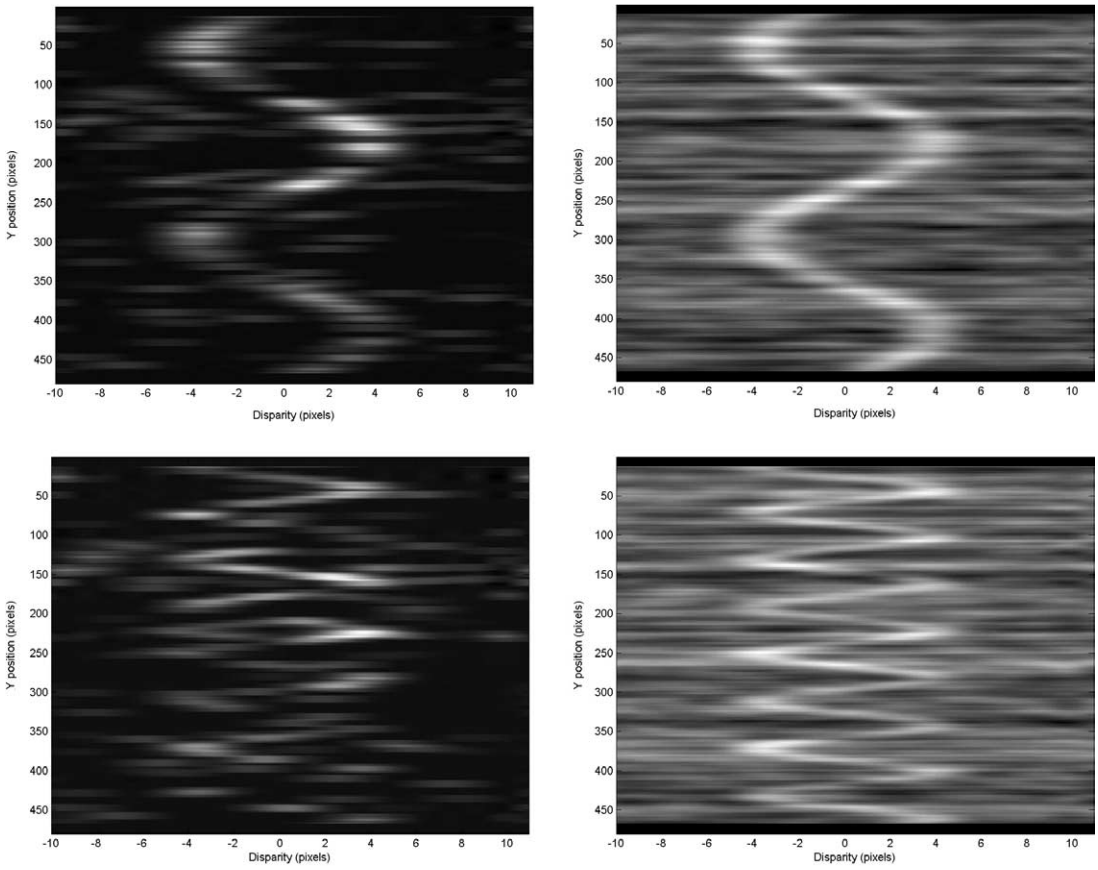

Fig. 10. Effect of dot density on the output of our windowed correlator when RDS had 80% image decorrelation. All displays had a corrugation amplitude of 2 arcmin. The top two panels represent the outputs for a 0.22 cpd display, whereas the bottom two outputs represent the outputs for a 0.88 cpd display. The output on the left was produced for a sparse 23 dots/deg^2 RDS, whereas the output on the right was produced for a dense 676 dots/deg^2 RDS.

issues and/or spatial frequency sensitivity differences were responsible for the spatial frequency effects observed in the current paper were both consistent with the findings of a recent modeling study by Palmisano et al. (2000). This earlier study found that human sinusoid detection in the presence of Gaussian distributed additive disparity noise was spatial frequency dependent. However, a template-matching ideal observer failed to demonstrate any spatial frequency dependency when presented with the same stimuli (see Fig. 11). In principle, this null find-

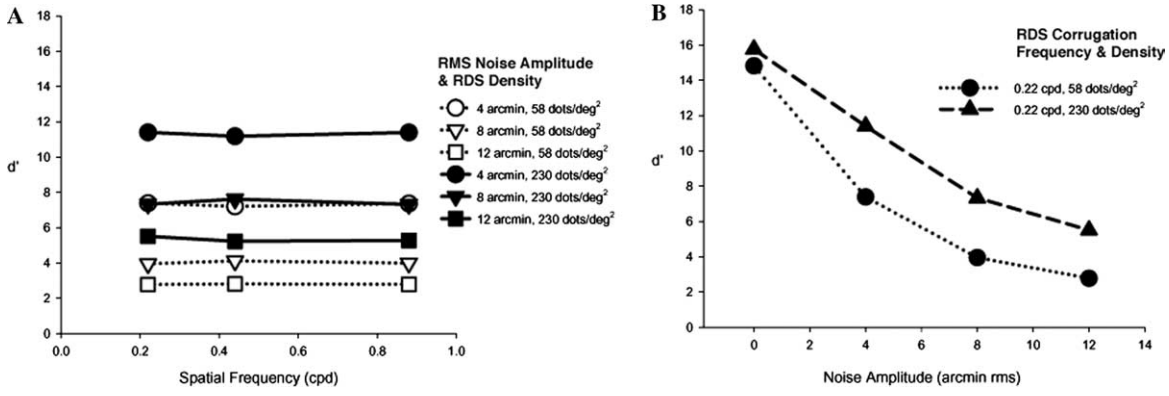

Fig. 11. Effects of RDS density and spatial frequency on an ideal observer's ability to detect noisy sinusoidal depth gratings (an algorithm used by Palmisano, Allison and Howard, 2000). All of the dots in these static RDS were correlated (represented a 2′ amplitude grating) and contained Gaussian-distributed additive disparity noise (similar to the impulse depth noise produced by image decorrelation in our RDS). The ideal observer was able to perform the matching task perfectly. After extracting the ideal disparity map, it then compared its output to three matched filters (one corresponding to each signal spatial frequency) to make its decision about whether a sinusoidal surface was present or not. (A) Ideal observer detection performance (d') as a function of RDS density (58 or 230 dots/deg^2), corrugation spatial frequency (0.22, 0.44 or 0.88 cpd) and the RMS amplitude of the additive disparity noise (4, 8 or 12 arcmin). (B) The effects of RMS noise amplitude and display density on ideal observer detection performance in finer detail for 0.22 cpd RDS.

ing might indicate that the ideal observer used was not an appropriate model for human stereopsis—since it failed to demonstrate spatial frequency selectivity. However, since the ideal observer was able to match all of the available signal dots perfectly, this null finding could also be taken as support for the notion that the reduced human noise tolerance for 0.88 cpd RDS was due to sampling issues. The latter interpretation was partially supported by the output of our window correlator model in Fig. 10—where the signal for a sparse 0.88 cpd, 80% decorrelated RDS appeared to be poorly represented. In either case, the disparity noise produced by spurious matches should have been more disruptive for 0.88 cpd RDS (irrespective of whether the underlying cause of the spatial frequency effect was due to differences in sensitivity or sampling issues).

Several sources of evidence suggest that the benefits obtained by increasing the effective/physical density of a RDS were strictly limited: (i) while reducing dot lifetimes from 1.6 s to 80 ms substantially improved tolerance to image decorrelation for sparse displays, further reductions had little effect on noise tolerance (Experiment 3); and (ii) while detection performance with dense-static displays was always more tolerant to image decorrelation than that with sparse-static displays, detection performance with sparse-dynamic and dense-dynamic displays was very similar (in some cases performance was actually superior with sparse-dynamic displays). It appears that the above manipulations of physical and effective density were sufficient to bring detection performance and noise tolerance to near ceiling levels.

In conclusion, the present experiments have shown that stereoscopic surface detection can tolerate substantial image decorrelation. However, it appeared that greater numbers of correlated dots were required to detect a 3-D surface than were required to detect the presence of interocular correlation (as reported by Cormack et al., 1997). Detection performance was found to tolerate greater levels of image decorrelation as either the physical density of the RDS increased or its dot lifetimes decreased, because both manipulations rendered the observer more resistant to consequences of spurious matches (they prevented effective undersampling, increased the effective signal-to-noise ratio and reduced the impact of depth noise on surface reconstruction). The remarkable noise tolerance observed in this study provides further evidence that: (i) the visual system can match dots in the two eyes' images in a highly proficient manner; and (ii) stereoscopic surface detection can often tolerate substantial disparity noise when errors in binocular matching occur.

Acknowledgments

The authors thank Toni Howard and two anonymous reviewers for their helpful feedback and suggestions. We would also like to thank the following naive observers who ran numerous experimental sessions in studies related to this paper: Xueping Fang, Desley Hennessy, Melinda Hinton, Michelle Hudoba and Heather Jenkin. This work was partly supported by Defence R&D Canada, Contract W7711-7-7393. The support of NSERC (Canada) to IPH and RA is also gratefully acknowledged. Correspondence should be addressed to Stephen Palmisano, Department of Psychology, University of Wollongong, Wollongong, NSW 2522, Australia.

References

Allison, R. S., & Howard, I. P. (2000). Stereopsis with persisting and dynamic textures. *Vision Research, 40*, 3823–3827.

Banks, M. S., Gepshtein, S., & Landy, M. S. (2004). Why is spatial stereoresolution so low? *The Journal of Neuroscience, 24*(9), 2077–2089.

Bradshaw, M. F., & Rogers, B. J. (1999). Sensitivity to horizontal and vertical corrugations defined by binocular disparity. *Vision Research, 39*, 3049–3056.

Bradshaw, M. F., Rogers, B. J., & De Brun, B. (1995). Perceptual latency and complex random-dot stereograms. *Perception, 24*, 749–759.

Christophers, R. A., Rogers, B. J., & Bradshaw, M. F. (1993). Perceptual latencies, vergence eye-movements and random-dot stereograms. *Investigative Ophthalmology & Visual Science*, 1438 (ARVO Abstracts).

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, 14*, 723–730.

Cormack, L. K., Stevenson, S. B., & Schor, C. M. (1991). Interocular correlation, luminance contrast and cyclopean processing. *Vision Research, 31*, 2195–2207.

Cormack, L. K., Stevenson, S. B., & Schor, C. M. (1994). An upper limit to the binocular combination of stimuli. *Vision Research, 34*, 2599–2608.

Harris, J. M., & Parker, A. J. (1994). Constraints on human stereo dot matching. *Vision Research, 34*, 2761–2772.

Howard, I. P., & Rogers, B. J. (1995). *Binocular vision and stereopsis*. New York: Oxford University Press.

Julesz, B. (1960). Binocular depth perception of computer generated patterns. *Bell System Technical Journal, 39*, 1125–1162.

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

Julesz, B. (1971). *Foundations of cyclopean perception*. Chicago, Ill: University of Chicago Press.

Lankheet, M. J. M., & Lennie, P. (1996). Spatio-temporal requirements for binocular correlation in stereopsis. *Vision Research, 36*, 527–538.

Livingstone, M. S. (1996). Differences between stereopsis, interocular correlation and binocularity. *Vision Research, 36*, 1127–1140.

MacMillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide*. New York, NY: Cambridge University Press.

Palmisano, S., Allison, R. S., & Howard, I. P. (2000). Effect of disparity noise on stereoscopic surface perception in humans and ideal observers. *Proceedings of ICSCC on Intelligent Systems & Applications, 1*(2), 1006–1012.

Palmisano, S., Allison, R. S., & Howard, I. P. (2001). Effects of horizontal and vertical additive disparity noise on stereoscopic corrugation detection. *Vision Research, 41*, 3133–3143.

Rogers, B. J., & Graham, M. (1982). Similarities between motion parallax and stereopsis in human depth perception. *Vision Research, 22*, 261–270.

Shannon, C. E. (1949). Communication in the presence of noise. *Proceedings of Institute of Radio Engineers, 37*, 10–21.