Pergamon

PII: S0042-6989(97)00132-6

Vision Res., Vol. 38, No. 3, pp. 387–401, 1998

Printed in Great Britain
0042-6989/98 $19.00 + 0.00

Luminance Contrast Affects Motion Coherency in Plaid Patterns by Acting as a Depth-from-occlusion Cue

GENE R. STONER,† THOMAS D. ALBRIGHT*

Received 13 December 1996; in revised form 2 May 1997

Moving plaid patterns composed of component gratings that differ in luminance contrast tend not to cohere perceptually. Plaid patterns configured to mimic one occlusive grating overlying another also fail to cohere. We hypothesized that plaids constructed of components with different luminance contrasts fail to cohere because these components are interpreted as occlusive surfaces lying in different depth planes. It is known that when depth-from-occlusion and depth-from-binocular disparity cues support the same depth-ordering, both segregation in depth and motion non-coherency are more likely to be perceived than when these two cues conflict. We exploited this interaction and tested our hypothesis by introducing horizontal binocular disparity between two superimposed component gratings of different luminance contrasts. We found that both depth segregation and motion non-coherency were much more likely when the high-contrast grating was stereoscopically in front of the low-contrast grating. From these results we infer that luminance contrast acts as a depth-cue in plaid patterns, with higher contrast gratings appearing to lie in front of lower contrast gratings. Perceptual motion coherency parallels these depth-ordering judgments. We conclude that luminance contrast affects motion coherency by acting as a depth-from-occlusion cue.

Motion coherency Plaid Luminance contrast Depth

INTRODUCTION

A central problem facing the visual system is the decomposition of complex dynamic images into component motions. The computational strategies used to solve this problem can be investigated in the laboratory using visual stimuli formed by superimposition of two one-dimensional gratings (Fig. 1). When these "plaid patterns" are moved within a stationary aperture, the two component gratings can be seen to move together ("coherent motion") or independently ("non-coherent motion") (Adelson & Movshon, 1982). Although experiments employing plaid patterns have shed light on the stimulus parameters that govern this selective motion decomposition, the nature of the underlying neural mechanism is still controversial (see Stoner & Albright, 1994; Albright & Stoner, 1995, for reviews). At the center of this debate lies the question of whether motion processing is determined exclusively by "image-based" parameters (i.e., primary stimulus dimensions such as luminance contrast and contour orientation) or whether "scene-based" interpretation (i.e., the interpretation of image features in terms of object properties such as surface reflectance and depth) also plays a critical role. Image-based and scene-based explanations have been associated historically with two different classes of experimental manipulations employing plaid patterns. Image-based mechanisms have been invoked typically to explain the observation that components tend to cohere less when they are dissimilar along stimulus dimensions such as spatial frequency, luminance contrast, and binocular disparity (Adelson & Movshon, 1982, 1984; Movshon *et al.*, 1985; Krauskopf & Farell, 1990; Kooi *et al.*, 1992; Dobkins *et al.*, 1992; Kim & Wilson, 1993). The dominant mechanistic account that has been advanced to explain these effects of component similarity is a simple feed-forward, labeled-line model, in which the motions of dissimilar gratings are processed by separate motion channels (Adelson & Movshon, 1982; Movshon *et al.*, 1985).

Whereas image-based explanations are generally restricted to the problem of extracting the motions of retinal image features, scene-based accounts address the larger problem of interpreting these elemental motions in terms of the motion of objects in the visual scene. The distinction between the motions of retinal image features and environmental objects is a critical one. Members of

*Vision Center Laboratory, The Salk Institute for Biological Studies, La Jolla, CA 92037, U.S.A.
†To whom all correspondence should be addressed at: The Salk Institute, P.O. Box 85800, San Diego, CA 92138, U.S.A. [*Tel*: 619-453-4100 x1483; *Fax*: 619-546-8526; *Email*: gene@salk.edu].

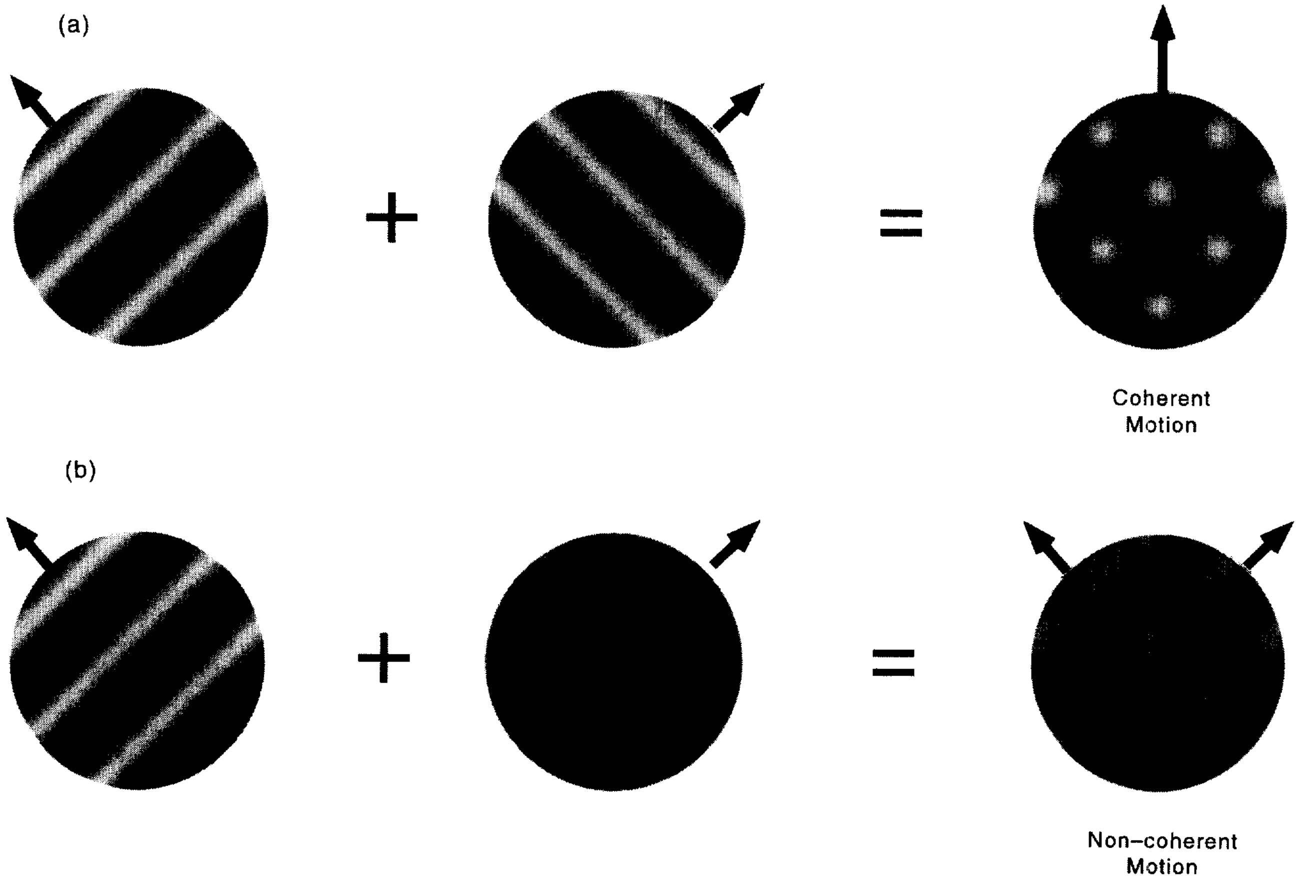

FIGURE 1. Schematic depiction of plaid pattern construction from gratings that are of the same or different luminance contrast. (a) Linear addition of two sine-wave gratings of equal contrast yields a plaid pattern that is nearly always perceived to move coherently. (b) Linear addition of two sine-wave gratings of sufficiently different contrasts yields a plaid pattern that is nearly always perceived to move non-coherently.

one class of retinal image features ("intrinsic" features) have a direct correspondence with the features of environmental objects. A second class of features ("extrinsic" features) have no such correspondence, being instead a consequence of the incidental overlap of objects in the formation of the retinal image (Shimojo *et al.*, 1989). The movements of intrinsic and extrinsic retinal image features must be distinguished if a veridical representation of object motion is to be achieved. Scene-based accounts of motion processing stress the role of depth-from-occlusion cues (Fig. 2) in making this distinction. Evidence implicating scene-based mechanisms includes the observation that motion non-coherency is more likely when plaid patterns are configured using such cues to resemble two overlapping surfaces (Stoner *et al.*, 1990; Stoner & Albright, 1996; Trueswell & Hayhoe, 1993; Vallortigara & Bressan, 1991; Stoner & Albright, 1996).

It is conceivable (though not parsimonious) that scene-based mechanisms need only be called upon when depth-ordering cues are explicitly introduced, but that image-based mechanisms suffice to explain the effects of component similarity. Alternatively, the influence of component similarity on motion coherency may reflect the inadvertent triggering of mechanisms sensitive to depth-from-occlusion cues (Stoner & Albright, 1993, 1994). Here we provide evidence in support of the latter position as it applies to luminance contrast. This evidence points to a conceptual framework that unifies the seemingly diverse effects of component similarity and depth-ordering cues.

Luminance contrast and occlusion

A potential link between one dimension of component similarity (luminance contrast) and image segmentation is illustrated in Fig. 2. Figure 2(a) contains two overlapping rectangles in which sub-regions labeled A–D define the elemental unit of opaque occlusion—the "T-junction". These terminated lines of luminance contrast are extrinsic retinal image features that result from occlusion of one opaque surface by another. Not surprisingly, T-junctions generally give rise to a percept of opaque occlusion. It is obvious—but notable for arguments that follow—that the foreground surface (B) can be uniquely identified by the fact that it possesses the least luminance contrast (zero, in this case) with the region of overlap (A).

Figure 2(b) illustrates the elemental unit of transparent occlusion—the "X-junction". This extrinsic retinal image feature arises when two surfaces overlap and the

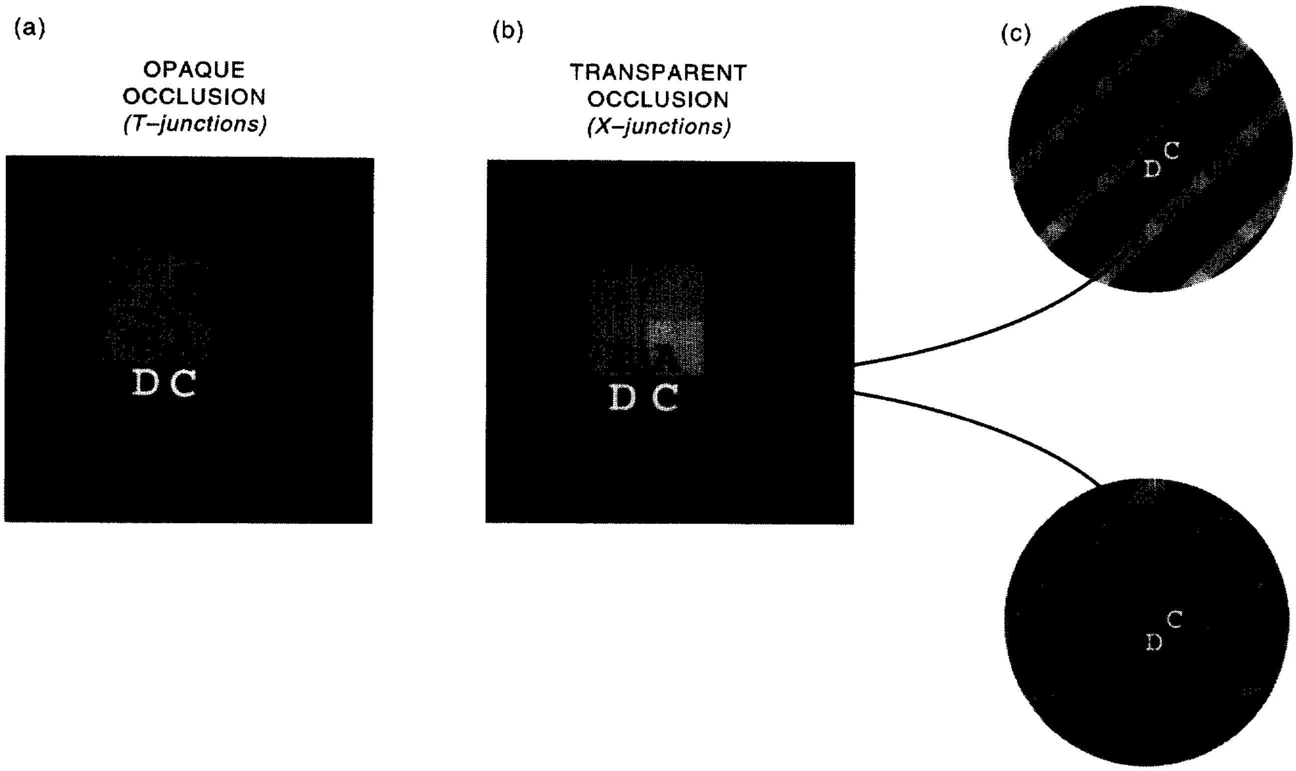

FIGURE 2. Luminance conditions associated with opaque and transparent occlusion yield characteristic image features known as T- and X-junctions. These two basic types of occlusion can be simulated using simple physical rules for reflection and transmittance of light by surfaces (Matelli, 1974; Beck *et al.*, 1984; Stoner & Albright, 1996). Appropriate luminance ratios convey a sense of depth-ordering in a pattern devoid of other depth cues. (a) In this "cascading rectangles" configuration, the rectangle designated B can be interpreted as an opaque foreground surface partially occluding rectangle C. Both rectangles can be seen to lie in front of a larger rectangle D. Luminance conditions that simulate opaque occlusion are such that rectangle B reflects light but transmits no light from rectangle C. Because foreground transmittance is nil, the luminance of the region of surface overlap (A) is simply equal to the luminance of the foreground rectangle (B). These luminance relationships define an image feature known as a "T-junction," which is a ubiquitous cue for opaque occlusion. A T-junction thus implies opaque occlusion and the depth-ordering of the overlapping surfaces is revealed by luminance identity with the point of overlap (foreground surface luminance always equals the overlap luminance). (b) Luminance conditions that simulate transparent occlusion are such that rectangle B reflects light and transmits some light from rectangle C. Because foreground transmittance is non-zero, the luminance of the region of surface overlap (A) is equal to the luminance of the foreground rectangle (B) plus the product of foreground transmittance and the luminance of rectangle C. These luminance relationships define an image feature known as a "X-junction," which is a ubiquitous cue for transparent occlusion. An X-junction thus implies transparent occlusion. Surface depth-ordering can be determined by a generalization of the luminance identity rule for opaque occlusion: the foreground surface (B) can be identified by the fact that it possesses the least luminance contrast with the point of surface overlap (A). (c) The rule for perceptual depth-ordering of surfaces in the presence of X-junctions implies that depth-ordering should occur in rectangular-wave (top) and sine-wave (bottom) plaids constructed from gratings of different luminance contrasts.

foreground surface has some transmittance. The principle underlying depth-ordering for X-junctions can be considered a simple extension of that underlying depth-ordering for T-junctions (Trueswell & Hayhoe, 1993): the surface having the least contrast (B) with the region of overlap (A) is usually seen as part of the foreground occluder. Depth-from-transparent-occlusion emerges, Trueswell and Hayhoe suggest, from a mechanism whose primary function is to detect depth on the basis of opaque occlusion.

The particular intensities of the X-junction shown in the middle panel of Fig. 2 were chosen to match those of a plaid pattern composed of two component gratings of different contrasts [Fig. 2(c)]. This geometrical correspondence is an incidental and previously unappreciated consequence of additive superimposition and constitutes the basis of our prediction that occlusion-based rules for depth-ordering account for the well-documented effects of dissimilar component contrasts on perceptual motion coherence (Adelson & Movshon, 1982). If plaids created by additive superimposition have components of different contrast, the perceived region of overlap (A) will always possess less contrast with respect to the high-contrast grating (B) than with respect to the low-contrast grating (C). From the logic outlined above, it follows that the high-contrast grating should appear to occlude the low-contrast grating. Using plaid stimuli of this sort we sought to determine: (1) whether a particular depth-ordering based on an occlusion interpretation is promoted in plaid patterns with component gratings of dissimilar

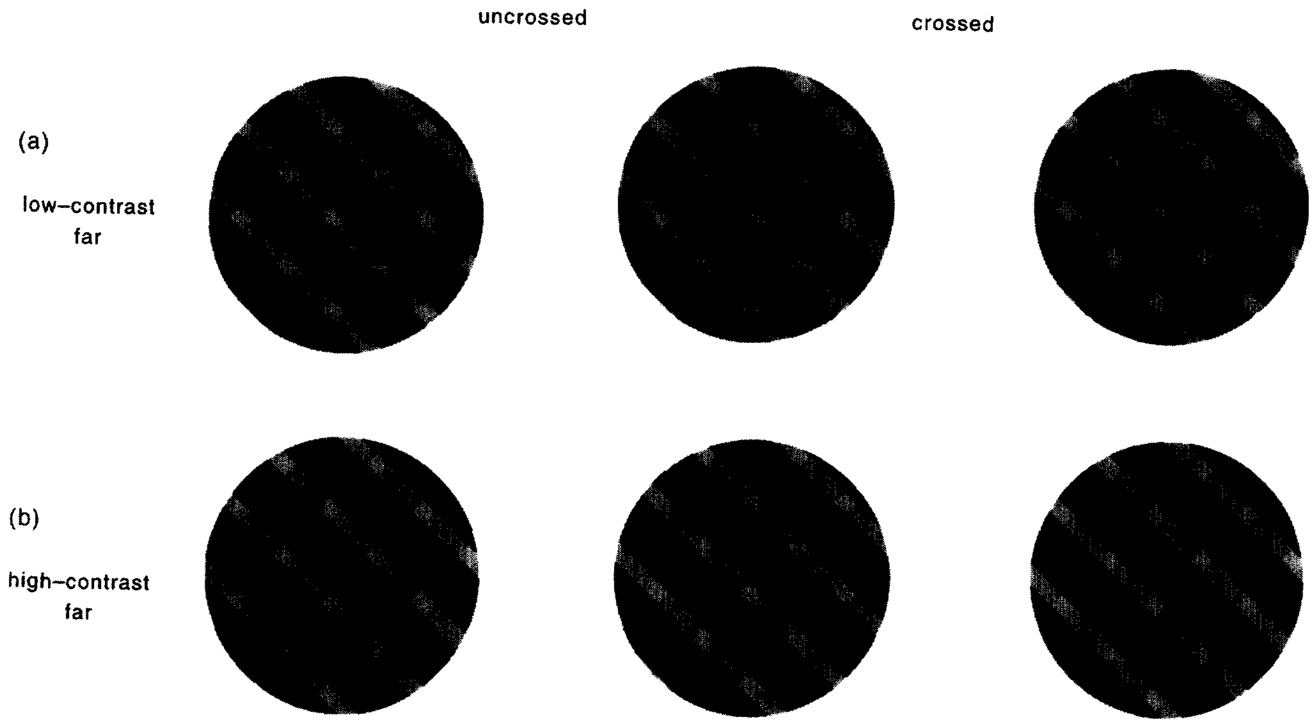

FIGURE 3. Stereoscopic illustration of rectangular-wave plaids used in Experiments 1 and 2. Two different binocular disparity sign conditions are shown, each at a single disparity magnitude. Stereo image pairs are presented for both crossed and uncrossed viewing. (a) Low-contrast–far condition. In this case the sign of binocular disparity is such that the low-contrast grating in the plaid is placed stereoscopically behind the high-contrast grating. The occlusion-based depth-ordering hypothesis argues that disparity and occlusion cues are consistent for this condition, which should lead to unambiguous depth-ordering. (b) High-contrast far condition. In this case the sign of binocular disparity is such that the high-contrast grating is placed stereoscopically behind the low-contrast grating. The occlusion-based depth-ordering hypothesis argues that disparity and occlusion cues are inconsistent for this condition, which should disrupt perceptual depth-ordering.

contrasts; and (2) whether this depth segregation, if significant, underlies the non-coherence of such plaids. To do so we conducted three experiments employing a simple and direct strategy: we evaluated the contribution of the hypothesized depth-from-occlusion cue by pitting it against the known depth-cue of horizontal binocular disparity.

Experiment 1 confirmed that high-contrast gratings generally appear to lie in front of low-contrast gratings. Experiment 2 confirmed that this depth-ordering is highly correlated with perceptual motion coherence. Experiment 3 verified that the visibility of X-junctions is critical for both the depth-segregation and motion coherency effects. Taken together, these results demonstrate that luminance contrast affects motion coherency by acting as a depth-from-occlusion cue.

GENERAL METHOD

Subjects

Three human subjects participated in these experiments. Two subjects were completely naïve with regard to the goals of these experiments; the third subject was one of the authors (GS). All subjects had normal or corrected-to-normal acuity and normal stereopsis.

Apparatus

All stimuli were generated using a high-resolution graphics controller (Pepper SGT, Number Nine Computer Corporation: 640 × 480 pixels, analog RGB output, 8 bits/gun) operating in a microcomputer. Left and right stereo images were displayed on a 14′ analog RGB video monitor (Zenith ZCM-1490, flat technology CRT, 60 Hz frame rate, non-interlaced). Photometric linearization tables were computed and used to reform the non-linear voltage-luminance relationship characteristic of the monitor. Stereo pairs were viewed through a mirror haploscope from a distance of 228 cm. Pixels were square with an angular extent of 1.10 min of arc, which defines the maximal spatial resolution of luminance modulated patterns, as well as the grain of all binocular disparity manipulations. The ambient light level in the experimental room was approx. 0.1 cd/m^2 and the mean luminance of the screen during the inter-trial interval was 25.0 cd/m^2. A chin and forehead rest were used for head stabilization.

Visual stimuli

Stereoscopic plaid patterns similar to those depicted in Fig. 3 were used in Experiments 1, 2 and 3. Both rectangular- and sine-wave gratings (upper and lower right panels of Fig. 2, respectively) were used in Experiments 1 and 2. Only rectangular-wave gratings were used in Experiment 2. The duty cycle (narrow bar width/[narrow bar width + wide bar width]) of the rectangular-wave gratings was 0.25 for Experiments 1 and 2. These asymmetrical duty cycles were used to bias the interpretation of foreground/background. Additional stimulus constraints led to the use of a duty cycle of 0.1875 in Experiment 3. These issues are discussed in more detail in the specific Methods sections that follow. Sine-wave gratings were phase symmetric. Plaids were formed by simple additive superimposition of two components differing in orientation by 90 deg (±45 deg relative to vertical). The two component gratings also differed in luminance contrast: one grating was always of high luminance contrast (0.8 Michelson) and the other was always of low contrast (0.2). The plaid pattern viewed by each eye appeared within a circular software aperture with a diameter of 2.35 deg. The positions of these apertures were fixed and centered on the points of ocular fixation. Plaids used in Experiment 1 were static during each presentation; plaids used in Experiment 2 moved within the viewing aperture. Both static and moving plaids were used in Experiment 3.

There were two principal independent variables in each experiment. The first was the sign of horizontal binocular disparity between the two gratings of different contrast: either the low- or the high-contrast grating was placed in the background. We refer to these as the "low-contrast-far" and "high-contrast-far" conditions, respectively. The second independent variable was the magnitude of horizontal binocular disparity between component gratings. The foreground grating was always placed at 0 disparity (i.e., in the fixation plane) and the background grating varied over a range that depended upon stimulus conditions (see below).

For Experiments 1 and 2, a third independent variable was the type of grating (rectangular- vs sine-wave). Each stimulus condition in Experiments 1 and 2 was thus a conjunction of disparity sign, disparity magnitude, and grating type. Experiment 3 involved additional manipulations, described below.

Procedure

The dependent variable in Experiment 1 was the perceived depth segregation of the component gratings: subjects were asked to report whether the gratings were seen to lie in the same plane or in different planes. The dependent variable in Experiment 2 was perceived motion coherence: subjects were asked to report whether the gratings were seen to move independently (non-coherence) or as a single pattern (coherence). In Experiment 3, subjects were asked to make both depth segregation and motion coherence judgments on separate sets of trials.

Excepting the type of perceptual report required, procedures for Experiments 1–3 were identical. Data were collected with a two-alternative forced-choice procedure, using the method of constant stimuli. Subjects were asked to fixate a red square (4.40 by 4.40 min of arc) at the center of the display before and during each stimulus presentation. Human subjects are capable of reliable fixation under these conditions (Murphy *et al.*, 1975). Stimulus presentation was initiated by a keyboard press by the subject once fixation was achieved. A stereoscopic plaid pattern then appeared, constructed with a specific combination of the independent variables described above. Low-contrast–far and high-contrast–far conditions were randomly interleaved within blocks. Sine-wave and rectangular-wave plaid conditions were presented in separate blocks. In addition, to ensure that biases based on orientation were not confounded with luminance contrast manipulations, the orientations of the low- and high-contrast gratings were reversed for half of the blocks for all experiments. Magnitude of binocular disparity was randomized from trial to trial within each block. Following stimulus presentation, subjects indicated perceptual judgment with an appropriate key-press. For subjects RD and GS, the data reported for all three experiments are based on six blocks of trials with five trials per condition per block, yielding a total of 30 trials/condition. Subject EC's data are based on 10 blocks of 10 trials per condition per block, yielding a total of 100 trials/condition. Prior to collection of these data, each subject was initially presented with a series of practice trials, which were continued until performance became stable and the subject expressed a clear understanding of the requirements of the task.

EXPERIMENT 1: EFFECT OF LUMINANCE CONTRAST ON DEPTH-ORDERING IN PLAID PATTERNS

The general goal of these experiments was to determine whether the motion non-coherency previously observed for plaid patterns composed of gratings of different contrasts arose as a result of depth-ordering based on the contrast cues associated with X-junctions. The specific goal of Experiment 1 was to assay the potency of contrast-based depth-ordering by examining its interactions with depth-from-binocular disparity. We also wished to compare depth-ordering elicited using rectangular-wave vs sine-wave plaids. Whereas previous studies of the effects of component contrast similarity employed sine-wave plaids exclusively (Adelson & Movshon, 1982), our use of rectangular-wave plaids in the present study forms a critical link with earlier work on occlusion-based depth-ordering in plaid patterns, which has heretofore only been demonstrated for rectangular-wave patterns (e.g. Stoner *et al.*, 1990; Stoner & Albright, 1992, 1996; Trueswell & Hayhoe, 1993). Employing rectangular-wave patterns, Stoner & Albright (1996) demonstrated that whether an X-junction is interpreted as two overlapping surfaces depends critically on the perceptual assignment of foreground/background to the sub-regions of that X-junction. Preliminary questioning of several naïve and non-naïve subjects established that sine-wave plaids composed of gratings of different contrast were spontaneously parceled into foreground and background, as indicated in Fig. 2 (i.e., region D as the background, regions B and C as overlapping surfaces, and region A as their region of overlap). Asymmetric duty cycles were used to construct rectangular-wave plaids, as we have previously found the resulting areal size differences to render reliable perceptual assignment of foreground/background and depth (Stoner & Albright, 1996). The rectangular-wave duty cycles used strongly biased the foreground/background assignment to conform with that indicated in Fig. 2. Agreement between experiments employing rectangular-wave and sine-wave plaid patterns would support our assertion that surface segmentation mechanisms interpret the two types of stimuli similarly.

Methods

Visual stimuli and psychophysical procedure. Subjects viewed stereoscopic plaid stimuli like those depicted in Fig. 3. The mean luminances for rectangular-wave and sine-wave plaids were 13.9 and 24.7 cd/m^2, respectively. Disparity between components of rectangular-wave plaids ranged from 0 to 5.5 min arc at intervals of 1.1 min arc. Disparity between components of sine-wave plaids ranged from 0 to 16.5 min arc at intervals of 3.3 min arc. Based on data acquired during subject training, stimulus durations were chosen to insure that, for a particular stimulus, each subject reported a reasonable balance of coherent and non-coherent motion. For rectangular-wave conditions, viewing duration during each trial was 2.5 sec for subjects RD and GS and 1.5 sec for subject EC. Viewing duration of sine-wave conditions was 3 sec for all subjects. Subjects were instructed to report whether the component gratings appeared to be separated in depth or whether they appeared to lie in the same depth plane.

Results

Previous experiments documented an interaction between depth-ordering elicited by X-junctions and binocular disparity (Trueswell & Hayhoe, 1993). On the basis of these findings, we anticipated that perceptual depth segregation would be most robust when occlusion and disparity cues were in agreement. Because our hypothesis argues that the high-contrast grating resembles an occluder overlying the low-contrast grating, depth-cue agreement should be associated with the low-contrast–far condition. In terms of our independent variables, we expected an interaction between the effects of binocular disparity magnitude and sign, such that perceptual depth segregation would be perceived more often for the low-contrast–far condition. This prediction is borne out by our results.

Results obtained using static rectangular-wave plaids are shown in Fig. 4(a). Plots for each of the three subjects indicate the proportion of trials on which the component gratings were judged to lie in separate depth planes. This proportion is plotted separately for blocks of trials in which the low-contrast–far (filled triangles) and high-contrast–far (filled circles) stimuli were viewed. For low-contrast–far plaids, the proportion of trials on which depth segregation was reported increased as a function of the magnitude of binocular disparity between the two gratings. For high-contrast–far stimuli, however, the probability of reporting depth segregation either failed to increase with increasing binocular disparity, or increased at a much lower rate than seen for low-contrast–far plaids. For any given non-zero disparity, subjects were much more likely to report segregation in depth for the low-contrast–far plaids than for the high-contrast–far plaids.

A qualitatively similar pattern of results was obtained for sine-wave plaids, plotted in Fig. 4(b) (note that the x-axes of lower and upper panel of plots in Fig. 4 are of different scales). Once again, frequency of depth segregation reports increased rapidly with increasing disparity for the low-contrast–far condition. By comparison, segregation reports increased slowly for the high-contrast–far condition. As was true for the rectangular-wave plaids, for all non-zero disparities, all subjects were much more likely to report segregation in depth for the low-contrast–far than for the high-contrast–far conditions.

The principal difference between results obtained using rectangular-wave vs sine-wave plaids is the relatively greater disparity required to see the components of sine-wave plaids segregated in depth. There are at least three possible explanations for this difference. First, it could be due in part to the relative absence in sine-wave plaids of high spatial frequencies, which may be important for establishing binocular correspondence. Second, it is possible that sine-wave components are less likely to be interpreted as "surfaces" owing to their diffuse appearance. A third possibility is that the different results are attributable not to the sine- vs rectangular-wave difference, but to the component duty cycle difference that covaried with grating type (sine-wave = 0.5; rectangular-wave = 0.25). According to this hypothesis, sine-wave plaids may have yielded a less stable impression of depth-from-occlusion owing to the absence of size-cues for foreground/background interpretation. This interpretation is supported by the fact that the perceived motion of plaids constructed from 0.5 duty cycle rectangular-waves is more likely to be coherent than for plaids made of asymmetrical rectangular-waves (Stoner & Albright, 1996). In any case, the interaction between binocular disparity sign and magnitude on perceptual depth segregation can be taken as evidence for an implicit depth-from-occlusion cue that is present in both rectangular- and sine-wave plaids, such that high-contrast gratings tend to be seen in front of low-contrast gratings.

EXPERIMENT 2: EFFECT OF LUMINANCE CONTRAST ON MOTION NON-COHERENCY IN PLAID PATTERNS

The results of Experiment 1 suggest that the contrast-related principles of depth-from-occlusion, identified by Trueswell & Hayhoe (1993), apply to the simple case of plaids constructed from additive superimposition of

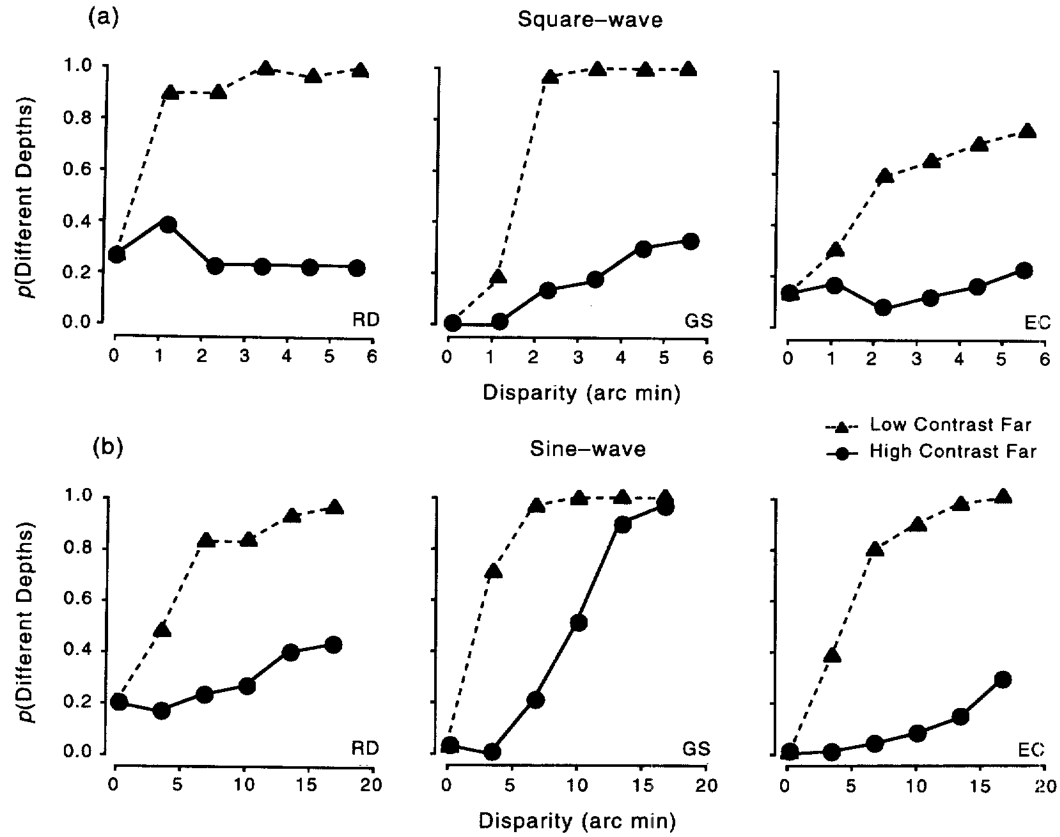

FIGURE 4. In Experiment 1, the influence of contrast-based occlusion on perceptual depth-ordering was evaluated by calibrating it against binocular disparity cues. Subjects viewed a static plaid on each trial (see Figs 1–3 for stimulus description and design principles). Two independent variables were manipulated: (i) sign of binocular disparity, i.e., low-contrast grating stereoscopically in front of high-contrast grating, or vice versa; (ii) magnitude of binocular disparity. All other stimulus properties, except grating type (rectangular- vs sine-wave; see below), were constant. Following each brief stimulus presentation, subjects were required to report whether the component gratings appeared to lie in same or different depth planes. Data are shown for each of three subjects (columns) who viewed each of two grating types (rows). (a) Proportion of trials on which subjects reported depth segregation of rectangular-wave gratings is plotted as a function of binocular disparity magnitude for each disparity sign (filled triangles = low-contrast–far; filled circles = high-contrast–far). For low-contrast–far stimuli, probability of depth segregation reports increased dramatically with increasing disparity magnitude. By contrast, depth segregation was rarely reported for high-contrast–far stimuli, regardless of disparity magnitude. The potency of contrast-based occlusion cues for depth-ordering is thus revealed by interaction with disparity cues. (b) The same conditions were repeated using sine-wave gratings. Perceived depth-ordering exhibited a similar dependence upon the conjunction of disparity sign and magnitude. For subjects RD and GS each data point is based on a total 30 trials. For subject EC, each data point is based on 100 trials. See text for details.

components of differing contrast. We hypothesized that this apparent depth-ordering leads to the motion non-coherence associated with such stimuli (Adelson & Movshon, 1982). To test this hypothesis, we conducted Experiment 2, the specific goal of which was to evaluate the contribution of contrast-based depth-ordering to motion coherence by examining the interaction of the contrast cue with a binocular disparity cue.

Method

Visual stimuli and psychophysical procedure. The plaid stimuli used for Experiment 2 were constructed by the same rules and had the same geometry and luminance values as those used in Experiment 2. The only difference was that the plaids used for Experiment 2 moved within the viewing aperture. Both component gratings moved at 1.17 deg/sec, which yielded a pattern speed of 1.65 deg/sec. Plaid patterns moved either up or down on a random schedule. Independent variables and trial blocking procedures were also identical to those for Experiment 1, as were viewing durations. Subjects were instructed to report whether the component gratings appeared to drift across one another in separate directions (motion non-coherency) or to move together in a common direction (motion coherence).

Results. In view of the results obtained in Experiment 1, and adopting the assumption that perceptual motion non-coherence should parallel perceptual segregation in depth, we expected to see an interaction between the effects of binocular disparity magnitude and sign. Specifically, we predicted that reports of perceptual motion non-coherence would be greatest in the low-contrast–far condition. This prediction is supported by our results.

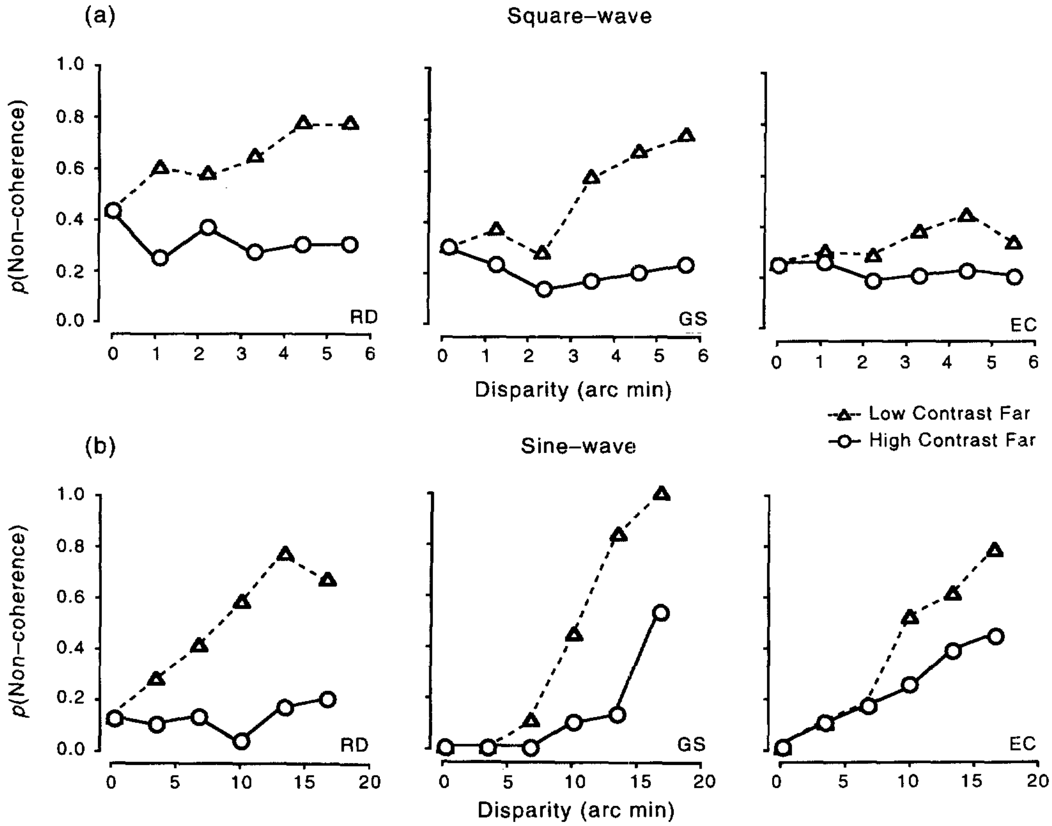

FIGURE 5. In Experiment 2, the influence of contrast-based occlusion on perceptual motion coherence was evaluated by calibrating it against binocular disparity cues. Subjects viewed a moving plaid on each trial (see Figs 1–3). Two independent variables were manipulated: (i) sign of binocular disparity; (ii) magnitude of binocular disparity. All other stimulus properties, except grating type, were constant. Following each brief stimulus presentation, subjects were required to report whether the component gratings appeared to move coherently as a single surface or move non-coherently across one another. Data are shown for each of three subjects (columns) who viewed each of two grating types (rows). (a) Proportion of trials on which subjects reported non-coherent motion of rectangular-wave gratings is plotted as a function of binocular disparity magnitude for each disparity sign (open triangles = low-contrast–far; open circles = high-contrast–far). For low-contrast–far stimuli, the probability of non-coherent motion judgments increased reliably with increasing disparity magnitude. By contrast, non-coherent motion was rarely reported for high-contrast–far stimuli, regardless of disparity magnitude. (b) The same conditions were repeated using sine-wave gratings. Perceived motion non-coherence exhibited a similar dependence upon the conjunction of disparity sign and magnitude. For subjects RD and GS each data point is based on a total of 30 trials. For subject EC, each data point is based on 100 trials. See text for details.

Results obtained using moving rectangular-wave plaids are shown in Fig. 5(a). Plots for each of three subjects indicate the proportion of trials on which the component gratings were judged to drift non-coherently across one another. As for the depth judgments of Experiment 1, this proportion is plotted separately for blocks of trials in which the low-contrast–far (filled triangles) and high-contrast–far (filled circles) stimuli were viewed. For all three subjects, these motion judgments parallel the depth judgments (cf. Figure 4). For all non-zero disparities, the proportion of trials on which motion non-coherence was reported was greater for low-contrast–far then for high-contrast–far stimuli. A similar pattern of results was obtained for sine-wave plaids [Fig. 5(b), once again mirroring the depth segregation judgments. In concert with the results of Experiment 1, these results strongly suggest that relative contrast affects motion non-coherency by acting as a cue for depth-ordering.

EXPERIMENT 3: DEPTH-ORDERING AND MOTION NON-COHERENCY IN PLAID PATTERNS—ROLE OF "INTERSECTIONS"

We have postulated that the different results obtained using low-contrast–far vs high-contrast–far plaids in Experiments 1 and 2 are a consequence of the "accidental" creation of X-junctions in which the high-contrast grating resembles a foreground occluder. An alternative explanation is that luminance contrast differences, independently of the X-junctions, account for these results. Luminance contrast could conceivably affect depth-from-disparity judgments in a number of ways. Two possibilities seem worth considering. First, it may be that luminance contrast is acting not as a depth-from-occlusion cue, but as a "depth-from-aerial perspective" cue. Aerial perspective refers to the fact that retinal images of distant objects are of lesser contrast than those of near objects, due to atmospheric scattering of light

(Egusa, 1982; O'Shea *et al.*, 1994; Vallortigara & Bressan, 1991; Stoner & Albright, 1993). Consistent with this interpretation, various investigators have demonstrated that low-contrast figures appear far relative to high-contrast figures (e.g. O'Shea *et al.*, 1994; Schor & Howarth, 1986; Rohaly & Wilson, 1993). In support of this idea, Vallortigara & Bressan (1991) have reported that occlusion and relative luminance contrast interact in their influence on motion coherency. Based on their findings, they have suggested that luminance contrast impacts both depth-ordering and motion coherency in plaids patterns by virtue of the visual system's tendency to assume that objects with higher contrast are closer.

A second possibility is that low-contrast gratings simply provide a weaker depth-from-disparity signal than do high-contrast gratings. The validity of our depth-from-occlusion hypothesis can be weighed against these alternative explanations by exploiting the fact that they depend differentially upon figural overlap and full viewing of all four plaid sub-regions (A–D, defined in Fig. 2, right panels). Clearly, it is the contrast relationships between all four of these subregions that define an X-junction; hence all four must be visible to render the hypothesized depth-ordering based on transparent occlusion. Conversely, if simple differences in luminance contrast between the components—such as those upon which the alternative accounts are founded—are responsible for the observed depth-ordering effects, visibility of regions B–D should be sufficient to reproduce those effects—the region of perceived overlap should not be essential. This assertion follows from the observation that subjects typically perceive region D (defining the intersection of the dark phase of both gratings) as background and hence region D should constitute the yardstick against which relative contrast is measured. An explanation based on depth-from-aerial perspective thus predicts that region B (the bright phase of the high-contrast grating) should appear closer than region C (the bright phase of the low-contrast grating) because it possesses higher contrast relative to region D. This depth-ordering is in the same direction as that promoted by the putative depth-from-occlusion cue and thus constitutes a potential alternative explanation of our findings. Conversely, if region A were perceived as the background against which the contrasts of regions B and C were measured, the opposite depth-ordering should prevail as region B contrasts less with region A than does region C. Because this depth-ordering is in the wrong direction to explain our findings, removal of region A, should, if anything, increase the contribution of luminance contrast acting as a depth-from-aerial perspective cue.

Experiment 3 was designed to evaluate the merits of these different hypotheses by determining whether the visibility of all four plaid sub-regions is critical to the effects seen in Experiments 1 and 2 (as required by the depth-from-occlusion hypothesis) or whether sensitivity to the luminance relationships between regions B, C, and D is sufficient (as implied by the alternative hypotheses).

Method

Visual stimuli and psychophysical procedure. Plaid subregion A was "erased" by superimposing static filled circles, or "patches" having zero horizontal binocular disparity [Fig. 6(a)]. To ensure that any differences between the percepts elicited by these stimuli and those employed in Experiments 1 and 2 were not simply due to the addition of static circles, we also included a condition in which unfilled circles, or "rings," were used in place of patches [Fig. 6(b)]. The presence of patches vs rings thus constituted a third independent variable in Experiment 3, in addition to the sign and magnitude of binocular disparity. Patches/rings were 0.50 deg diameter. (For comparison, region A was 0.19 deg along the vertical diagonal.) Two of three subjects (GS and RD) viewed patches/rings that were of the same luminance as region D (3.6 cd/m^2). The remaining subject (EC) viewed dark patches/rings of an intensity that could be distinguished from all plaid subregions (0.1 cd/m^2). Plaids possessed a mean luminance of 11.2 cd/m^2 (neglecting the contribution from superimposed "patches" and "rings"). In all other respects, the method of plaid construction in Experiment 3 was identical to that used for the rectangular-wave plaids in Experiments 1 and 2. A combination of spatial frequency (1.38 c/deg) and duty cycle (0.1875) was chosen so that regions A, B and C were of the same size used for the rectangular-wave plaids of Experiments 1 and 2. The larger spatial frequency (relative to Experiments 1 and 2) insured that patches did not obscure regions other than A. In one set of experimental conditions, the plaid stimuli were static, as they were in Experiment 1; in a second set of conditions, plaids moved. Plaid motion was up and down, such that regions corresponding to those labeled A in Fig. 2 were never revealed. This oscillatory motion was sinusoidal with a total angular excursion of 0.22 deg and cycle period of 0.67 sec. Average component speed was 0.47 deg/sec yielding an average plaid speed of 0.66 deg/sec.

The presence or absence of motion was an additional independent variable, which was associated with different psychophysical task requirements. The dependent variable for static plaid conditions was perceptual depth segregation of the component gratings. Subject instructions and task requirements were identical to those used in Experiment 1. The dependent variable for moving plaid conditions was perceptual motion coherence. For this condition, subject instructions and task requirements were identical to those used in Experiment 2.

Results

The depth-from-occlusion hypothesis predicts that the invisibility of plaid subregion A should prevent differential component contrast from contributing to depth segregation and motion non-coherency. In our experiments, this effect should be manifest as a loss of the interaction between disparity sign and magnitude, such that perceptual judgments (depth segregation and motion coherence) of low- and high-contrast–far conditions

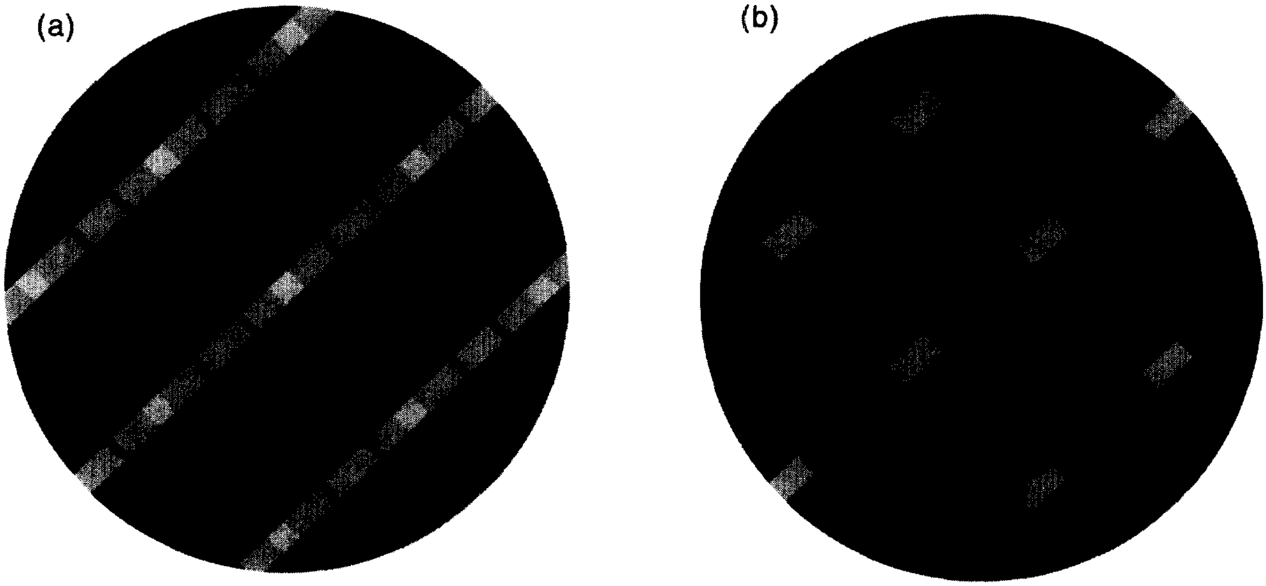

FIGURE 6. Illustration of plaid designs used for Experiment 3 (actual stimuli were stereoscopic pairs based on these designs). Component gratings and plaid construction were identical to rectangular-wave plaids used in Experiments 1 and 2, with the exception that grating duty cycles were slightly more asymmetric in Experiment 3 (0.1875 vs 0.25). Plaids were either static or moved depending on the required perceptual report (static: depth segregation; moving: motion coherence). Motion was along the same axis as in Experiment 2 (i.e., equal component speeds) but was sinusoidal (up and down) with a total angular excursion of 13.2 arc min and cycle period of 0.67 sec. Peak component speed was 2.34 deg/sec yielding a peak plaid speed of 3.31 deg/sec. The most important difference between these stimuli and those used in Experiments 1 and 2 was the addition of "patches" or "rings." These features were static and placed over the points of intersection of the bright phases of each grating, in positions corresponding to the intersection centers (for static plaids) or the mid-point of the oscillatory path of each intersection (for moving plaids). For two of three subjects (RD and GS) tested, the luminance of patches and rings was identical to the luminance of the darkest sub-region of the plaid (i.e., the points of intersection of the dark phases of each component grating: 3.6 cd/m^2), as depicted here. For the remaining subject (EC), the luminance of patches and rings was darker than all plaid sub-regions (approx. 0.1 cd/m^2). (a) Patch placement and appearance. For static plaids, patches occluded all intersection regions, thereby eliminating visibility of X-junctions. For moving plaids, patch diameter just exceeded peak-to-peak oscillatory excursion plus the intersection diagonal, thus occluding all intersection regions at all times, and likewise eliminating the visibility of X-junctions. (b) Ring placement and appearance. Rings were identically placed and of the same diameter as patches. They did not occlude X-junctions and simply served as a control for the presence of static features unrelated to plaid geometry.

show equal sensitivity to disparity magnitude. The alternative explanations addressed above predict that the interaction between disparity sign and magnitude should be unaffected by the presence of patches. The prediction of the depth-from-occlusion hypothesis is strongly supported by our results.

Depth segregation. Results obtained using static plaids with patches are shown in Fig. 7(a). Plots for each of three subjects indicate the proportion of trials for which the component gratings were judged to lie in different depth planes. Compared with the results of Experiment 1 [cf. Figure 4(a)], the difference between reports elicited by the low-contrast–far vs high-contrast–far condition was much reduced. This was true regardless of whether subjects viewed patches that were the same luminance as subregion D (GS and RD) or a different luminance (EC). Results obtained using stimulus conditions that were identical save the use of rings instead of patches (thus possessing visible X-junctions) are shown in Fig. 7(b). This experimental condition served as a control for the influence of superimposed stationary features. The asymmetry found in Experiment 1 between the two disparity sign conditions was restored for these stimuli.

These findings demonstrate that the results of Experiment 1 cannot be accounted for by hypothesizing that low-contrast gratings simply render a weak binocular disparity signal: all subjects reported depth segregation in the low-contrast–far condition when the X-junction was obscured by occlusion of region A. The absence of any systematic difference between the percepts elicited by the two disparity sign conditions when patches were present also rules out the possibility that depth-from-aerial perspective plays a major role in these experiments. Indeed, these results exclude explanations of any type that are based on luminance contrast differences without reference to figural overlap and X-junctions. Only when the X-junction was fully visible did subjects have difficulty perceiving the high-contrast grating to lie behind the low-contrast grating.

Motion coherence. The second set of stimulus conditions in Experiment 3 was designed to confirm that full viewing of the X-junctions was critical for motion coherency asymmetries, as well as depth-ordering. Results obtained using moving plaids with patches are shown in Fig. 8(a). Plots for each of three subjects indicate the proportion of trials on which the component

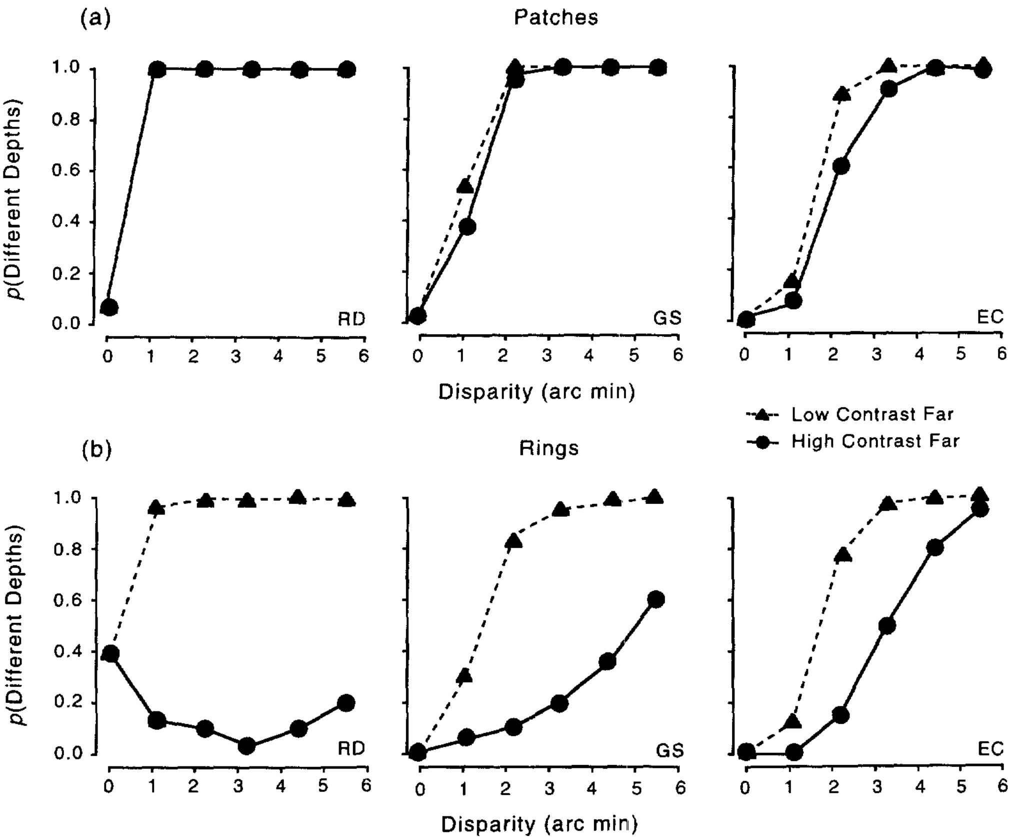

FIGURE 7. One set of conditions in Experiment 3 examined the extent to which depth segregation judgments of the sort made in Experiment 1 were dependent upon the visibility of X-junctions. The method of stimulus construction is illustrated in Fig. 6, and stimulus details are described in the legend to Fig. 6. As in Experiment 1, independent variables included both disparity sign and magnitude. The visibility of X-junctions, as manipulated by the presence of "patches" (non-visible X-junctions) or "rings" (visible X-junctions), constituted a third independent variable. All other stimulus properties were constant. Following each brief stimulus presentation, subjects were required to report whether the component gratings appeared to lie in the same or different depth planes. Data are shown for each of three subjects (columns) who viewed plaids adorned with either patches or rings (rows). (a) Data obtained for the condition in which X-junctions were occluded by patches. Plotted are proportions of trials on which subjects reported depth segregation as a function of binocular disparity magnitude for each disparity sign (filled triangles = low-contrast–far; filled circles = high-contrast–far). Probability of depth segregation reports increased dramatically with increasing disparity magnitude, regardless of disparity sign. Because disparity sign did not interact with perceptual depth-ordering under these conditions (cf. Figure 4), it appears that the effects of differential component contrast are due to the occlusion cue (i.e., the X-junction). (b) The same conditions were repeated with X-junction visibility restored (rings instead of patches). Once again, perceived depth-ordering exhibited dependence upon the conjunction of disparity sign and magnitude, implicating contrast-based occlusion cues. For subjects RD and GS each data point is based on a total of 30 trials. For subject EC, each data point is based on 100 trials. See text for details.

gratings were judged to drift non-coherently across one another. Compared with the results of Experiment 2 [cf. Figure 5(a)], the difference between reports elicited by low-contrast–far and high-contrast–far conditions was much reduced. These results parallel depth segregation judgments in that subjects were more likely to report motion non-coherency as a function of increasing binocular disparity, regardless of the disparity sign. Again, this was true regardless of whether subjects viewed patches that were the same luminance as subregion D (GS and RD) or a different luminance (EC). Results obtained using stimulus conditions that were identical save the use of rings instead of patches (thus possessing visible X-junctions) are shown in Fig. 8(b). The asymmetry found in Experiment 2 between the two disparity sign conditions was restored for these stimuli. These results strongly support the hypothesis that X-junctions, serving as a cue for depth-from-occlusion, are the dominant factor in determining depth segregation and motion non-coherency of plaid patterns.

GENERAL DISCUSSION

Our findings confirm that luminance contrast affects motion coherency in plaid patterns by stimulating mechanisms sensitive to depth-from-occlusion. In doing so, these results expose the insufficiency of image-based accounts. In the remainder of the Discussion, we will address the relevance of these findings to: (1) luminance contrast as a depth-cue; (2) models of motion non-coherence; and (3) image- vs scene-based mechanisms.

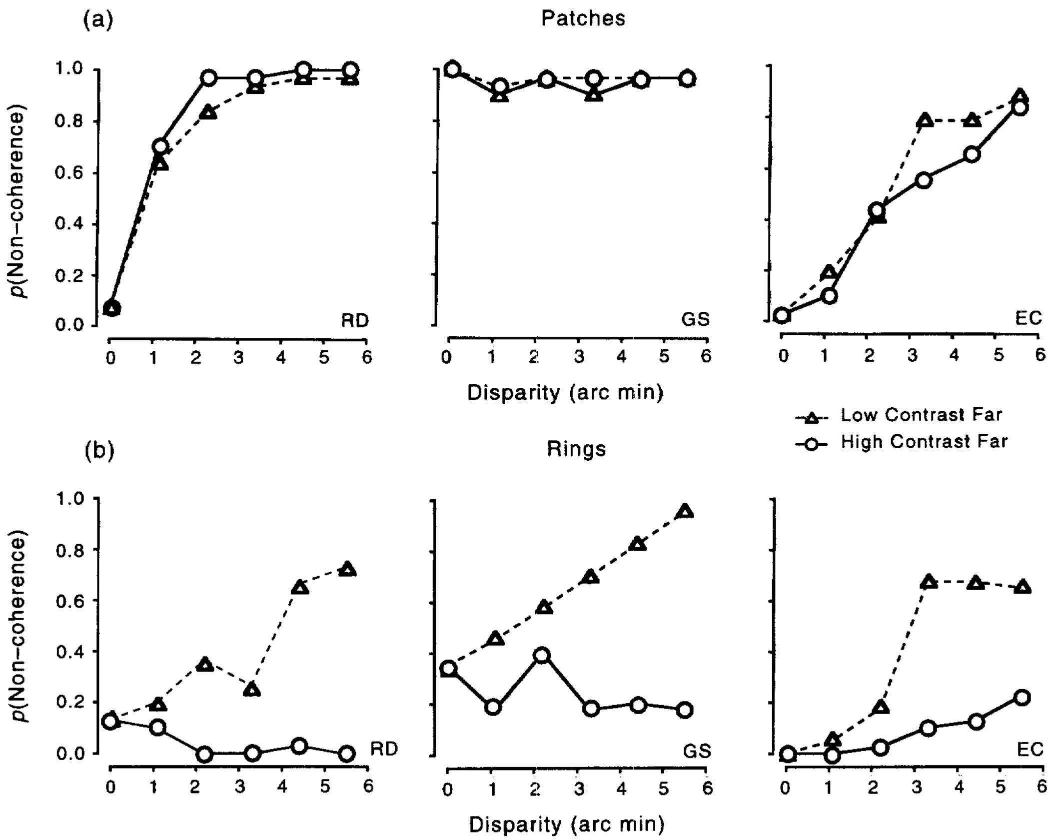

FIGURE 8. A second set of conditions in Experiment 3 examined the extent to which motion coherence judgments of the sort made in Experiment 2 were also dependent upon the visibility of X-junctions. The method of stimulus construction is illustrated in Fig. 6, and stimulus details are described in the legend to Fig. 6. As in Experiment 2, independent variables included both disparity sign and magnitude. The visibility of X-junctions, as manipulated by the presence of "patches" (non-visible X-junctions) or "rings" 2 (visible X-junctions), constituted a third independent variable. Stimuli were moved in an oscillatory fashion (see legend to Fig. 6). Following each brief stimulus presentation, subjects were required to report whether the component gratings appeared to move coherently as a single surface or move non-coherently across one another. Data are shown for each of three subjects (columns) who viewed plaids ornamented with either patches or rings (rows). (a) Data obtained for condition in which X-junctions were occluded by patches. Plotted are proportions of trials on which subjects reported motion non-coherence as a function of binocular disparity magnitude for each disparity sign (open triangles = low-contrast–far; open circles = high-contrast–far). As for perceptual depth segregation (cf. Figure 7), the probability of motion non-coherence reports increased dramatically with increasing disparity magnitude, regardless of disparity sign. (b) The same conditions were repeated with X-junction visibility restored (rings instead of patches). Once again, perceptual motion coherence exhibited dependence upon the conjunction of disparity sign and magnitude, implicating contrast-based occlusion cues. For subjects RD and GS each data point is based on a total of 30 trials. For subject EC, each data point is based on 100 trials. See text for details.

Luminance contrast as a depth-cue

Several studies have demonstrated that luminance contrast affects depth perception. Of special pertinence are those studies that used stereoscopic methods to assay the potency of contrast cues (Fry *et al.*, 1949; Schor & Howarth, 1986; Rohaly & Wilson, 1993). In each case, these studies found that low-contrast stimuli tend, under some circumstances, to be seen as more distant than high-contrast stimuli of the same disparity. This propensity has been attributed to a mechanism incorporating tacit "knowledge" of the fact that distant objects frequently contrast less with the background than do closer objects as a result of atmospheric scattering of light (Egusa, 1982; O'Shea *et al.*, 1994). Two of three subjects in our experiments (GS and EC) did, in fact, exhibit a slight tendency to report separation in depth more frequently for the low-contrast far conditions (Fig. 7, top row). This depth-ordering bias was, however, much smaller than that observed when the occlusion cue was present (Fig. 7, bottom row). Our findings thus demonstrate that, under the conditions of our experiments, the contribution of depth-from-aerial perspective is minor compared with that of depth-from-occlusion. Moreover, this contribution, if indeed genuine, had no significant impact on the perception of motion coherency (Fig. 8).

Notably, our conclusions regarding the contribution of depth-from-aerial perspective appear at variance with those of Vallortigara & Bressan (1991); see also Bressan *et al.* (1993), who have suggested that depth-from-aerial perspective influences both perceptual depth segregation and motion non-coherency. This seeming discrepancy deserves closer examination. The general approach adopted by Vallortigara and Bressan was similar to our own and involved measuring the strength of the putative depth-cue of luminance contrast by placing it in opposition to, or in support of, a second cue. These investigators constructed plaid patterns in which depth-from-aerial perspective was pitted against depth-from-occlusion (recall that, for our experiments, aerial perspective and occlusion cues in Experiments 1 and 2 supported the same depth-ordering and that the occlusion cue was absent for the "patches" condition of Experiment 3). They then compared the degree of perceptual motion coherency resulting from two different depth-ordering conditions: (1) "lighter in front" (corresponding to a low-contrast grating occluded by a high-contrast grating); and (2) "darker in front" (corresponding to a high-contrast grating occluded by a low-contrast grating). The "darker in front" condition was found to elicit more perceptual non-coherency than the "lighter in front" condition. This result was interpreted as reflecting an interaction between the two depth-cues of luminance contrast and occlusion, such that when the two depth-cues conflicted (i.e. "lighter in front"), depth-segregation, and hence motion non-coherency, was less likely. We, on the other hand, found little evidence for contrast *per se* having a substantial role as a depth-cue in plaid patterns. Without the occlusion cue (i.e., the "patches" conditions of Experiment 3), the high-contrast–far and low-contrast–far conditions (analogous to the "lighter in front" and "darker in front" conditions employed by Vallortigara and Bressan) elicited equivalent levels of non-coherency (Fig. 7). This lack of interaction between luminance contrast and binocular disparity indicates that depth-from-luminance contrast was not a significant factor.

There are several procedural differences between the Vallortigara and Bressan experiment and our own that may account for the divergent findings. For example, although Vallortigara and Bressan did not characterize their plaid stimuli in terms of component grating contrast, it may be that the contrast differences they used were more effective at promoting depth-ordering. An alternative hypothesis asserts that results of Vallortigara and Bressan are due to factors unrelated to depth-ordering. In their experiments the depth-ordering promoted by the occlusion cue was manipulated by adjusting the intensity of specific plaid sub-regions. This manipulation is known to alter the distribution of moving Fourier components (see Stoner & Albright, 1996). Thus, while the two occlusion conditions of Vallortigara and Bressan are identical in terms of implied surface depth-ordering (in the sense that only the color—a quality incidental to depth—of the depth-ordered surfaces changes), they do have different distributions of Fourier components. The distribution of Fourier energy is well known, in turn, to play a role in motion coherency, as shown by experiments in which the angle between components gratings was varied (Movshon *et al.*, 1985). If the low-contrast-occluding–high-contrast condition used by Vallortigara and Bressan were to have possessed a greater balance of Fourier components moving in the coherent direction than did the complementary condition, that condition would be expected (for reasons completely unrelated to depth-ordering) to produce higher levels of coherency.

Unlike Vallortigara and Bressan's depth-from-occlusion manipulation, the depth-from-disparity manipulations employed in the experiments reported here are not accompanied by variation in the distribution of Fourier components. Whether the different answers yielded by these two studies can be accounted for by the above "Fourier components" explanation might be resolved by directly assaying perceptual depth-ordering (as done in our study) as well as by submitting the stimuli of Vallortigara and Bressan to Fourier analysis (Stoner & Albright, 1996).

Mechanisms of motion coherence

The discovery by Adelson & Movshon (1982) that plaid patterns were more likely to be perceived as non-coherent when the component gratings were of different spatial frequencies led those investigators to hypothesize that the motions of dissimilar gratings are processed by independent channels. This channel hypothesis was given some plausibility by the finding that neurons in cortical visual motion processing area MT are selective for spatial frequency (Newsome *et al.*, 1983). Adelson and Movshon also reported that similarity along the dimension of luminance contrast influenced motion coherency. While one would be tempted to invoke the channel hypothesis in this case as well, there exists little evidence for neurons selective for specific luminance contrasts. Moreover, an explanation based on luminance contrast channels cannot easily account for the observation reported herein; the level of motion coherency associated with superimposed gratings of different contrasts depends upon their stereoscopic depth-ordering. It could be countered that the underlying channels are preferentially tuned to specific conjunctions of contrast and stereo disparity. This elaborated channel model is invalidated by the findings of Experiment 3, however, which demonstrate that the interactive influences of contrast and stereo disparity on motion coherency disappear with the removal of the occlusion cue.

An alternative to the channel hypothesis has been offered by Wilson & Kim (1994). These investigators have championed an model of motion detection in which a separate "non-Fourier" motion pathway supplements the standard "Fourier" pathway. This model accounts for the effect of luminance contrast on motion coherency by observing that the amplitude of the non-Fourier signal moving in the coherent pattern direction is greater for same-contrast plaids than for different-contrast plaids. Kim and Wilson's innovative scheme does not include any provision for the role of depth-ordering and feature classification and hence does not address the findings presented here. On the basis of both experimental findings and theoretical considerations, we have previously argued (Stoner & Albright, 1995, 1996) that neurons responsive to non-Fourier motion might serve more generally to detect depth-from-occlusion. According to this hypothesis, neurons responsive to non-Fourier image variation may influence motion coherency—not by supplying information about motion *per se*, but rather by providing information about depth-ordering.

Image- vs scene-based mechanisms

Visual processing has typically been portrayed as involving image- and scene-based stages (e.g. Marr, 1982; Barlow, 1972). Image-based mechanisms, it is presumed, extract elementary image features and higher-order mechanisms, in turn, interpret these image primitives in terms of objects in the physical world. Debate about a particular problem in vision frequently centers around whether that problem is best investigated as an image- or a scene-based process. At stake are fundamental issues of both experimental design and interpretation. If one's object of study is presumed to be an image-based process, one is inclined to ignore any perceived correspondence between retinal image properties and visual scene attributes, because the latter are thought to be inferred at a later processing stage. Investigators are thus free to design stimuli and build theories without considering the problem of how a representation of the visual scene is constructed from the retinal image. Whether a sine-wave grating, for example, resembles a corrugated surface becomes an irrelevant consideration. The channel hypothesis advanced by Adelson & Movshon (1982), discussed above, is a classic example of this type of image-based perspective in that it makes reference, not to scene elements (e.g. surfaces and objects), but solely to image primitives (i.e., oriented components of a particular spatial frequency). If, on the other hand, the involvement of a scene-based mechanism is suspected, the perceptual interpretation of one's visual stimulus in terms of scene attributes (in other words, what the stimulus looks like) cannot be ignored.

The possibility that the extent to which a plaid pattern resembles two superimposed surfaces might influence perceptual motion coherency was investigated by Stoner *et al.* (1990). It was found that plaid patterns constructed to appear as one transparent grating overlying another tend not to cohere. This was taken as evidence for the role of surface segmentation mechanisms in motion processing. The need to invoke scene-based mechanisms to account for these results has been challenged and alternative image-based explanations offered (e.g., Kim & Wilson, 1993). However, the weight of evidence from recent experiments designed to explicitly evaluate these alternative accounts rather conclusively favors the scene-based explanation (Trueswell & Hayhoe, 1993; Stoner & Albright, 1996; Lindsey & Todd, 1996).

The results presented herein carry the implications of Stoner *et al.* one step further by offering an important object lesson: scene-based processes may intrude regardless of whether they are anticipated. Our stimuli were not designed originally with the intent to simulate occlusion and are, moreover, not ideally suited to stimulate depth-from-occlusion mechanisms. The visual system apparently does its best to interpret such artificial images in terms of real-world scene elements. While this conclusion is not terribly surprising when considered in terms of the functions for which the visual system evolved, it has far-reaching ramifications. Our findings suggest that image-based processes cannot easily be investigated independently of scene-based mechanisms.

REFERENCES

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

Adelson, E. H. & Movshon, J. A. (1984). Binocular disparity and the computation of two-dimensional motion. *Journal of the Optical Society of America, 1*, 1266.

Albright, T. D. & Stoner, G. R. (1995). Visual motion perception. *Proceedings of the National Academy of Sciences USA, 92*, 2433–2440.

Barlow, H. B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception, 1*, 371–394.

Beck, J., Prazdny, K. & Irvy, R. (1984). The perception of transparency with achromatic colors. *Perception & Psychophysics, 35*, 407–422.

Bressan, P., Ganis, G. & Vallortigara, G. (1993). The role of depth stratification in the solution of the aperture problem. *Perception, 22*, 215–228.

Dobkins, K. R., Stoner, G. R. & Albright, T. D. (1992). Oculomotor responses to perceptually coherent and non-coherent plaids. *Society of Neuroscience Abstracts, 18*, 1034.

Egusa, H. (1982). Effect of brightness on perceived distance as a figure–ground phenomenon. *Perception, 11*, 671–676.

Fry, G. A., Bridgman, C. S. & Ellerbrock, V. J. (1949). The effects of atmospheric scattering on binocular depth perception. *American Journal of Optometry, 26*, 9–15.

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

Kooi, F. L., De Valois, K. K. & Switkes, E. (1992). Higher order factors influencing the perception of sliding and coherence of a plaid. *Perception, 21*, 583–598.

Krauskopf, J. & Farell, B. (1990). Influence of colour on the perception of coherent motion. *Nature, 348*, 328–331.

Lindsey, D. T. & Todd, J. T. (1996). On the relative contributions of motion energy and transparency to the perception of moving plaids. *Vision Research, 36*, 207–222.

Marr, D. (1982) *Vision: A computational investigation into the Human representation and processing of visual information*. New York: Freeman.

Metelli, F. (1974). The perception of transparency. *Scientific American, 230*, 91–95.

Movshon, J. A., Adelson, E. A., Gizzi, M. & Newsome, W. T. (1985). The analysis of moving visual patterns. In C. Chagas, R. Gattass & C. G. Gross (Eds), *Study group on pattern recognition mechanisms* (pp. 117–151). Vatican City: Pontifica Academia Scientiarum.

Murphy, B. J., Kowler, E. & Steinman, R. M. (1975). Slow oculomotor control in the presence of moving backgrounds. *Vision Research, 15*, 1263–1268.

Newsome, W. T., Gizzi, M. S. & Movshon, J. A. (1983). Spatial & temporal properties of neurons in macaque MT. *Investigative Ophthalmology and Visual Science Supplement, 24*, 106.

O'Shea, R. P., Blackburn, S. G. & Ono, H. (1994). Contrast as a depth-cue. *Vision Research, 34*, 1595–1604.

Rohaly, A. M. & Wilson, H. R. (1993). The role of contrast in depth perception. *Investigative Ophthalmology and Visual Science, 34*, 1437.

Schor, C. M. & Howarth, P. A. (1986). Suprathreshold stereo-depth matches as a function of contrast and spatial frequency. *Perception, 15*, 249–258.

Shimojo, S., Silverman, G. H. & Nakayama, K. (1989). Occlusion and the solution to the aperture problem for motion. *Vision Research, 29*, 619–626.

Stoner, G. R. & Albright, T. D. (1992). Neural correlates of perceptual motion coherence. *Nature, 358*, 412–414.

Stoner, G. R. & Albright, T. D. (1993). Image segmentation cues in motion processing: implications for modularity in vision. *Journal of Cognitive Neuroscience, 5*, 129–149.

Stoner G. R. and Albright, T. D. (1994). Visual motion integration: a neurophysiological and psychophysical perspective. In A. T. Smith & R. J. Snowden (Eds), *Visual detection of motion*. London: Academic Press.

Stoner, G. R. & Albright, T. D. (1995). Depth from occlusion and non-Fourier motion. *Society of Neuroscience Abstracts, 21*, 511.

Stoner, G. R. & Albright, T. D. (1996). The interpretation of visual motion: evidence for surface segmentation mechanisms. *Vision Research, 36*, 1291–1310.

Stoner, G. R., Albright, T. D. & Ramachandran, V. S. (1990). Transparency and coherence in human motion perception. *Nature, 344*, 153–155.

Trueswell, J. C. & Hayhoe, M. M. (1993). Surface segmentation mechanisms and motion perception. *Vision Research, 33*, 313–328.

Vallortigara, G. & Bressan, P. (1991). Occlusion and the perception of coherent motion. *Vision Research, 31*, 1967–1978.

Wilson, H. R. & Kim, J. (1994). A model for motion coherence and transparency. *Visual Neuroscience, 11*, 1205–1220.

Acknowledgements—We thank L. Croner, R. Duncan, and Y.-X. Zhou for valuable discussions and comments on the manuscript, J. Costanza for technical assistance, and M. Mitchell for administrative assistance. Rodel Divina and Erika Crawford deserve special thanks for their technical and intellectual contributions. This work was supported by NIH grant MH-50706.