

Hysteresis, cooperativity, and depth averaging in dynamic random-dot stereograms

BARTON L. ANDERSON
Vanderbilt University, Nashville, Tennessee

Experiments were performed to assess the response of the human visual system to dynamic random-dot patterns composed of disparity mixtures. In Experiment 1, the perceived depth and relative stability of two patterns were compared; one pattern depicted two transparent layers of dots, and the other depicted a volume of dots. Two effects were found: (1) the volume pattern exhibited a large degree of disparity averaging; and (2) asymmetries were observed in the relative stability of these two patterns. Experiment 2 was designed to determine whether these findings could be attributed to spatially localized processes occurring at the location of disparity discontinuities. This was accomplished by introducing unpaired noise points localized either along the disparity discontinuities or in the center of the layered and volume patterns. The amount of depth averaging and the direction of the asymmetry did not appear to depend on processes localized along the disparity discontinuities. Results of these experiments, taken in conjunction with those of previous studies, suggest that hysteresis is independent of cooperative persistence mechanisms.

Two general types of network models have emerged to represent stereoscopic space. Historically, the first class of models to emerge were *cooperative* models. Within the field of neural networks, cooperativity has been associated with excitatory feedback between detectors tuned to the same (Marr & Poggio, 1976; Prazdny, 1985) or nearby (Pollard, Mayhew, & Frisby, 1985) disparities. This feedback is generally ascribed two functions: to aid in the process of defining binocular correspondence, and to endow the network with hysteretic stability. Indeed, the phenomenon of hysteresis has been taken to be a signature of cooperative interactions (Chang & Julesz, 1984; Julesz, 1971, 1983). The putative benefit of hysteresis is that it maintains the organization of a percept once correspondence has been defined.

Two experimental results have been used to support the concept of cooperativity. The first of these is the phenomenon of fusional hysteresis originally documented by Fender and Julesz (1967). They found that the transformation between fusion and the unfused state was not reversible; an initially fused stimulus had a larger disparity limit than did a stimulus that was preceded by decorrelation. Fender and Julesz argued that this asymmetry was due to cooperative interactions that maintained the stability of the fused percept, causing an enlargement of the fusional zone. The second result cited as support for cooperativity was the "pulling" effect described by Julesz and Chang (1976). In this paper, Julesz and Chang demonstrated that a few percent of bias dots can "pull" an

ambiguous stereogram to the depth of the unambiguous points. Pulling "strength" was strongest when the disparity of the bias dots was the same as the disparity of one of the ambiguous planes, and it became stronger for longer exposure durations. Thus, cooperativity is putatively greatest between detectors tuned to similar disparities, but some processing time is required for the feedback to be effective.

A number of predictions can be made by identifying cooperativity with excitatory feedback. Cooperative feedback has been described as a mechanism that preserves the stability of arrays of similar disparities (see, e.g., Chang & Julesz, 1984; Grossberg, 1987; Julesz, 1971; Prazdny, 1985); that is, cooperative feedback putatively endows stereopsis with hysteretic stability. The thesis of cooperative feedback therefore suggests that a disparity field composed of similar disparities should be more stable than one composed of a mixture of disparities. Thus, a more *ordered* pattern (similar disparities) should exhibit greater stability than a more *disordered* pattern (a mixture of disparities).

Note that the thesis of cooperative hysteresis was invoked to explain the stability of organized percepts, such as fusion. But there is a paradox generated by the prediction that ordered percepts are stable. This paradox can be understood best through consideration of the experiments of Julesz and Tyler (1976; Tyler & Julesz, 1976, 1978), who measured sensitivity for detecting changes in the correlation of dynamic random-dot stereograms. The general result that they reported was that it took longer to detect disorder-order transitions than to detect order-disorder transitions; in other words, the disordered state (decorrelation) was more stable than the ordered state (correlation). This finding is consistent with the physical


I would like to acknowledge the time and support of J. Lappin in the preparation of an earlier version of this manuscript. Correspondence should be addressed to Barton L. Anderson, Harvard University, William James Hall, 33 Kirkland St., Cambridge, MA 02138.

principle of entropy, which attributes the greatest stability to disordered states. Julesz and Tyler recognized this connection and dubbed their phenomenon *neuronentropy*. The paradox of the cooperative prediction is that it suggests that the correlated distribution should be more stable than the uncorrelated distribution, but the opposite result was obtained.

It is relatively easy to understand why an array of similar disparities (i.e., an ordered pattern) should be stable under the thesis of cooperativity. But what is the physical reasoning behind the stability of *disordered* states? In statistical mechanics, the most stable distribution is the one that can be realized in the greatest number of ways (i.e., the most "probable" state). Clearly, there are more ways to realize an uncorrelated stereogram than a perfectly correlated stereogram. Indeed, there is only one way to realize perfect correlation (assuming that one of the half images is fixed). The fact that this is related to stability can be seen if one considers the impact of perturbations on the two patterns. Any small perturbation of the decorrelated stimulus will leave the pattern decorrelated; but small perturbations of correlated patterns will generate decorrelation. Thus, correlation should be physically unstable.

One argument that may help rationalize this paradox is to note that the cooperative thesis holds only for displays that have computable disparities for all of the points in the image, which was clearly not the case for the neuronentropy studies. Indeed, Tyler (1991) has suggested that the finding of neuronentropy may be unique to the percept of "lustre" generated by decorrelated displays. Therefore, a fair test of cooperative hysteresis requires that all of the points in the image have matches; after all, cooperativity is hypothesized to be a form of interaction between disparity-sensitive mechanisms, and disparity computations require matches by definition. In the experiments described below, two patterns were designed for such a comparison: a layered pattern and a volume pattern (see Figure 1). The far and near extremes of the volume pattern were assigned the same disparities as were the two surfaces of the layered pattern, but half the points within the volume were assigned disparities that fell between these values. Thus, the layered and volume patterns differed in the distribution of disparities present, but the total range of disparities selected was identical for both. Note that the layered pattern has only two disparity values, whereas the volume has a random distribution of disparities. Cooperative hysteresis would predict the layered pattern to be more stable than the volume.¹ In contrast, statistical mechanics leads us to conclude that the opposite should be true: the volume patterns (the disordered pattern) should be more stable than the layered pattern (the ordered pattern). Again, this argument hinges on the observation that the most "probable" distribution is physically the most stable. As with decorrelation, there are more ways to realize a random distribution of disparities than two discrete layers of disparity. Furthermore, any small perturbation of the random distribution will generate a very similar distribution, while small perturbations

Layered Pattern



Volume Pattern



Figure 1. Schematic illustration of the cyclopean percept of the two patterns used in the experiments. The layered pattern consisted of dots with two discrete values of disparity, giving rise to a percept of two transparent sheets of dots displaced symmetrically in depth around the ground plane. The volume pattern was similar, except that half of the dots were assigned random disparities between the values of the front and back planes of the layered pattern. These disparities were assigned on the basis of a random selection from a uniform probability distribution. The displays used in the experiment were dynamic random-dot stereograms with a frame duration of 20 msec.

of disparity layers will lead to more volume-like distributions.² One of the goals in these experiments was to determine which of these two patterns is the most stable.

The second class of stereopsis models that originally emerged solved correspondence in a coarse-to-fine manner. Instead of using local interactions among disparity mechanisms to solve correspondence, coarse-to-fine models solve correspondence by matching the outputs of spatial filters in a coarse-to-fine manner. The first model of this type was described in a paper by Marr and Poggio (1979). The motivation for this model was similar to the motivation behind its cooperative predecessors: to find a strategy that facilitates the process of establishing binocular correspondence. The coarse-to-fine strategy matches large spatial filter outputs first; the outputs of these matches are then used to constrain the search on smaller scales. This strategy effectively reduces the number of possible false matches that the visual system needs to consider.

Like cooperative models, coarse-to-fine strategies make predictions about the temporal dependence of stereopsis. However, these predictions focus on the computation of disparity as a function of time, not on the relative stability of disparity distributions. For example, Watt (1987) has suggested that temporal improvements in acuity can be characterized by coarse filters successively "shutting off" over durations of about 1,000 msec. Watt's model is consistent with a number of studies in which it has been demonstrated that stereoscopic acuity increases over periods of 500–1,000 msec (McKee, Levi, & Bowne, 1990; Ogle & Weil, 1958; Watt, 1987). From this perspective, depth interactions like depth averaging (Parker & Yang, 1989; Schumer, 1979) should *decrease* with time. Averaging should be greatest for the largest filters, and as filters successively "shut off," disparity averaging should diminish.

In contrast, if the coarse filters remain on while the smaller scales are "turned on," depth averaging should get slightly *larger* during this progression, because the scales being turned on would also average over some range of disparities. The thesis of cooperative feedback makes a similar prediction. Recall that the "pulling effect" revealed *stronger* depth interactions as the duration of the stimulus presentation was increased. This makes sense for a putative feedback mechanism; the feedback should exhibit a temporal latency, becoming stronger over some reasonable time frame. Thus, if mixed-disparity patterns exhibit depth interactions, these interactions would be expected to increase as a function in time.


The experiments described below were designed to assess properties of both cooperative and coarse-to-fine models of stereopsis. There were a number of specific goals in the studies presented here. First, we were interested in documenting the visual system's response to a mixed-disparity target: Do mixtures of disparity interact? Second, we were interested in measuring the (relative) temporal stability of the two patterns. There were two aspects to the question of stability. The first of these concerns the stability of depth computations: Do estimates of depth vary with exposure duration in a mixed-disparity target? If so, do they increase or decrease? The second question involves a comparison of the stability of the two organizations (layered or volume): Is there a difference between the stability of disparity layers and that of a volume of disparities?

GENERAL METHOD

We determined both the depth interactions and the relative stability of the layered and volume patterns by measuring the sensitivity of observers to transitions between the two patterns (i.e., from layers to volume, and from volume to layers). The use of this method to document disparity averaging is novel, and it requires some discussion. A variant of the method has been successfully used to study the stability of patterns (Julesz & Tyler, 1976; Tyler & Julesz, 1976, 1978). This method may be thought of as a *perturbation* methodology (see the patterns illustrated in Figure 2).

The basic idea is rather simple: the (relative) stability of a pattern should be given by its resistance to change. We have extended this methodology to provide an objective means of studying depth interactions in dynamic random-dot stereograms. The study of depth averaging has been fraught with the difficulty of establishing an objective criterion for when two surfaces are averaged (see, e.g., Parker & Yang, 1989). We have adopted a different strategy that allows us to use a simple yes/no discrimination task. The rationale behind this method requires a discussion of pilot observations reported by the subjects of these experiments. The most perceptually salient aspect of the transitions between the layered and volume

Layered Fixation



Volume Fixation




Figure 2. Schematic illustration of the perturbation methodology used in all of the experiments for the layered fixation and the volume fixation. The pictures represent sequences of frames. The first and last picture of each temporal sequence represent the fixation pattern and the backward mask. Each picture represents 50 frames shown in succession (zero ISI) of 20 msec each. The middle picture of each temporal sequence represents the perturbation sequence, which varied in number for each subject.

patterns was an apparent change in overall depth. Specifically, the transition from layers to a volume appeared as a dramatic reduction in the perceived depth interval of the whole pattern, while the transition from a volume to layers produced an equally compelling expansion.

Our first task was to determine the cause of this apparent depth difference. One possibility was that the volume pattern exhibited *depth averaging*, or *depth attraction*. On the other hand, the difference in relative depth of the layered and volume patterns may have been due to a repulsion that tends to push the layered pattern apart, generating an exaggerated or expanded depth interval for the layered form. Both attractive and repulsive interactions have been found for simple line and dot stereograms (Mitchison & Westheimer, 1984; Westheimer, 1986; Westheimer & Levi, 1987), as well as for dynamic random-dot stereograms (Stevenson, Cormack, & Schor, 1991), so it is not immediately apparent whether the depth difference represents an attraction or a repulsion. However, the disparities used for the layered pattern (25 arc min separation) in these experiments were well outside the range of either the attractions identified by Stevenson et al. (.5–1.0 arc min attraction for surfaces of 3–4 arc min separation) or the repulsions (.5–1.0 arc min repulsions for separations up to about 8 arc min). In contrast, the volume pattern was composed of disparities that should be subject to both attractive and repulsive interactions. It therefore seemed more likely that the depth interactions were occurring in the volume pattern.

We informally determined the cause of the depth difference between the layered and volume patterns by comparing the apparent depth of the front and back planes of the layered pattern with the depth of each plane displayed alone. Observers reported no noticeable difference between the apparent depth of the front or back planes of the layered pattern and the depth of the front or back plane shown in isolation (respectively). In contrast, the front and back surfaces of the volume pattern appeared to be displaced only half the depth of either plane shown in isolation. We therefore rejected the hypothesis that the transparent layers were repulsing each other, in favor of the conclusion that the perceived depth difference was due to a form of disparity averaging occurring in the volume pattern. Indeed, we chose disparities in such a way that the layered pattern exhibited no apparent depth interactions.

Our goal was to document disparity averaging by having observers discriminate transitions between these two patterns. Discrimination accuracy could then be used as an index of disparity averaging. We accomplished this goal by rescaling the disparities of the layered pattern to match the perceived depth interval of the volume. The disparities of layered pattern were therefore reduced to produce a depth interval that was perceived to be equal to the depth interval of the volume. The disparities of the volume pattern were not changed. By examining observers' discrimination performance over a wide range of scale values (where *scale* represents the proportion of

the [unscaled] disparities of the layered pattern), we would expect that the scale values that generate perceptually equivalent depth intervals would produce the worst discrimination performance. The goal of this method was to eliminate differences in the overall depth of the two patterns, making the discrimination task more difficult. Observers would therefore have to discriminate the patterns solely on the basis of differences in the distribution of dots in depth in the absence of any apparent changes in the overall depth interval. If the layered pattern is rescaled beyond the point of perceptually equivalent depth intervals (i.e., the layered pattern now appears to have a smaller depth interval than the volume does), observers should again find the task relatively easy because the perturbation would appear as an expansion (for the layers to volume transition) or a contraction (for the volume to layers transition). If observers are tested over a range of scale values, this procedure should produce V-shaped performance functions. The minimum of the V would correspond to the scale value for which the layered and volume patterns are perceived to be equal in their overall depth range.³ This method is illustrated in Figure 3.

Because we wished to measure the temporal dependence of disparity interactions in the two patterns, observers were required to perform this task for a number of temporal durations. This experiment is described below.

EXPERIMENT 1

The Temporal Dependence of Depth Averaging

Method

Stimulus patterns. Two types of dynamic random-dot stereograms were presented. One pattern depicted two "transparent" layers of dots at two distinct depth planes (i.e., two values of disparity in the same visual direction), symmetrically displaced around a zero-disparity ground. This pattern will be described as the *layered* form. The second pattern was similar (the *volume*), except that half of the dots were assigned disparities that fell between the planes of the layered pattern. These patterns are illustrated in Figure 1. The zero-disparity ground was created by randomly distributing dots within an 8 × 8 cm square region and presenting the same pattern to both eyes. This region subtended a visual angle of 4° × 4°. In the center of the (zero-disparity) ground, a circular region 4 cm in diameter (2° visual angle) was used to create the transparent layers and the volume figures. The figures were therefore foveal targets. For the layered pattern, each point in the 2° circular area had a 50% chance of being assigned a crossed or an uncrossed disparity of a fixed magnitude. Consequently, this procedure produced a percept of two flat circular planes of dots, with half of the dots on each depth plane, and each displaced the same distance (but opposite in sign) from the surrounding ground. To minimize monocular cues, each eye was assigned half of the crossed or uncrossed disparity (*crossed* and *uncrossed* disparity refer to the zero-disparity ground). For example, to create a percept of a plane behind the ground plane, the left eye's view was displaced to the left by half the total disparity, and the right eye's view was displaced to the right by half the total disparity. The image disparities (i.e., the difference in positions on the two display scopes) were 12.5 arc min for both of the transparent planes—that is, 12.5 arc min for the plane with crossed disparity, and 12.5 arc min for the plane with uncrossed disparity. This produced an overall disparity difference of 25 arc min for the central figure. This value represents the unscaled dis-


Layered Pattern**Volume Pattern**

Figure 3. Schematic illustration of the scaling method used in the experiments. The top figures represent the percepts of the layered and volume patterns when they both had a maximal disparity difference of 25 arc min. The volume pattern appeared to have a depth interval that was only about half that of the layered pattern. The amount of depth averaging exhibited by the volume pattern was evaluated by rescaling the disparities of the layered pattern and measuring discrimination performance for each scale value used. Note that the discrimination performance should be worst when the apparent change in overall depth of the patterns is removed as a cue (middle figures), but should be comparatively easy when the differences in perceived depth are large (top and bottom figures). Thus, discrimination performances should be V shaped when plotted as a function of the scale factor used (the scale factor is the proportion of the 25-arc-min disparity difference used in the unscaled layered pattern).

parities of the layered pattern—a scale factor of 1. A scale factor of .5 therefore represents a total disparity difference of 12.5 arc min, and is distributed in the same symmetric manner as are the unscaled disparities (6.25-arc-min crossed and uncrossed disparities).

The volume pattern was similar to the layered pattern, except that half of the dots in the circular figure were randomly assigned disparities between the extremes of the unscaled layered form. The probability of any given disparity assignment was uniform over the 25-arc-min disparity range (12.5-arc-min crossed and uncrossed disparities). Thus, half of the dots in the disparity volume were assigned 12.5-arc-min crossed or uncrossed disparity, and the other

half were randomly assigned disparities that fell between the values of the layers (< 12.5 arc min). There were 400 possible disparity values for these points. Each disparity step was 3.75 arc sec. Frame durations were 20 msec, and each frame contained 100 dots/eye. This produced a low-density random-dot stereogram, which was necessary to produce clear percepts of transparency (see, e.g., Akerstrom & Todd, 1988). These patterns appeared as dynamic visual noise: the distribution of dots varied from frame to frame in the x, y plane, but the range of disparities did not change within a given pattern type. Each fixation pattern consisted of 50 frames shown in a fixed order. These frames would be displayed as a continuous loop until an observer depressed a key to begin a trial. At least one entire loop of the fixation sequence was displayed before and after each perturbation sequence.

Apparatus. The patterns were displayed on two Tektronix 608 cathode-ray tube (CRT) monitors with P-31 phosphor and controlled by a PDP-11/73 computer with a direct-memory-access 12-bit digital-to-analog (D/A) interface (Data Translation DT2771). The two CRT displays were haploscopically combined by prisms and mirrors adjustable for each observer to match individual vergence tendencies and interocular separations. These displays were viewed from a distance of 114.6 cm, with the head position constrained by the two prisms mounted in front of each eye. The displays were positioned at eye level. The CRT screens were seen through a circular aperture 10 cm (5°) in diameter, with the outside borders of the monitors occluded by black plastic baffles mounted immediately in front of the monitors.

The CRTs were viewed in a dimly lit room. The background luminance on the CRT screens was approximately 0.056 cd/m². Luminance averaged over a circular area 6 arc min in diameter containing an individual point was approximately 1.8 cd/m² as measured by a Pritchard spot photometer from the observer's eye position. The diameter of an individual point was approximately 0.25 mm (45 arc sec). Corresponding points on the two CRTs were alternately refreshed with an interval of approximately $16 \mu\text{sec}$ between successive pulses from the D/A interface. Thus, the interval between successive refreshes of a given point on one CRT was approximately $16 \mu\text{sec} \times 2 \times 100$, or 3.2 msec.

Task. The method of this experiment was a yes/no discrimination task. The observers fixated a zero-disparity fixation marker embedded in a dynamic random-dot stereogram that was one of the two patterns described above: either two transparent layers or a volume of dots in a circular region surrounded by a zero-disparity ground. This fixation pattern was followed by a sequence of frames (the perturbation sequence) that either contained the complementary pattern (layers or volume) or remained the same as the original fixation pattern. The perturbation sequence was then immediately followed by the original fixation sequence (these conditions are illustrated in Figure 2). Thus, there were both forward and backward masks for all displays. The observers were forced to decide whether or not a change had occurred in the perturbation sequence; 50% of the trials within a given block of trials were of each type. To avoid the possibility that observers were using some cue other than disparity differences as the basis of their judgments, the same random distribution of dots was used for the frame sequence of both the volume and layered patterns within a block of trials. A new random distribution was chosen for each new block of trials (50 trials/block), thus eliminating the possibility that the observers could memorize any given pattern of dots. The only difference between the layered and volume patterns consisted in the disparities assigned to the central figures. This ensured that only the disparity differences between the two patterns could serve as reliable information for the discrimination task. Within a block of trials, the fixation pattern did not change. Two separate blocks of trials were run, one for each fixation pattern.

Sensitivity was assessed over a range of temporal durations and scale factors. The perturbation durations were chosen on the basis of individual performance in pilot work. The strategy was to find

a relatively short perturbation duration such that the observers could perform the discrimination only when they perceived a contraction or an expansion, and a number of longer durations that allowed them to discriminate the two patterns when the depth intervals were perceived to be equivalent. The observers were required to perform this discrimination for two fixation patterns (layers and volume), six scale values, and four (F.N. and L.M.) or five (B.A.) perturbation durations.

Results

The results of this experiment for the 3 observers are shown in Figure 4. Each graph represents the performance of one of the 3 observers, averaged over the two types of fixation patterns (layers and volume). An individual function represents the observer's performance in detecting transitions between the two patterns for a specified duration of the perturbation sequence. Asymmetries in detecting transitions between the two fixation patterns will be described in a later section of this paper. In Figure 4, each data point represents 200 trials for observers L.M. and F.N., and each point for B.A. represents 300 trials. The abscissa represents the scale factor used to rescale the disparities of the layered pattern. This value represents the proportion of the 25-arc-min disparity difference of the unscaled layered pattern. Thus, a scale value of .5 indicates that the disparity difference of the layered pattern was 12.5 arc min. It is important to remember that the layered pattern was the only pattern that was rescaled; there was always a disparity difference of 25 arc min for half of the points in the volume pattern.

As expected, all of the observers' data are V shaped, with the minimum of this function falling within a range of .4-.6 of the scale factor. This means that the apparent collapse in the depth interval produced by placing dots with intermediate disparities between the two layers was about half of the total depth interval—that is, a reduction in perceived depth of about 10-15 arc min of the volume pattern. This is about three times larger than the largest amount of depth averaging reported in previous studies with static random-dot stereograms (Parker & Yang, 1989). The change in the perceived depth of the two patterns was the most salient cue for performing the discriminations. This can be seen by the fact that observers' performance dropped to chance (50% correct) when the depth intervals of the two patterns were perceived to be identical (for short temporal durations), but was nearly perfect when changes in the apparent range in depth were visible (high and low values of the scale factor). This implies that observers are not making their judgments simply on the basis of correlation changes in the patterns, because such changes exist even for scale factors in the .4-.6 range.

As the perturbation duration was increased, the minima of the V-shaped discrimination functions were elevated from chance performance (for Observers B.A. and F.N.; L.M. required further training, as is discussed below). The observers reported that they could begin to see the way in which the dots were distributed in depth in addition to the overall depth interval. The location of the V's minimum did not depend on the duration of the perturba-

tion sequence; it fell consistently within scale factors of .4-.6 for all durations of the perturbation sequence and for all observers. This indicates that the depth averaging observed in the volume pattern was not a transient response of the visual system, as has been observed elsewhere for various forms of disparity interpolation (Mitchison & McKee, 1987a, 1987b). The ability to resolve the structure of the displays did improve with time, however. This finding is consistent with other research documenting an improvement in stereoscopic resolution over durations of about 1,000 msec (McKee et al., 1990; Ogle & Weil, 1958; Tyler, 1991). This will be discussed in greater detail below.

Before turning to the topic of stability, we will examine the cause of depth averaging more carefully. A number of studies revealing depth averaging or depth interpolation have ascribed a primary role to disparity discontinuities in scaling depth, especially during the early stages of stereoscopic processing (Gillam, Chambers, & Russo, 1988; Mitchison & McKee, 1987a, 1987b). It is possible that the depth averaging observed in the volume pattern resulted from a spatially localized estimate of disparity at the boundaries of the figure, and that the central region of the volume inherited the estimates of disparity computed at the discontinuity. In Experiment 2, two forms of noise were added to the perturbation sequences to determine the putative significance of the disparity discontinuities in the layered and volume patterns: unpaired points were distributed either along the boundaries of the central figures (*edge* noise), or within the interior of the form, with no noise along the boundary of the figures (*center* noise). Experiment 2 was performed to document the impact of these two forms of noise on the magnitude of depth averaging observed for the volume pattern. The topic of stability will be discussed in the section that follows this experiment.

EXPERIMENT 2

Disparity Discontinuities and Depth Averaging

Method

The stimuli, apparatus, and task for this experiment were identical to those in Experiment 1, except for the presence of spatially localized noise points in the perturbation sequence. In this experiment, *noise* refers to points that were not assigned a match in the other eye's view. Two kinds of such noise were introduced into the perturbation sequences: edge noise was concentrated at the location of the disparity discontinuities, and center noise was concentrated in the interior of the circular figure. The total number of noise points was statistically identical for both noise distributions. For Observers B.A. and F.N., the noise points accounted for one third of the dots in the central figure; for L.N. (the least sensitive observer), this proportion was half this value, or one sixth. The density of the noise was varied over the central figure in one of two ways, corresponding to the edge and center noise conditions. For the edge noise condition, the density of the noise was maximal at the cyclopean edges of the figure and decreased monotonically toward the center of the circle. The density of the noise was varied monotonically (rather than as a step function) to prevent the formation of a new boundary or edge. For the center noise

Depth Averaging Data

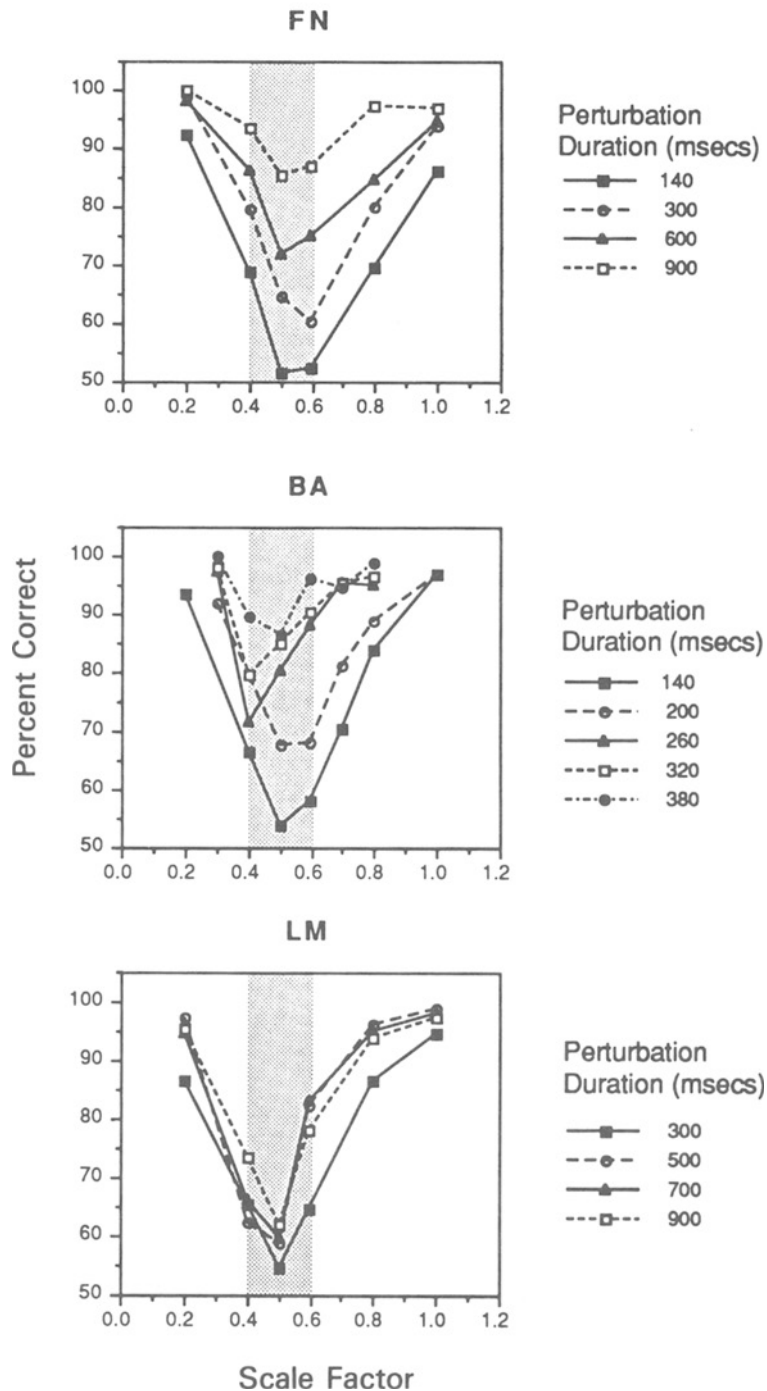


Figure 4. Results of Experiment 1 for 3 observers. Each graph represents the performance of an individual observer, and each function corresponds to the discrimination accuracy of an observer for a given duration of the perturbation sequence. These durations were varied to match individual sensitivities. Note that the minimums of the V-shaped discrimination functions fell within .4-.6 of the scale factor for all perturbation durations and all observers. This amounts to 10-15 arc min of disparity averaging in the volume pattern.

condition, the density of the noise was maximal in the center of the figure; it decreased monotonically toward the edges of the circular region.

To create a pattern with noise concentrated at the edge of the figure, each dot was passed through a *probability filter*, in such a way that the likelihood of a dot's being a noise dot increased as the square of the distance from the center of the circle. The total radius of the central circular figure was normalized so that the distance from the center (D_{Center}) of the circle was scaled on the interval [0,1], with 0 being the center of the circular region, and 1 the boundary of the circle. In the convention used here, the probability P_{Noise} refers to the likelihood of a point's being designated a noise point, which is given by $P_{\text{Noise}} = (D_{\text{Center}})^2$. The probability filter for edge noise was designed so that a point on the boundary of the circular region would have a 100% chance of being a noise point. A similar strategy was used for creating center noise, but the probability of being designated as noise was $P_{\text{Noise}} = 1 - (D_{\text{Center}})^2$. This strategy ensured that the total number of noise points was statistically identical for the edge and center noise conditions. This was confirmed computationally by comparing the number of points computed as noise for both distributions on a number of trial runs.

Once a dot was designated as a noise dot, the dot that matched it in the complementary eye was randomly repositioned to ensure that no density differences were introduced (no changes were perceivable monocularly), as well as to ensure that the likelihood of a match for the noise point in the other eye was relatively small.

The task in this experiment was identical to that in Experiment 1. The observers were required to detect transitions (yes/no) between the layered and the volume patterns in the presence of the two types of noise. There were two noise conditions (edge and center), two fixation conditions (layered and volume), two (F.N. and L.M.) or four (B.A.) perturbation durations, and six (L.M.), seven (F.N.), or nine (B.A.) scale factors. One block of trials consisted of a single fixation type, a single noise type, and one scale factor.

Results

The results of this experiment are shown in Figure 5. As in Experiment 1, each graph represents the performance of 1 of the 3 observers, averaged over the two types of fixation patterns (layers and volume), with one function representing the observer's performance in detecting transitions between the two patterns for a specified duration of the perturbation sequence. Note that the magnitude of the depth averaging observed in the volume pattern does not appear to be affected by the spatial location of the incoherent noise. As in Experiment 1, the minima of the V-shaped discrimination functions were in the .4-.6 range of the scale factor. Thus, the same extent of depth reduction was observed at both the boundaries and the interior of the volume figure for all observers.

DISCUSSION

The Temporal Stability of Depth Averaging

Both cooperative network interactions and the coarse-to-fine model described by Watt (1987) predict that disparity interactions, like depth averaging or depth attraction, should vary as a function of stimulus duration. Watt's prediction is straightforward: averaging should decrease as stimulus exposure increases. The thesis of cooperativity may be interpreted in a number of different ways. If cooperativity is associated with a mechanism that main-

tains the stability of depth computations (i.e., a persistence mechanism), we would expect that the amount of depth averaging exhibited in the layered to volume transition should be less than that for the volume to layered transition. The logic of this prediction follows from the conceptualization of cooperativity as a form of feedback that maintains the stability of a depth computation. Thus, if the recent history of stimulation generates a representation of a small depth interval (e.g., the volume), then, for short durations, cooperative interactions should act to maintain the stability of this depth representation, making it more difficult to detect expansions. This would be reflected in the psychometric functions as a shift in the relative positions of the minima for the V-shaped discrimination functions, such that the volume to layered transition should have a minimum shifted toward smaller scale factors than those for the layered to volume transitions.

In Figure 6, we have replotted the data in Experiment 1 for three exposure durations, showing the discrimination functions of the layered to volume transition and the volume to layered transition separately for the 2 observers who exhibited temporal improvements in sensitivity (B.A. and F.N.). Note that there was no difference in the magnitude of depth averaging for the two transitions for any temporal duration: the minima of the discrimination functions were the same for the layered to volume and the volume to layered transitions. In other words, the expansion/contraction transformation was perfectly *reversible*; there was no perceptual "advantage" to either the pattern exhibiting disparity averaging (the volume) or the layered pattern in terms of the stability of their overall depth interval. This reversibility weakens the putative role of cooperative feedback as a persistence mechanism, at least with respect to the form of disparity averaging observed here.

In addition to the reversibility of depth averaging, the overall magnitude of depth averaging remained essentially invariant with increases in exposure duration. This can be seen in Figure 4. Note that the minima of the V-shaped discrimination functions occur at roughly the same scale factors for all exposure durations (.4-.6). Watt's (1987) model attributes the improvement in spatial discrimination tasks to the sequential deactivation of spatial filters in a coarse-to-fine fashion. A model of this kind would predict that the extent of spatial averaging should be reduced for longer display durations, a prediction not supported by the present data. Note that perceptual improvements *are* occurring over these time intervals, despite the fact that the amount of averaging remains virtually constant. If the stereo system is utilizing a coarse-to-fine strategy, it does not seem to deactivate the coarse filters with time, but rather just turns the fine mechanisms on. Thus, we have found no evidence to support the hypothesis that cooperative feedback causes estimates of relative depth to persist.⁴

Factors Influencing Disparity Averaging

A number of recent studies have documented disparity averaging with static (Parker & Yang, 1989) and dynamic

Depth Averaging with Spatially Localized Noise

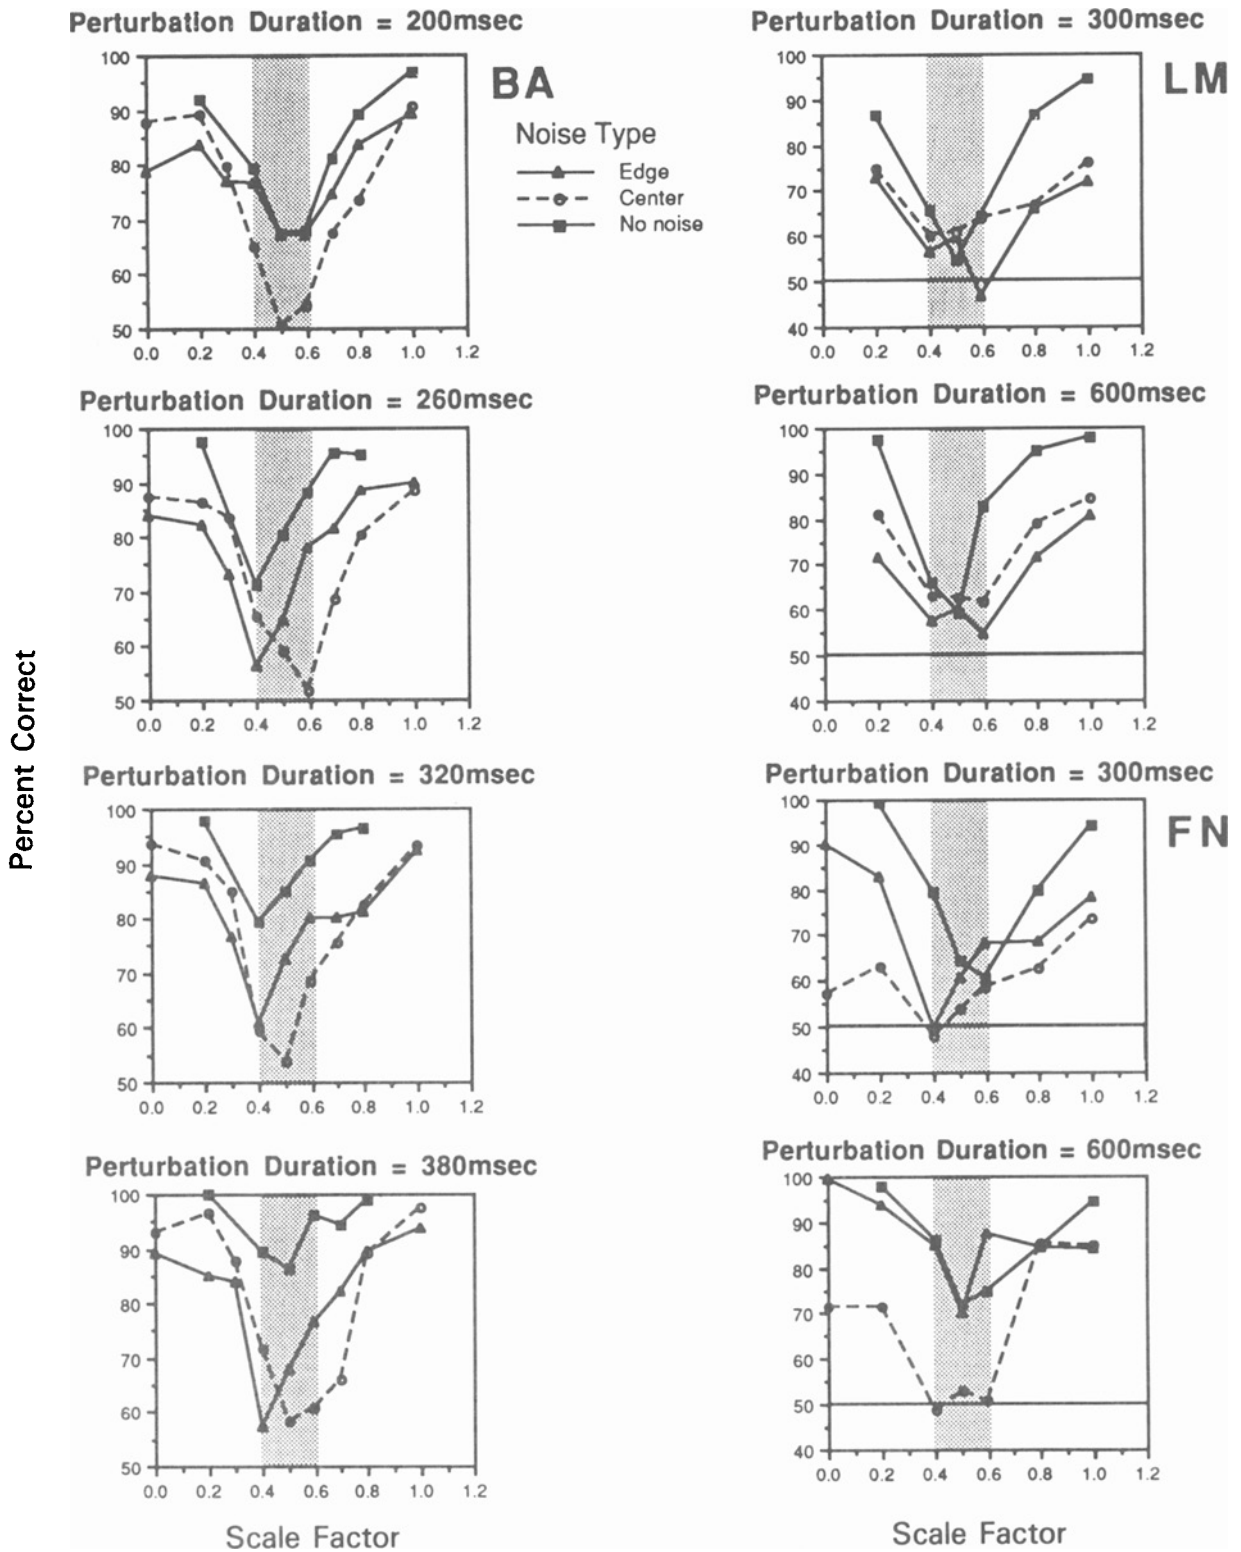


Figure 5. Results of Experiment 2 for Observers B.A., L.M., and F.N. Each graph plots the performance of 1 observer for a given perturbation duration and for the two types of noise (edge and center; see text for details). The patterns used in Experiment 1 are plotted for comparison. Note that the amount of depth averaging is roughly identical (.4-.6) for both types of noise and for the noiseless patterns, suggesting that depth is represented explicitly throughout the displays and not just at the disparity discontinuities.

Reversibility of Depth Averaging

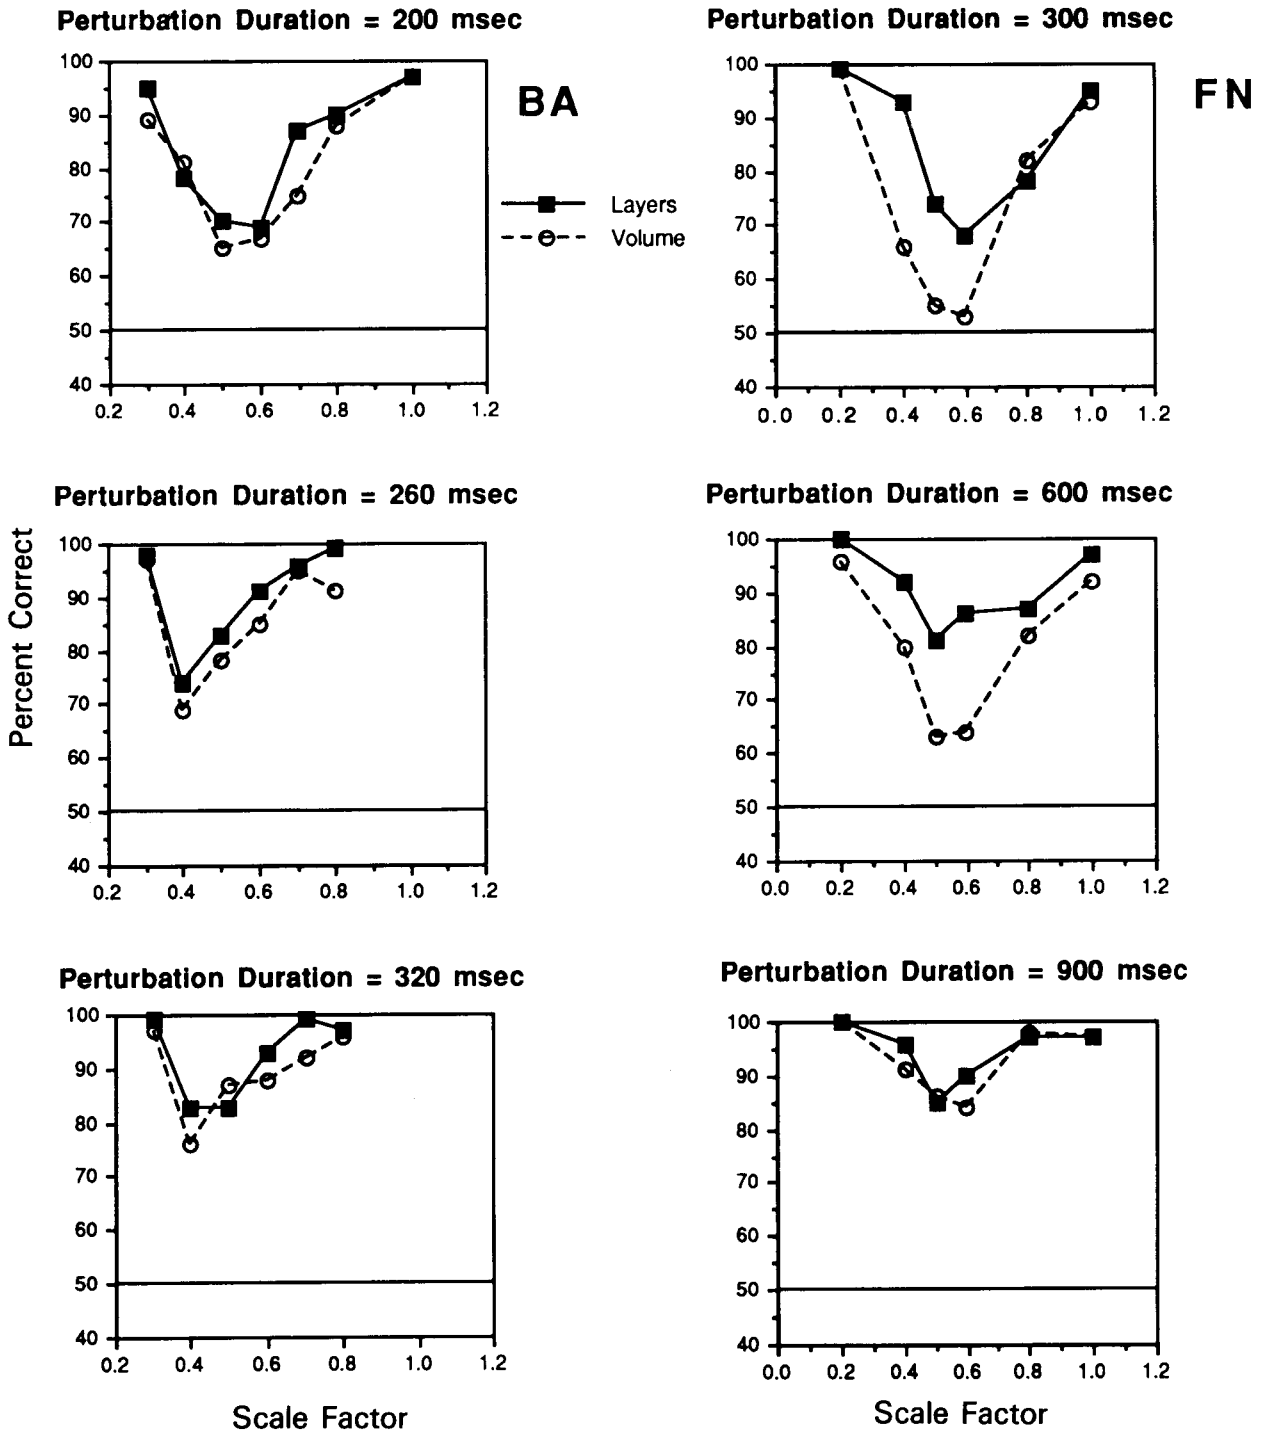


Figure 6. Data demonstrating the reversibility of the expansion/contraction cue for the 2 observers who showed temporal improvements in sensitivity. Note that the minima of the two functions occur at the same scale factors. This implies that the amount of depth averaging was the same for the two fixation patterns.

(Stevenson et al., 1991) random-dot stereograms, but the magnitude of averaging reported here is three to four times larger than that reported previously. It is not clear what caused this difference. The attractions reported by Stevenson et al. were also performed with dynamic random-dot stereograms with comparable temporal parameters; but the density of their patterns was much higher, and the range of disparities was much smaller than those reported here. To determine the critical variables leading to the differences in our results, we would have to perform these experiments over a range of densities and disparity values. Equipment limitations prevented us from performing a study of that kind. We therefore focused our energies on the role of disparity discontinuities in depth averaging.

The results of Experiment 2 suggest that relative depth is explicitly represented throughout the central figure, not simply at the location of the disparity discontinuities. The patterns used in these studies were rather sparse, however. Spatial and/or temporal averaging was required to make the pattern perceptually transparent. This meant that there were some disparity discontinuities that occurred locally (spatially and temporally) throughout the layered figure. The expression *disparity discontinuities* in this paper therefore refers to the positions of discontinuities that are spatially and temporally stable over the duration of the stimulus. Indeed, it is possible that the failure to find clear differences between the edge and center noise conditions was due to the transient discontinuities that were present throughout the layered pattern. However, given that the edges of the central figure were by far the strongest discontinuities in the image, we would expect some tendency for shifts in the magnitude of spatial averaging if discontinuities played a crucial role in this phenomenon. We found no systematic shifts in the amount of depth averaging for either type of noise. We therefore conclude that the disparity discontinuities present in these patterns did not play a significant role in the depth averaging results found in Experiments 1 and 2.

The Temporal Stability of Depth Organizations

In this section, we will consider the relative stability of the layered and volume depth *organizations*. The phenomenon of depth averaging complicates this comparison. Within the context of current computational models, a difference in relative depth indicates that different disparity-sensitive mechanisms were activated by the two patterns. By using the scale factor method described above, we were able to determine what disparity values generated equivalent percepts of relative depth for the two patterns. This ensured that differences in the perceived organization of the dots in depth was the only source of information available to perform the discriminations (at these scale factors). Theoretically, this implies that there were minimal differences in the disparity detectors excited by the two patterns (under the assumption that relative depth is coded by the disparity detectors that are active).

Hysteresis Data

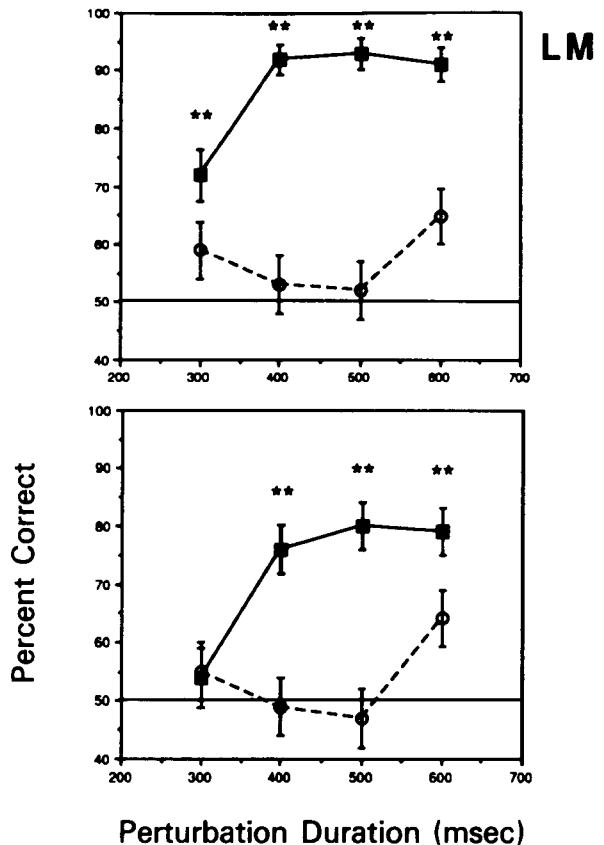
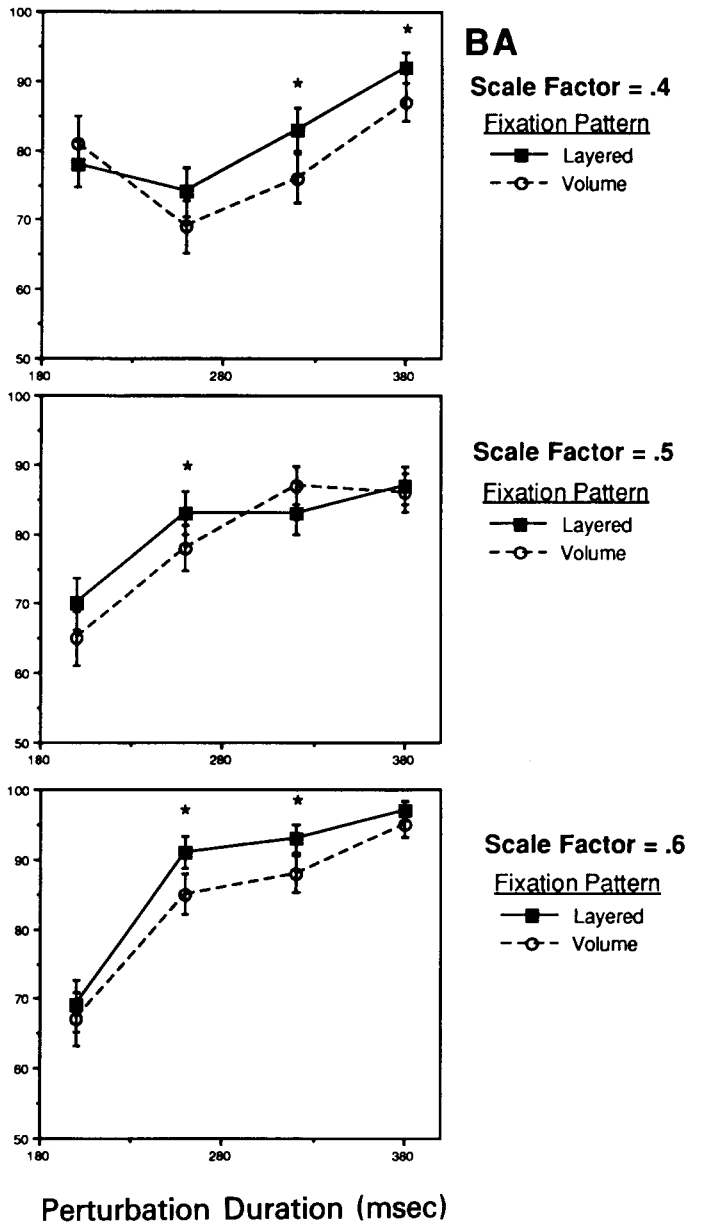
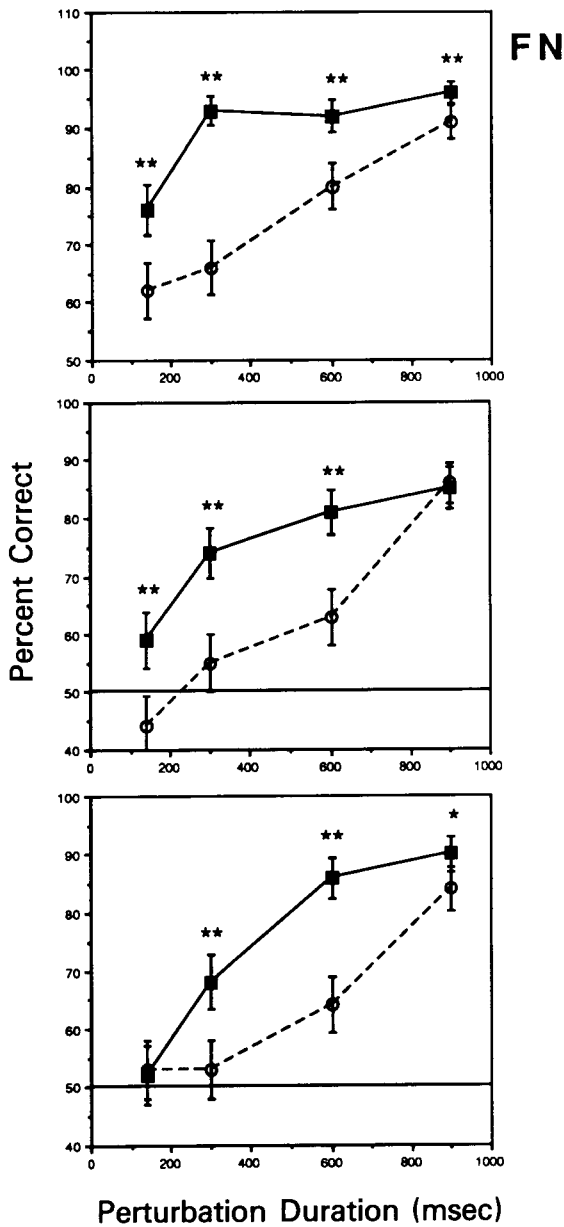


Figure 7 (Above and next page). Hysteresis data for L.M., F.N., and B.A., for values of the scale factor that led to equivalent percepts of depth. Each graph represents the performance of 1 observer for a given value of the scale factor plotted as a function of the perturbation duration. The two functions represent the two fixation patterns used in the experiments. *Significant difference of $p < .05$. **Significant difference of $p < .01$. Observers typically find it easier to detect perturbations of the layered pattern than of the volume pattern. (Scale factors = .4 (top) and .5. Fixation pattern legends given on next page.)

In Figure 7, the data collected in Experiment 1 have been replotted, showing accuracy as a function of perturbation duration for scale factors near the minimum of the discrimination functions (.4-.6)—that is, scale factors generating equivalent percepts of overall depth. The two functions on each graph correspond to the layered and volume fixation patterns. Error bars represent standard deviations computed as $[p(p-1)/n]^{1/2}$, where p represents the proportion correct and n the number of trials. If observers could detect both perturbations with equal ease, these functions should be superimposed on each other. In general, this was not the case. The single asterisk in the graphs represents a z score significance level of $p < .05$, and the double asterisks indicate a significance level of $p < .01$. The data reveal that observers found it more difficult to detect volume to layered transitions

Hysteresis Data

Hysteresis Data



than layered to volume transitions, as expected from the asymmetry between volume to plane and plane to volume transitions described by Julesz and Tyler (1976) (the correlated sequences of Julesz and Tyler's studies were perceived as a plane, and the decorrelated sequence appeared as a lustrous volume). Thus, unlike the reversibility of depth averaging, there is evidence of an asymmetry (hysteresis) in detecting perturbations between the layered and volume patterns: it was harder to detect transitions from the volume pattern to the layered pattern than vice versa. There were individual differences in the magnitude of hysteresis observed for the 3 observers, however. The most probable explanation for these differences was that there

were variations in the stereoscopic experience of the observers. L.M. was an inexperienced viewer of stereograms of any kind; indeed, L.M. did not initially demonstrate the ability to perform the discrimination task at these scale factors for any of the perturbation durations tested. It seemed likely that this difference was due to her relative inexperience as a stereoscopic observer. If so, her poor performance could be remedied with further training. Because her data revealed well-localized minima at scale factors of .4 and .5 (see Figure 4), she restricted her practice to these two values. The data for Subject L.M. in Figure 7 therefore represent a retest of her sensitivity for these scale values following further train-

ing. L.M. showed the largest hysteresis of all observers tested; she never could reliably detect perturbations of the volume pattern at these scale factors, even after prodigious amounts of practice. Observer F.N. also exhibited a rather large hysteresis, but it was not as large as L.M.'s. F.N. was an experienced stereoscopic observer, but most of this experience involved static random-dot stereograms. B.A. (the author), who exhibited the smallest hysteresis of all observers, was the most experienced viewer of dynamic random-dot stereograms and of these patterns in particular. The difference in experience is reflected in the performances for the 3 observers in discriminating the two patterns. Observer B.A. had the best discrimination performance for all values of the perturbation sequence (note that he reached ceiling performance for much shorter perturbation durations and all scale values); F.N. had the second best performance; and L.M. performed the worst. Thus, differences in the perceptual experience of observers was reflected by their relative skill in discriminating the patterns, which was in turn reflected in the magnitude of hysteresis exhibited by the respective observers.

This suggests that the comparatively small hysteresis observed in B.A.'s data may be due to the insensitivity of the perturbation method in distinguishing psychometric functions that rise so rapidly; B.A. reached ceiling performance very rapidly, effectively masking differences between these functions. If this was true, B.A. would reveal a larger hysteresis if the task was made more difficult. The addition of noise points in Experiment 2 provided a good test of this hypothesis; this noise led to substantial decrements in the gain of B.A.'s performance function with perturbation duration. As with the hysteresis data of Experiment 1, the same scale factors (.4-.6) were used to assess hysteresis. Recall that the amount of depth averaging did not depend on the spatial distribution of noise, so that these values represent the amount of scaling needed to remove an expansion or contraction cue. As shown in Figure 8, the kind of asymmetry that was observed in Experiment 1 was observed for both forms of noise in Experiment 2: it was typically more difficult to detect perturbations of the volume fixation than it was to detect perturbations of the layered pattern. The magnitude of the hysteresis is somewhat larger in the edge noise condition than it is with center noise, but this was probably because there seems to be a floor effect limiting performance in the center noise condition. Note that all of the signal dots were closer to the fixation point when noise was distributed along the edge, whereas all of the signal dots in the center noise condition were about 1° away from the fixation point. The floor effect operating in the center noise condition may therefore be due to the high density of noise that surrounds the fixation point in the center noise condition. Regardless of the cause of the floor effect observed with center noise, it is important to underscore the fact that hysteresis effects can only be observed if the discrimination functions are not limited by ceiling or floor effects. All observers revealed hysteresis effects in the same direction when the psychometric func-

tions rose relatively slowly over the temporal durations tested. More critically, no observer found it easier to perform the volume to layered discrimination, the prediction that follows from cooperative feedback.

Hysteresis and Cooperativity

The concept of cooperativity emerged as one of the most influential theoretical constructs to explain the visual system's ability to disambiguate the types of surfaces represented in random-dot stereograms. Historically, the phenomenon of hysteresis has played a prominent role in justifying the thesis of cooperativity. The fusional hysteresis initially observed by Fender and Julesz (1967) seemed to provide evidence that the fusional limit had expanded when the stereo pair was initially fused. The observed hysteresis was believed to arise from the stability of a binocularly fused state, such that once fusion was obtained, cooperative interactions would act to prevent the fused state from being broken. This finding was promoted as one of the strongest pieces of evidence that stereopsis should be considered a cooperative process (see Julesz, 1971, and Julesz's commentary on Grossberg).

Replications of Fender and Julesz's (1967) study have also revealed hysteresis, although the point of refusion has generally been much larger than 6 arc min (Diner, 1978; Diner & Fender, 1987; Erkelens, 1988; Hyson, Julesz, & Fender, 1983; Piantanida, 1986). There has therefore been a history of controversy surrounding the observation of fusional hysteresis, but surprisingly, its putative relationship with cooperativity has emerged unquestioned. Indeed, hysteresis effects continue to be interpreted as evidence for cooperative interactions, and they are often the sole rationalization for cooperative interactions (Williams & Phillips, 1987; Williams, Phillips, & Sekuler, 1986). The conceptual link between cooperativity and hysteresis was shaped by the belief that hysteresis effects result from the persistent stability of *ordered* perceptual organizations. In the example of fusional hysteresis described above, the ordered percept corresponds to the binocularly fused state. Therefore, it was not sufficient to simply demonstrate hysteresis to sustain the thesis of cooperativity; rather, it had to be demonstrated that hysteresis resulted from the persistence of an ordered organization.

Recently, Erkelens (1988) performed an experiment that questions this interpretation of fusional hysteresis. Erkelens measured three fusional limits: the limit of diplopia for an initially fused random-dot stereogram, where the disparity is slowly increased; the limit for fusion when the random-dot stereogram half images are initially diplopic; and the static limit for fusion of a large random-dot stereogram with no previous history. The interesting result of this procedure was that the static limit of fusion was slightly larger than the limit in either of the conditions in which disparity was manipulated in real time. Rather than providing evidence that the fusional limit was expanded for an initially fused stereo pair, Erkelens demonstrated that the fusional zone actually contracts when

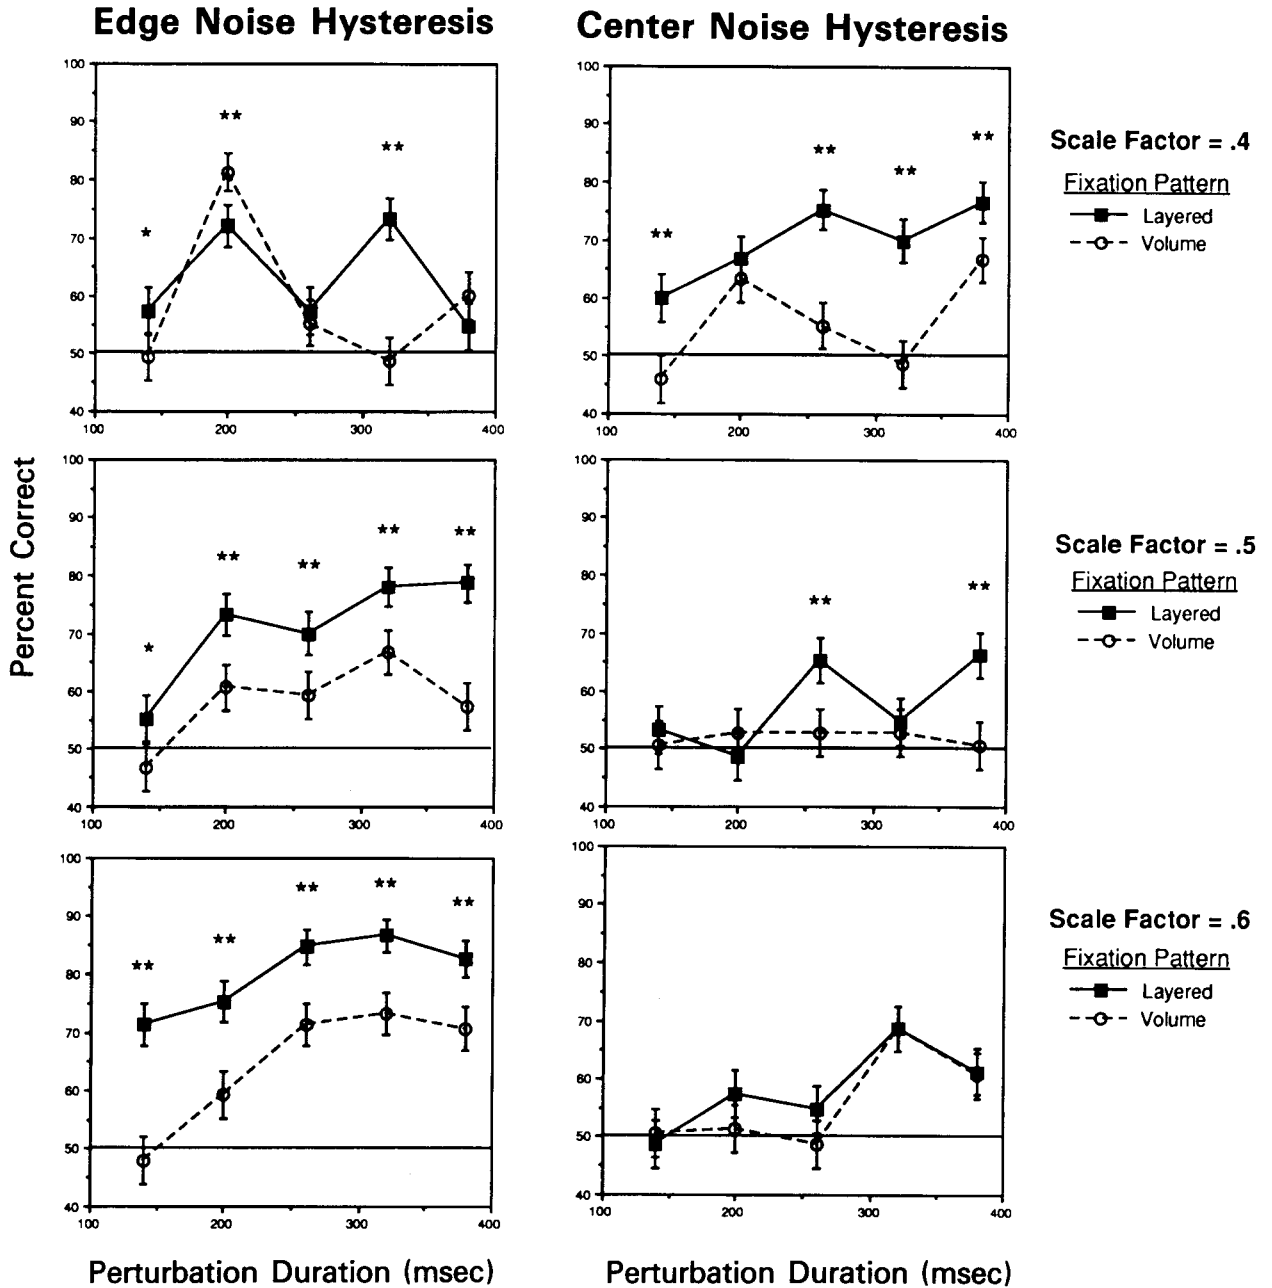


Figure 8. Hysteresis data of Observer B.A. for the two types of noise in Experiment 2. A number of the graphs reveal the presence of a floor effect limiting performance, especially in the center noise condition. However, when these floor effects are not present, the data reveal the same pattern of hysteresis as was observed in Experiment 1.

preceded by a state of diplopia. Erkelens argued that the state of diplopia interfered with the process of regaining fusion, causing a temporal contraction of the fusional zone. This interference was attributed to the binocular rivalry generated by the diplopic images (Erkelens, 1988).

Erkelens's interpretation of fusional hysteresis is consistent with research performed by Julesz and Tyler (1976) and Tyler and Julesz (1976, 1978), mentioned briefly in the introduction of this paper. In these studies, dynamic

random-dot stereograms were generated that were identical (100% correlation) or uncorrelated (0% correlation), or were complements of each other (-100% correlation). Duration thresholds were measured for detecting transitions between these states. They found that it generally took much longer to detect transitions from a disordered state (uncorrelated random-dot stereogram) to an ordered state (correlated random-dot stereogram) than vice versa. This, too, is a form of hysteresis, an irreversibility or

“lag” in the transition between different states. Yet it is interesting that this irreversibility is in the opposite direction of that predicted by a cooperative interpretation of hysteresis. This interpretation of cooperativity would predict that the more ordered state should show a resistance to change; once formed, the coherent state (the correlated random-dot stereogram) should persist. In fact, the opposite form of hysteresis was observed. Similar effects were mentioned by Chang and Julesz (1984) for random-dot cinematograms in an ambiguous motion display. Here too, it was found that it was more difficult for observers to detect a transition from a disordered state to an ordered state than conversely. Julesz and Tyler dubbed the correlation/decorrelation asymmetry *catabolism*, for it seemed to be a neural analogue reflecting the physical principle of entropy increase that occurs in (catabolic) biological processes.

Prior to the experiments of Erkelens (1988), the hysteresis observed for fusion and binocular correlation could not be interpreted within a single framework. Fusional hysteresis seemed compatible with the idea that ordered states resist change, whereas the extreme sensitivity to perturbations in correlated random-dot stereograms suggested that ordered states were highly sensitive to change. The hysteresis effects observed in the present studies also suggest that the ordered state is more sensitive to change. These discrepant findings may now be reconciled. Both studies reveal that it is more difficult to effect transitions from a disordered binocular state to an ordered state than it is to effect transitions from an ordered state to a disordered state, as originally proposed by Julesz and Tyler (1976). This suggests that the link between cooperativity and hysteresis is not as direct as has been previously thought (Julesz, 1971; Williams & Phillips, 1987; Williams et al., 1986). Indeed, the idea that hysteresis be considered a signature of cooperative systems was sustained by the belief that hysteresis effects reflected neural interactions that led to the persistence of organized percepts. Given the experiments that have been performed to date, it seems fair to conclude that hysteresis cannot be attributed to cooperativity, at least in the sense of persistence mechanisms. As in the experiments of Erkelens (1988), Julesz and Tyler (1976), and Tyler and Julesz (1976, 1978), the data presented here suggest that hysteresis effects can be observed in disorder-order transitions, but the state most resistant to change is the relatively disordered state.

The consistency of hysteresis effects described to date suggests that a single principle may connect these phenomena. Recently, Tyler (1991) has offered a theory that would link the hysteresis observed in detecting correlation changes with fusional hysteresis. Tyler proposed that fusional hysteresis, like the neurontropy asymmetry, could be understood (at least in part) as the consequence of a suppression of fusional mechanisms generated by the state of rivalry. Erkelens (1988) also rationalized fusional hysteresis with the putative suppression of fusional mechanisms engaged by the state of diplopia. A similar explanation of the neurontropy studies of Tyler and Julesz was

offered by Blake (1989), who suggested that decorrelated stimuli generate interocular suppression that slows the transition to a fused binocular state. Thus, there would seem to be agreement on the cause of the hysteresis effects observed for binocular fusion and changes in binocular correlation. Surprisingly, however, the thesis of cooperativity does not play a role in any of these models.

The results of the experiments described here are consistent with the other forms of hysteresis observed in binocular vision. Tyler (1991) has suggested that the neurontropy asymmetry might be unique to the stereoscopic percept of lustre, because no similar effect was previously found with a comparable stereoscopic task. The experiments presented here suggest otherwise; similar asymmetries were found, and the displays did not generate percepts of lustre. It therefore seems reasonable to conclude that the correlation/decorrelation hysteresis is not a unique property of lustre. Rather, it appears that hysteresis effects are a general property of binocular vision, and that they arise when there are sufficiently large differences in the degree of order in the binocular states. We are currently conducting research to determine whether similar effects can account for the hysteresis effects observed in motion perception which have been ascribed to cooperative persistence mechanisms (Williams & Phillips, 1987; Williams et al., 1986).

Physical Theories of Hysteresis

Objections have been raised that the data described here do not directly apply to the thesis of hysteresis because the task does not conform to the usual hysteresis paradigm.⁵ For example, consider fusional hysteresis. In the usual hysteresis paradigm, observers have to report the point at which a fused stimulus becomes diplopic, or a diplopic stimulus becomes fused. The idea is that “fusion” and “diplopia” constitute stable, mutually exclusive states that the system resists flipping between. In contrast, the experiments reported here suggest a resistance to moving up, rather than down, a continuum of relative order. Therefore, it has been suggested that these results should not be used to assess the relationship between hysteresis and cooperativity.

This criticism is a useful point of departure for a discussion of the relationship between hysteresis and cooperativity in physical systems. One of the central points of this paper is to question the usual hysteresis paradigm, in terms of both the prevailing methodology and the theoretical interpretation applied to the results. First, it is important to point out that the phenomenon of hysteresis that motivated Julesz’s thesis of cooperativity was that observed in the magnetization curves of ferromagnetic materials. Indeed, this metaphor played such a strong role in guiding his thinking that it led to the development of the spring-coupled magnetic-dipole model of stereopsis (Julesz, 1971; Julesz & Chang, 1976), one of the first cooperative models of stereopsis. It is therefore instructive to review the physical theory of this phenomenon to gain insight into the cause of hysteresis in ferromagnetic materials.

A ferromagnetic substance is a material that becomes magnetized when placed in a magnetic field. This magnetization takes time, but eventually the material becomes aligned with the external magnetic field. Imagine the following experiment. Start with a ferromagnet that is magnetized in a particular orientation. Shift the external magnetic field to the opposite orientation, and plot the magnetization curve. When the ferromagnet has asymptotically reached the orientation of the external field, reverse the external field back to the original orientation of the material and plot that curve. These curves are not the same; the transformation between the two orientations is irreversible. This hysteresis is attributed to a dissipation of energy created when individual magnetic domains get "hung up" on impurities in the material, causing the magnetic domains to "snap" over the impurities. This "snapping" dissipates energy in two ways—by generating eddy currents that heat the material, and by generating sound waves (see Feynman, 1964). Irreversibility is therefore attributed to the energy loss (entropy) generated by these processes, not to cooperative interactions. We contend that although the experiments reported here do not conform to the usual strategies used to document hysteresis, they are nonetheless consistent with the physical processes that putatively justify the importance of hysteresis studies in the first place. This leads us to question the putative link between hysteresis and cooperativity on theoretical grounds. We conclude with the observation that hysteresis effects have been understood to arise physically as the consequence of processes that degrade energy, not processes that maintain "ordered" states (i.e., cooperativity). We suggest that the same processes allow for an understanding of the relevant psychophysical data reported to date. This does not imply that the concept of cooperativity should be abandoned as a theoretical construct, but rather, that cooperativity should not be interpreted as a persistence mechanism that maintains the stability of organized (stereoscopic) percepts.

REFERENCES

- AKERSTROM, R. A., & TODD, J. T. (1988). The perception of stereoscopic transparency. *Perception & Psychophysics*, **44**, 421-432.
- BLAKE, R. (1989). A neural theory of binocular rivalry. *Psychological Review*, **96**, 145-167.
- CARERI, G. (1984). *Order and disorder in matter*. Menlo Park, CA: Benjamin/Cummings.
- CHANG, J., & JULESZ, B. (1984). Cooperative phenomena in apparent motion perception of random-dot cinematograms. *Vision Research*, **24**, 1781-1788.
- DINER, D. B. (1978). *Hysteresis in human binocular fusion: A second look*. Unpublished doctoral dissertation, California Institute of Technology, Pasadena.
- DINER, D. B., & FENDER, D. H. (1987). Hysteresis in human binocular fusion: Temporalward and nasalward ranges. *Journal of the Optical Society of America*, **4**, 1814-1819.
- DUWAER, A. L., & VAN DEN BRINK, G. (1981). What is the diplopia threshold? *Perception & Psychophysics*, **29**, 295-309.
- ERKELENS, C. (1988). Fusional limits for a large random-dot stereogram. *Vision Research*, **28**, 345-353.
- FENDER, D., & JULESZ, B. (1967). Extension of Panum's fusional area in binocularly stabilized vision. *Journal of the Optical Society of America*, **57**, 819-830.
- FEYNMAN, R. P. (1964). Ferromagnetism. In R. P. Feynman, R. B. Leighton, & M. Sands (Eds.), *The Feynman lectures on physics* (Vol. 2, pp. 36-1-36-15). Menlo Park, CA: Addison-Wesley.
- GILLAM, B., CHAMBERS, D., & RUSSO, T. (1988). Postfusional latency in stereoscopic slant perception and the primitives of stereopsis. *Journal of Experimental Psychology*, **14**, 163-175.
- GROSSBERG, S. (1987). Cortical dynamics of three-dimensional form, color, and brightness perception: II. Binocular theory. *Perception & Psychophysics*, **41**, 117-158.
- HYSON, M., JULESZ, B., & FENDER, D. H. (1983). Eye movements and neural remapping during fusion of misaligned random-dot stereograms. *Journal of the Optical Society of America*, **73**, 1665-1673.
- JULESZ, B. (1971). *Foundations of cyclopean perception*. Chicago: University of Chicago Press.
- JULESZ, B. (1983). Commentary on Grossberg, S., The quantized geometry of visual space: The coherent computation of depth, form and lightness. *Behavioral & Brain Sciences*, **6**, 625-692.
- JULESZ, B., & CHANG, J. (1976). Interaction between pools of binocular disparity detectors tuned to different disparities. *Biological Cybernetics*, **22**, 107-119.
- JULESZ, B., & TYLER, C. (1976). Neuronropy, an entropy-like measure of neural correlation, in binocular fusion and rivalry. *Biological Cybernetics*, **23**, 25-32.
- MARR, D., & POGGIO, T. (1976). Cooperative computation of stereo disparity. *Science*, **194**, 283-287.
- MARR, D., & POGGIO, T. (1979). A computational theory of human stereo vision. *Proceedings of the Royal Society of London: Series B*, **204**, 301-328.
- McKEE, S. P., LEVI, D. M., & BOWNE, S. F. (1990). The imprecision of stereo judgments. *Vision Research*, **30**, 1763-1780.
- MITCHISON, G. J., & McKEE, S. P. (1987a). Interpolation and the detection of fine structure in stereoscopic matching. *Vision Research*, **27**, 295-302.
- MITCHISON, G. J., & McKEE, S. P. (1987b). The resolution of ambiguous stereoscopic matches by interpolation. *Vision Research*, **27**, 285-294.
- MITCHISON, G. J., & WESTHEIMER, G. (1984). The perception of depth in simple figures. *Vision Research*, **24**, 1063-1073.
- OGLE, H. N., & WEIL, M. P. (1958). Stereoscopic vision and the duration of the stimulus. *Archives of Ophthalmology*, **48**, 4-12.
- PARKER, A. J., & YANG, Y. (1989). Spatial properties of disparity pooling in human stereo vision. *Vision Research*, **29**, 1525-1538.
- PIANTANIDA, T. P. (1986). Stereo hysteresis revisited. *Vision Research*, **26**, 431-437.
- POLLARD, S. B., MAYHEW, J. E. W., & FRISBY, J. P. (1985). A stereo correspondence algorithm using a disparity gradient limit. *Perception*, **14**, 449-470.
- PRAZDNY, K. (1985). Detection of binocular disparities. *Biological Cybernetics*, **52**, 93-99.
- SCHUMER, R. A. (1979). *Mechanisms in human stereopsis*. Unpublished doctoral dissertation, Stanford University.
- STEVENSON, S. B., CORMACK, L. K., & SCHOR, C. M. (1991). Depth attraction and repulsion in random-dot stereograms. *Vision Research*, **31**, 805-813.
- TYLER, C. W. (1991). Cyclopean vision. In J. H. Macmillan (Ed.), *Vision and visual dysfunction* (Vol. 9). New York: Cronly-Dillon.
- TYLER, C. W., & JULESZ, B. (1976). The neural transfer characteristic (neuronropy) for binocular stochastic stimulation. *Biological Cybernetics*, **23**, 33-37.
- TYLER, C. W., & JULESZ, B. (1978). Binocular cross-correlation in time and space. *Vision Research*, **18**, 101-105.
- WATT, R. J. (1987). Scanning from coarse to fine spatial scales in the human visual system after the onset of a stimulus. *Journal of the Optical Society of America A*, **4**, 2007-2021.
- WESTHEIMER, G. (1986). Panum's phenomenon and the confluence of signals from the two eyes in stereopsis. *Proceedings of the Royal Society of London*, **228**, 289-305.

- WESTHEIMER, G., & LEVI, D. (1987). Depth attraction and repulsion of disparate foveal stimuli. *Vision Research*, 27, 1361-1368.
- WILLIAMS, D. W., & PHILLIPS, G. (1987). Cooperative phenomena in the perception of motion direction. *Journal of the Optical Society of America A*, 4, 878-885.
- WILLIAMS, D. W., PHILLIPS, G., & SEKULER, R. (1986). Hysteresis in the perception of motion direction as evidence for neural cooperativity. *Nature*, 324, 253-254.

NOTES

1. It may be argued that early cooperative models are not stable under the conditions of transparency, so there is no reason to believe that a disparity mixture of one kind is more stable than another. The failure of Marr and Poggio's cooperative model in representing transparency arises from the inhibitory connections between disparity mechanisms tuned to different disparities along the same line of sight. This problem led Prazdny (1985) to develop a "purely" cooperative model, wherein only the facilitory connections between similarly tuned mechanisms are retained. Grossberg's model (1987) also included cooperative feedback that maintains the stability of "boundary contours." Thus, the prediction of differential stability between different types of disparity mixtures arises from cooperative models that are stable under transparency, not from the model of Marr and Poggio (1976).

2. The most widely used measure of order in physical systems was given by the entropy function first described by Boltzmann. This function was later generalized as a measure of uncertainty by Shannon, who used it to measure quantities such as channel capacity in a communication system. This function may be written

$$S = k \cdot \ln(P),$$

where P is the number of ways to achieve a given state. If the different states available to the system are equally probable (*state* refers here to the disparity of a given dot), it follows that P is proportional to the number of possible states available to the system (Careri, 1984). In the present study, the volume pattern had "access" to many more possible disparity values than did the layered pattern, and it therefore had a higher degree of "uncertainty" or "disorder." Any given dot in the layered pattern could have only two possible disparity values. In contrast, the volume pattern selected disparities from over 400 possible values. Because S is a monotonically increasing function of P , it follows that the volume pattern had a larger entropy than the layered pattern did; that is, it was more disordered. We wish to emphasize that only an ordinal difference in relative order of these patterns was needed to describe them as "ordered" and "disordered," respectively (see Careri, 1984, for a general discussion of order and disorder in physical systems).

3. Only the smallest scale factor used (.2) placed the layered pattern within the range of repulsive interactions reported by Stevenson et al. (1991). These interactions appeared to have little or no impact on perceived depth for the densities used in these experiments. A repulsion should make it more difficult to detect changes between the two patterns for small-scale factors: while the actual disparities of the layered pattern would be reduced by the scale parameter, repulsive neural interactions would putatively push the layers apart perceptually. This should make it more difficult to detect expansions and contractions at small-scale factors, which should be reflected in the data as a shallower slope on the left-hand side of the V-shaped discrimination functions. In fact, there is a small trend in the opposite direction: the slope of the discrimination functions at small-scale factors is slightly steeper than the slopes with large-scale factors. We therefore conclude that repulsive interactions within the layered pattern were not a limiting constraint in these experiments.

4. It has been suggested that the "pulling" effect described by Julesz and Chang (1976) could account for the perceptual "advantage" of the layered pattern over the volume form. This argument is based on the finding that a few percent of bias will pull an ambiguous stereogram to the depth of the unambiguous points. The effectiveness of the pulling strength decreases as the disparity of the bias dots differs from the

disparity of the ambiguous plane. In the present experiments, the greatest depth interactions occurred for the volume pattern having many different disparities. Recall that observers had to compare the depths of the front and back surfaces of the layered pattern to each surface presented in isolation. No depth differences were apparent. In contrast, the front and back surfaces of the volume pattern only seemed to have a depth interval half that of the front and back surfaces of the layered pattern presented in isolation. It therefore seems unlikely that the depth differences observed with the layered and volume patterns could be understood as cooperative interactions between disparity mechanisms tuned to similar disparities.

It may also be argued that the "pulling effect" does not necessarily support cooperative interactions between similarly tuned disparity detectors, but rather, some form of competition within a field of disparities. Consider the fact that as little as 2% bias is sufficient to overcome an observer's natural bias. Thus, the perceptual system makes a choice between correlation values separated by a very small percentage of correlation. How can this choice be made? Describing the phenomenon as a "pulling effect" suggests that the unambiguous dots are somehow computed first, and that they then "pull" the ambiguous pattern to that depth plane. Given their sparseness, however, this seems unlikely. Indeed, the effective "pulling strength" of these dots increases with exposure duration. These dots therefore affect observer bias relatively late in the processing of the ambiguous pattern. This indicates that there would be some initial build-up of "matching energy" specific to the natural bias of the observer that these unambiguous dots eventually overcome. This initial build-up would make it difficult to understand how such a small amount of bias can influence the organization of the stimulus. Using a dynamic systems analogy, the early processing appears to place the system within the basin of the attractor corresponding to the observer's natural bias. The question then becomes one of rationalizing how so few unambiguous dots can break the stability of this attractor. The results of the studies of Julesz and Tyler on neuroentropy offer a viable mechanism for just such a process. These papers demonstrated that observers are much more sensitive to a burst of decorrelation in a correlated pattern than to bursts of correlation in a decorrelated pattern. This finding suggests that it would be the stereo system's sensitivity to decorrelation that breaks the stability of the observer's biased state. Thus, it may be the suppression generated by the inability to form 100% matches at the observer's natural bias state that "pushed" rather than "pulled" the percept into the alternative organization. The point of this argument is that even the "pulling" effect is not definitive evidence for cooperative interactions among similarly tuned disparity mechanisms.

5. Many demonstrations of hysteresis involve the manipulation of some visual parameter in real time, requiring observers to report transitions from one pattern type to another in two possible directions. By comparing the reported transition "points" (from pattern₁ to pattern₂, and from pattern₂ to pattern₁), hysteresis is indicated if these two points occur at different values of the manipulated parameter. The significance of this method crucially depends on a number of distinct factors. First, it must be demonstrated that any observed difference is not simply the consequence of a reaction time problem. This potential confound has been recognized in all published reports of hysteresis, and since the observed lags are often an order of magnitude larger than reaction times, this does not represent a significant obstacle. A second problem exists with this method that is often not acknowledged, however. By allowing observers to specify the "point" of a transition, it must be the case that the transition between the two states is perceptually well defined in both directions of the transition. To illustrate, consider what could happen if only one of the transitions is well defined (i.e., perceptually abrupt), while the other transition occurs very gradually. It may still be possible to generate data that exhibit hysteresis simply because the observer has to apply different criteria when deciding whether a change has occurred or not. Thus, one could obtain a hysteresis simply because the observer is waiting for some percept to disappear "completely," a state that would not be perceptually well defined if the transition was relatively gradual. This is a potential explanation for a number of published experiments that have revealed hysteresis effects (Williams & Phillips, 1987; Wil-

liams et al., 1986). Indeed, Duwaer and van den Brink (1981) found large variability in the threshold for diplopia, depending on the criterion used by the observers.

It therefore seemed prudent to use a discrimination methodology to measure hysteresis to avoid the possible confounds of observer bias. In the first two experiments, observers were forced to discriminate a transition between two pattern types: a stereoscopic volume of dots versus two transparent layers. More specifically, observers were required to detect perturbations in the structure of the disparity fields of the two

patterns. This method has a number of advantages. First, observers' phenomenal reports were only used to supplement discrimination data, allowing for a more objective definition of what they could or could not see. Second, this discrimination performance allows the stability of these two patterns to be compared directly: the more stable the percept, the harder it should be to detect perturbations.

(Manuscript received July 22, 1991;
revision accepted for publication January 17, 1992.)