Importance of color in the segmentation of variegated surfaces

Andrea Li* and Peter Lennie†

Center for Visual Science and Department of Brain and Cognitive Sciences, University of Rochester, Rochester, New York 14627

Received August 17, 2000; revised manuscript received November 29, 2000; accepted December 4, 2000

We examined how variations in color and brightness are used by the visual system in distinguishing textured surfaces that differed in their first- or second-order statistics. Observers viewed a 32 × 32 array containing two types of square elements differing in chromaticity or luminance or both. The spatial distributions of the two kinds of elements were varied within the array until observers could distinguish two juxtaposed regions. At low but not at high contrast, observers are better able to distinguish regions when the elements differ only in chromaticity than when they differ only in luminance. The advantage of color at low contrasts results from the greater visibility of the arrays defined by color variation. An observer's capacity to distinguish textures defined by variations in first-order chromatic statistics is little affected by the addition of achromatic noise but is more affected by the addition of chromatic noise. The relative robustness of chromatic cues in the face of achromatic noise leaves the visual system well equipped to exploit color variations in segmenting complex scenes, even in the presence of variations in brightness. This capacity seems to depend on mechanisms that sum over large regions: When surfaces differ in their second-order statistics and cannot be distinguished by mechanisms that sum over large regions, the advantage of color is much diminished. © 2001 Optical Society of America

OCIS code: 330.1720.

1. INTRODUCTION

A good deal is known about the detection and discrimination of simple patterns defined by pure luminance variations or pure color variations: The visual system is relatively more sensitive to brightness variations at high spatial and temporal frequencies and to chromatic variations at low spatial and temporal frequencies.[1–9] Brightness variations are more effective than color variations in the discrimination of movement and depth.[10,11] On the other hand, in some stereo tasks color variations are used more efficiently than brightness variations.[12] In yet other tasks, such as detection of masked gratings,[13] texture discrimination,[14] and collinearity judgments,[15] we perform about equally well with stimuli defined by color variations or brightness variations. When stimuli contain variations in both color and brightness, observers may detect gratings[16] and make stereo discriminations[12] better than they do when stimuli are defined only by color variations or only by brightness variations.

Relatively little is known about the ways in which the visual system uses information about color and brightness in more complex tasks involving the articulation of textured surfaces or about the relative importance of color and brightness variations when both are present. It is generally agreed that brightness differences provide essential information about the structure and form of a visual scene. This can be demonstrated by removing the brightness variations from an image, leaving only color variations. The resulting images tend to look fuzzy, and objects are difficult to recognize. The importance of color differences is less clear: Images devoid of color variations are generally still easy to segment, as we can appreciate in watching black and white television, or in looking at black and white photographs. When an object appears against a background containing substantial local variations in brightness, particularly those accidental variations caused by shadows and highlights, one might expect color variations to become more important cues to surface structure.[17]

The experiments described here constitute an attempt to understand the importance of color variations in segmenting spatially complex scenes. We examine two issues: (1) the relative effectiveness of color and brightness variations as cues for segmenting surfaces composed of texture elements that can vary in color, brightness, or both and (2) the robustness of color and brightness variations as cues for segmenting textured surfaces in the presence of noise.

2. THE SEGMENTATION TASK

We define segmentation as the differentiation of regions by their surface characteristics, the characteristics of interest here being spatial variations in chromaticity and luminance. In the present experiments a surface is a two-dimensional array of texture elements. Elements can differ in chromaticity, luminance, or both. Surfaces can differ in the statistics that describe the color and luminance of the elements that compose them. The simplest possible surface would contain elements of a single chromaticity and luminance; a complex surface might contain elements sampled with specified statistics from the gamut available in a three-dimensional color space.

In the present experiments the surfaces to be distinguished are formed by the left and right halves of a

0740-3232/2001/061240-12$15.00

32×32 array of elements. The elements constituting the array fall into two classes, which for illustration we call A and B. Over the whole array the two classes of elements occur in equal numbers. We vary the proportions of A and B elements in the left and right halves of the array to find the least asymmetry that an observer can detect. Within each half, the elements that are present are randomly distributed. Figure 1 shows an example of a stimulus in which the A elements are black and the B elements white. We express the asymmetry as the proportion of the class A element on one side of the array (the class B element occurs in the same proportion in the other half), and call this dominance. At the lower limit, 50% dominance (Fig. 1, left), 50% of the elements on the left are class A, black (and therefore 50% on the right are class B, white). At this dominance the array appears as a single continuous surface. At 100% dominance (Fig. 1, right), all the elements on the left are class A (and therefore all on the right are class B). At this dominance, the stimulus unquestionably contains two regions. Each 1% difference in dominance corresponds to five elements of imbalance; i.e., five elements more of one class fall on one side of the midline and five elements more of the other on the other side of the midline.

The dominance at which the left and right sides of the array can just be distinguished we take as the observer's segmentation threshold. This metric allows us to compare performance when the observer distinguishes surfaces defined by elements whose properties differ in various ways.

The surfaces that we consider in this paper are composed of elements drawn from a single plane in color space, defined by the achromatic axis and the axis of purely L–M-cone excitation.[18,19] For the simplest cases

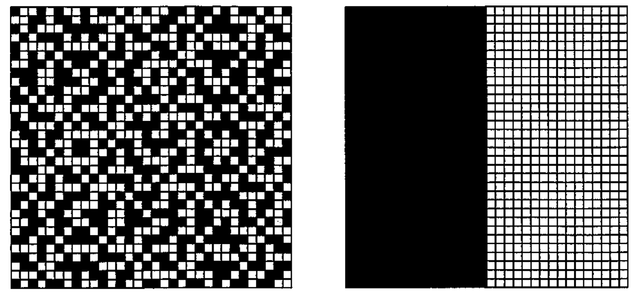

(a) (b)

Fig. 1. The arrangement of two types of elements (here shown as black and as white) in an array determines its dominance. Dominance is defined as the proportion of one element type on one side of the stimulus. At 50% dominance (left) half the elements on the left are black (and therefore half on the right are white). At 100% dominance (right) 100% of the elements on the left are black (and 100% on the right are white). Dominance was varied to the point at which observers could just distinguish two regions in the stimulus.

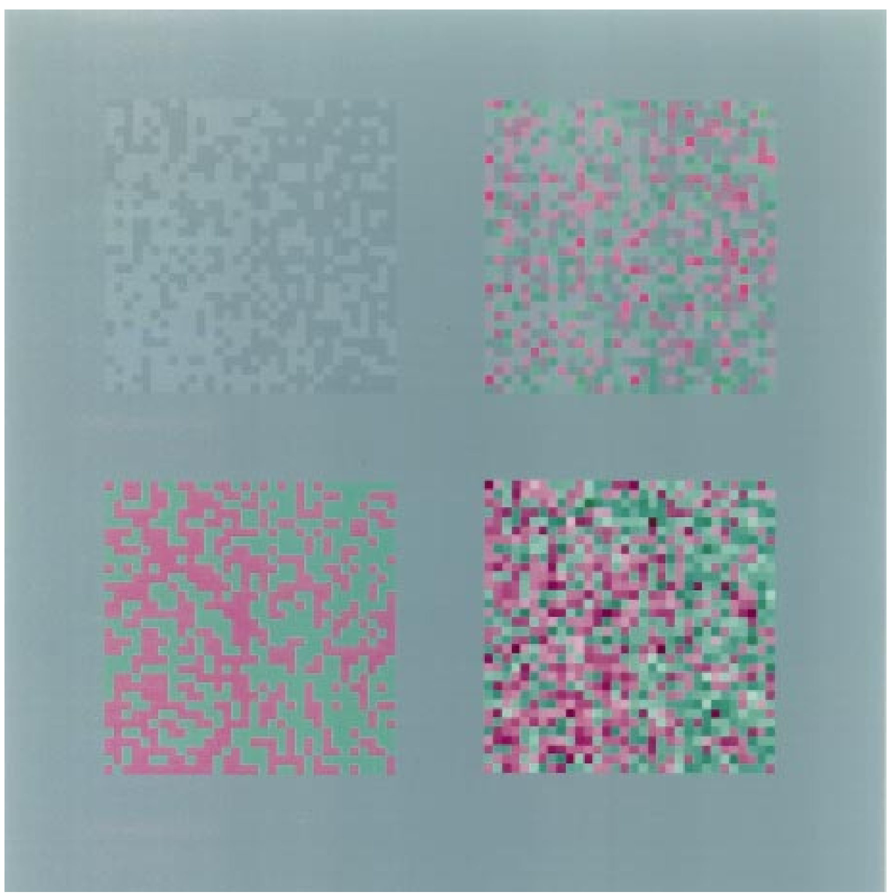

Fig. 2. Examples of simple arrays at 60% dominance. The two classes of elements differ only in luminance (top left) or only in chromaticity (bottom left). Panels on the right show the same arrays, now with element values perturbed by the addition of one-dimensional noise. At top right, chromatic noise has been added to the light- and dark-gray elements of the achromatic array; at the bottom right, achromatic noise has been added to the red and green elements of the chromatic array.

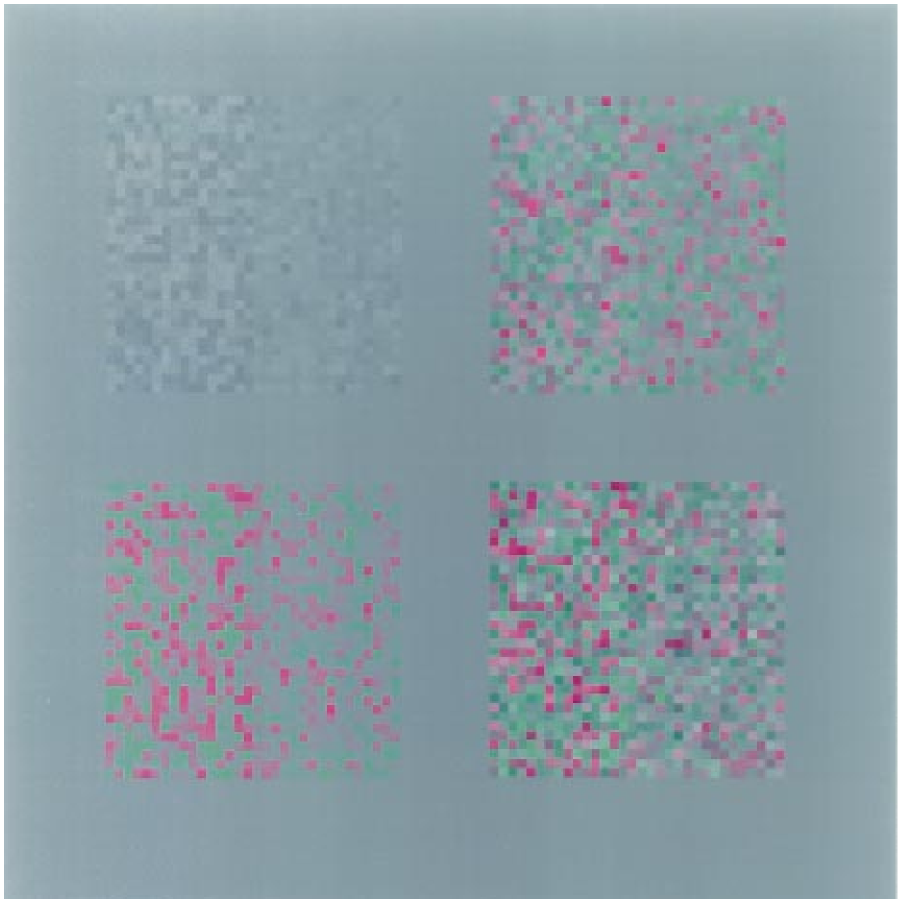

Fig. 3. Examples of complex arrays at 100% dominance. The two classes of elements differ in the statistics of the distributions of element luminances (top left) or chromaticities (bottom left). In both cases the distributions of elements have zero mean but differ in shape, one being Gaussian, the other uniform. The sides of the arrays on which elements are uniformly distributed (left) appear to have higher contrast. Panels on the right show the same arrays, now with element values perturbed by the addition of one-dimensional noise. At top right, chromatic noise has been added to the distributions of gray elements of the achromatic array; at bottom right, achromatic noise has been added to the distributions of red and green elements of the chromatic array.

that we explored, each class of elements from which the surfaces were assembled was defined by a single chromaticity and luminance (e.g., class A was red and class B was green). Figure 2 (left) shows examples of arrays at 60% dominance: In the top panels, A and B elements differ only in luminance; in the bottom panels, A and B elements differ only in chromaticity. For these cases, the mean chromaticity and luminance of the left and right sides of the array varies with dominance, so mechanisms that integrate across regions as large as half the array will be best suited to the task.

For the most complex cases that we explored, each class was defined by a distribution of chromaticities and luminances, with class and A and class B having different statistics. Figure 3 (left) shows examples of arrays at 100% dominance: In the top panels, the A and B elements are defined by distributions of luminance values with the same mean but different shape (values of A elements had a uniform distribution; values of B elements a Gaussian distribution); in the bottom panels, A and B elements are defined by corresponding distributions of chromaticities along the L–M axis. For these cases, the mean chromaticity and luminance of the two sides of the array do not change as a function of dominance. Mechanisms that integrate over regions smaller than each side of the array will thus be required for computation of local distributions of chromaticities and luminance values.

To simulate accidental variations in color and brightness within and across surfaces, we can add noise to perturb the values of individual texture elements. This noise provides an offset that displaces the position of an element in color space. By varying the properties of the noise, we can explore the selectivity of the mechanisms that distinguish the surfaces. Figure 2 (right) shows the array (left) perturbed by the addition to each element of an offset that displaces its position in color space: at the top, each element is changed in color by an offset along the L–M axis; at the bottom right, each element is changed in brightness by an offset along the achromatic axis. Similarly, Fig. 3 (right) shows the array (left) perturbed by the addition to each element of an offset that displaces its position in color space: At the top, each element is changed in color by an offset along the L–M axis; at the bottom right, each element is changed in brightness by an offset along the achromatic axis.

3. METHODS

A. Stimuli

Stimuli were generated by a PIXAR II computer and displayed on a 1280 × 1024 Nanao T560i monitor at a frame rate of 60 Hz. The Pixar provided for intensity quantization of 10 bits per gun; after correction for display nonlinearities, the effective range was reduced by more than 1 bit. Each element of the 32×32 array was 4 pixels square and subtended 12 arc min at a viewing distance of 57 cm. The whole array subtended 6.5×6.5 deg and was enclosed by a gray surround of the same space-average chromaticity and luminance ($x = 0.335$, $y = 0.341$, $35\,\mathrm{cd/m^2}$) that subtended 26 deg. Observers viewed the display binocularly with a free head in a dimly illuminated room. Under these conditions observers' pupils were ~4 mm in diameter.

The chromaticity and luminance of each element were specified by a vector in a color space defined by three axes: an axis along which the relative excitations of the L and M cones vary in opposition without change in luminance (the L–M axis), an axis along which the excitation of S cones varies (the S axis), and an achromatic axis along which all three cones are proportionally excited.[18,19] The intersection of these axes (the "white" point) lay at the mean chromaticity and luminance of the display. In all experiments, array elements were drawn from a single plane in the color space, differing in color (red and green) through variation along the L–M axis of the color space, and/or in brightness (light and dark) through variation along the achromatic axis of the color space.

Figure 4 shows how the chromaticity or luminance of elements was perturbed by adding chromatic or achromatic noise. At the top left, the values of elements differ only in brightness (along the achromatic axis). At the top right, perturbations along the L–M axis of color space have been added, thereby changing the color of each element. At the bottom left, the values of elements differ

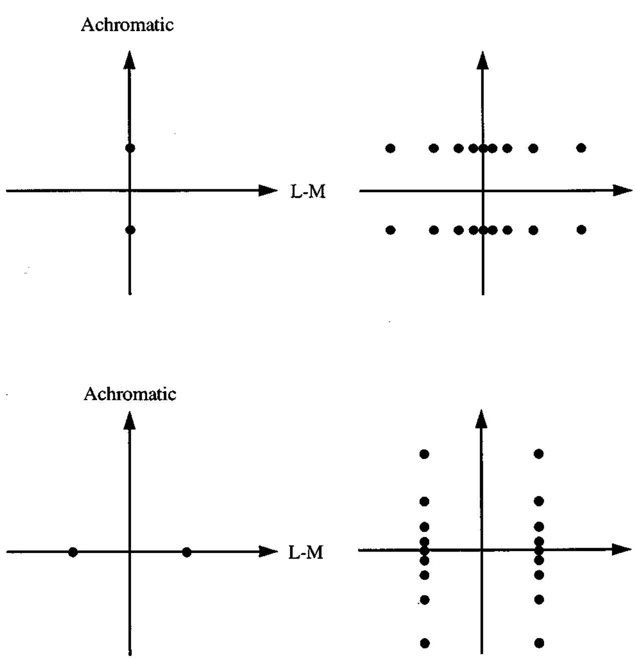

Fig. 4. Diagrams showing the locations in color space of the elements constituting the simple arrays shown in Fig. 2. In the absence of noise, arrays contained only two element values (left); when noise was added, the chromaticity (top right) or luminance (bottom right) of the elements was perturbed, so the values of elements were distributed along lines in color space. For additional information see text.

only in color (along the L–M axis). At the bottom right, these values have been perturbed by noise displacements along the achromatic axis, thereby changing the brightness of each element. Noise in the experimental stimuli was continuous with a Gaussian distribution.

Stimuli of this sort have been used previously[20] to examine the tuning of chromatic mechanisms underlying texture segmentation.

B. Determining Isoluminance

In preliminary measurements we established, for each observer, the locus of isoluminant values along the L–M direction in color space. The observer viewed a colored (red or green) target square subtending 2 deg alternating at 15 Hz with a gray square of the same size, with chromaticity and luminance fixed at the white point. The luminance of the colored square was adjusted to minimize perceived flicker. We made measurements along the L–M axis, in (rms cone) contrast steps of 5%, up to 20%, where the CIE coordinates were ($x = 0.397$, $y = 0.28$) and ($x = 0.237$, $y = 0.386$). The means of five measurements at each contrast value were used to estimate the locus of each observer's isoluminant L–M axis.

C. Determining Dominance Thresholds

The general procedure was to establish how variations in the properties of the elements in an array affected an observer's capacity to discern changes in dominance within the array. This was done through a forced-choice procedure. The observer was presented with two successive arrays of elements, one (in the null interval) containing statistically identical distributions of elements in its left and right halves (50% dominance) and the other (in the target interval) containing elements distributed unevenly between the halves. The task was to determine the interval in which the dominance of the array deviated from 50%. A single trial consisted of two 1-s intervals separated by a 1-s blank period during which the screen was gray (white point). In the target interval, dominance was varied according to a method of constant stimuli with a step size of 1%. An appropriate range of six dominance values was estimated in a preliminary session.

A session consisted of 210 trials (35 for each of 6 dominance values, randomly interleaved), preceded by 2 min of adaptation to the gray background. For each level of dominance, the positions of elements within each half of the array were shuffled randomly between trials to prevent the observer from using local clusters to identify the target interval. The dominance of the two element types randomly alternated sides in the array to prevent comparison of particular portions of the stimulus in the null and target intervals. By pressing one of two buttons on a mouse, observers indicated the interval in which the dominance of the array differed from 50%. Audible feedback was given. Each session lasted approximately 10 min.

Using a maximum-likelihood procedure, we fitted a Weibull function to the measurements from each session:

$$f(x) = 1 - (1 - c)\exp(-x/a)^b \tag{1}$$

where a is the threshold parameter, b is the slope parameter, and c is the chance level parameter that was set to

0.5 for a two-alternative forced-choice task. For each session, threshold was taken as the dominance that yielded 75% correct on the fitted curve. The final threshold estimate was the average of those obtained in three sessions.

Three observers served in the experiments: LB and AB were undergraduates ignorant of the purpose of the experiment; AL, an author, was an experienced observer. All had normal or corrected-to-normal acuity and normal color vision.

4. SEGMENTATION BY COLOR AND BY BRIGHTNESS

A. Segmentation by Color versus Segmentation by Brightness

To characterize the value of color as a cue for segmentation, we compared observers' capacities to segment an array of elements that differed only in chromaticity, only in luminance, or in both chromaticity and luminance. In the chromatic condition, elements were red and green and lay on the L–M axis; in the achromatic condition, they were light and dark and lay on the achromatic axis; in the mixed condition, elements were light red and dark green, with values chosen so that the ratio of cone contrast of the chromatic component to the cone contrast of the achromatic component was 0.5. For each of these conditions we measured two things: first the visibility threshold—the element contrast required for an observer to detect the presence of an array at 50% dominance; second, the segmentation threshold—the dominance required for an observer to segment the array—at a range of element contrasts. For each kind of measurement we used the forced-choice procedure described in Section 3 (for the detection measurements, one interval contained only the blank screen) and estimated thresholds from the fitted Weibull functions.

1. *Visibility Thresholds*

Figure 5 shows the rms cone contrasts required by each of two observers to detect arrays defined by the different kinds of elements. Vertical bars show 95% confidence intervals around the mean of three estimates. For both observers, thresholds for arrays defined by achromatic contrast are 2–2.6× those for arrays defined by chromatic contrast. Thresholds for the mixed condition lie at intermediate values.

2. *Segmentation Thresholds*

Figures 6 and 7 show, for the two observers from Fig. 5, how the dominance required for segmentation varies with element contrast. Segmentation thresholds for the different types of elements are represented by different symbols. In Fig. 6 the contrast is expressed in visibility units (from Fig. 5); in Fig. 7 the contrast is expressed as rms cone contrast. Vertical bars represent 95% confidence intervals around the mean of three estimates.

For all conditions, threshold dominance declines with increasing contrast and almost identically when contrast is expressed in visibility units. When contrast is expressed as rms cone contrast, observers perform best with elements differing in color and worst with elements differing in luminance. The visual system apparently handles chromatic signals more efficiently (visibility thresholds are lower), but, given adequately visible stimuli, the mechanisms that distinguish surfaces seem

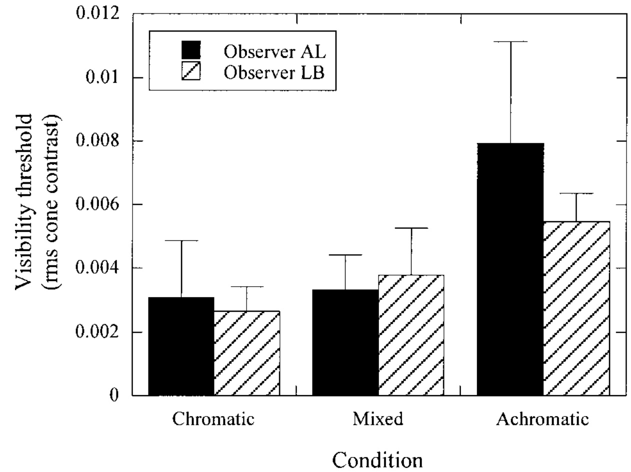

Fig. 5. Visibility thresholds in rms cone contrast for chromatic, mixed, and achromatic arrays of elements. Error bars represent 95% confidence intervals around the means, each based on three sets of measurements.

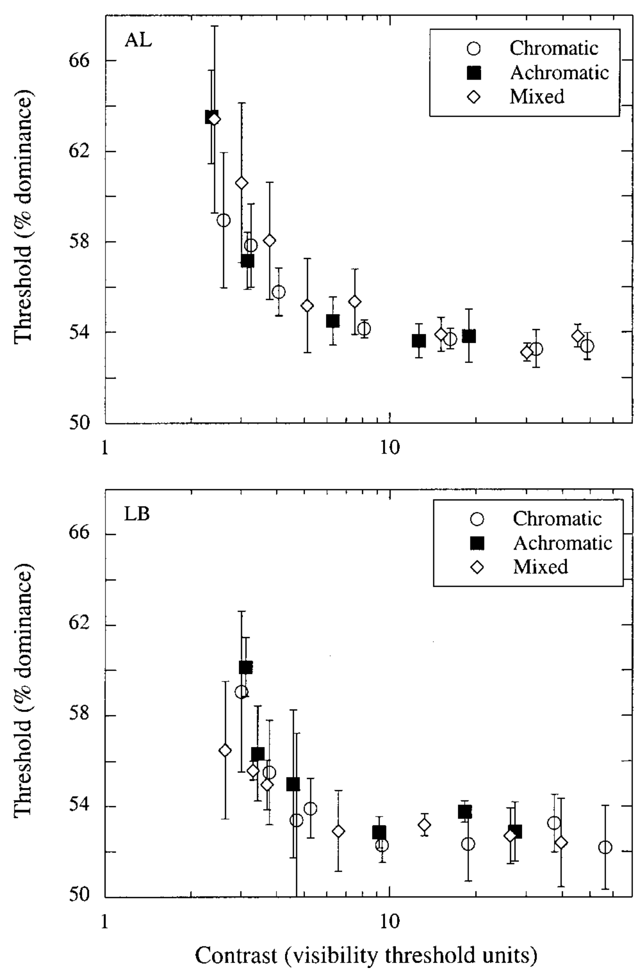

Fig. 6. Dominance required for segmentation, as a function of element contrast, where contrast is expressed in units of threshold visibility. Different kinds of elements (chromatic, achromatic, mixed) are shown by different symbols identified in the figure.

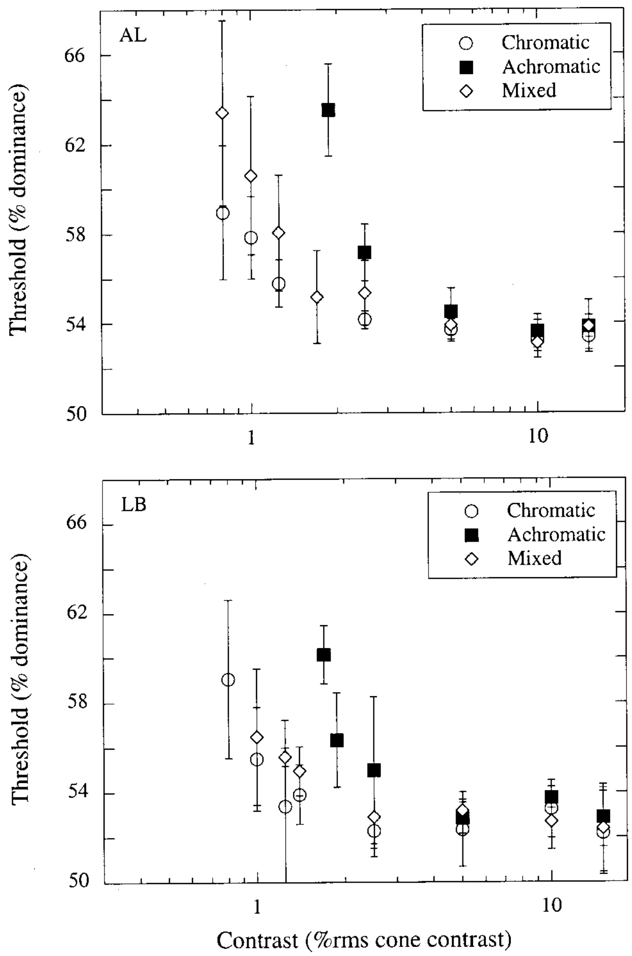

Fig. 7. Dominance thresholds from Fig. 6, replotted as a function of rms cone contrast. Other details as in Fig. 6.

to be indifferent to the chromatic dimensions along which elements differ.

A potential concern about the results in Figs. 5–7 is that the observer's capacity to distinguish the two halves of the display depends not on an appreciation of the different surface characteristics—the statistics of the two-dimensional arrays of elements—but on detecting a local discontinuity at the boundary between them. We have explored this in subsidiary experiments on one observer, using a task that was exactly as before except that the two halves of the array were separated by either 0.125 or 0.25 deg, producing a prominent gap between them. When a gap of either width separated the halves of the array, performance was reliably better than when the halves were juxtaposed. Evidently observers distinguish the juxtaposed surfaces not by detecting the boundary between them but by relying on some signal accumulated more broadly from within them: The array can be segmented only after the surfaces have been distinguished. Even when the task is to distinguish juxtaposed uniform fields that differ in chromaticity, observers apparently do not obtain useful information from the edge between them: Eskew and Boynton[21] found that sensitivity depended most on the area of the fields; for rectangular fields of equal area but different shape, performance became poorer as the common border between them was lengthened.

B. What Makes Color a Potent Cue to Surface Structure?

The results described so far make clear that when arrays of elements differ in their first-order statistics (mean values), those that differ in mean chromaticity are more efficiently distinguished than those that differ in mean luminance. The greater visibility of the chromatic patterns (Fig. 5) might have resulted from our use of arrays in which element sizes favored chromatic mechanisms, which are known to be relatively more sensitive to low spatial frequencies (the dominant spatial frequency in the arrays was 2.5 c/deg). To explore the sensitivity of our findings to the spatial characteristics of the array we undertook subsidiary experiments in which we varied the numbers and sizes of array elements that differed only in chromaticity. Within limits imposed by hardware, we explored the effect of varying element area over a factor of 4, allowing array area to vary commensurately, or holding area constant by varying the number of elements in the array. Table 1 shows the array configurations used (the configuration used in the original measurements is in boldface).

Figure 8 shows, for two observers, the effect of reducing element area by a factor of 4 while holding the number of elements constant or holding array area constant. These manipulations have little effect on the observers' capacity to segment the array. (For observer AL, increasing the total number of elements lowered threshold slightly.)

Figure 9 shows, for the same observers, the effect of increasing element area by a factor of 4 while holding the number of elements constant or by holding array area constant. Again the change in element size has little effect on performance. Where Figs. 8 and 9 show any effects of the spatial characteristics of arrays, it is such as to suggest improved performance with more and smaller elements. Were the generally superior performance with chromatic arrays due to their elements better matching the spatial properties of the mechanisms that analyze them, we would have expected the relative advantage of chromatic arrays to increase as element sizes were made larger. Instead, the advantage was, if anything, slightly diminished. Chromatic aberration will tend to introduce achromatic contrast into isoluminant patterns as spatial frequency is raised, conceivably compensating for some loss of sensitivity in chromatic mechanisms. For stimuli of the spectral and spatial composition used here, however, the effects of chromatic aberration should be inconsequential.[22]

These results suggest that observers' superior performance in distinguishing surfaces defined by chromatic

Table 1. Size and Number of Texture Elements Used to Explore Effects of Configuration on Performance

Element Size	Number of Elements		
	22 × 22	**32 × 32**	46 × 46
8.5 min		X	X
12 min		**X**	
17 min	X	X	

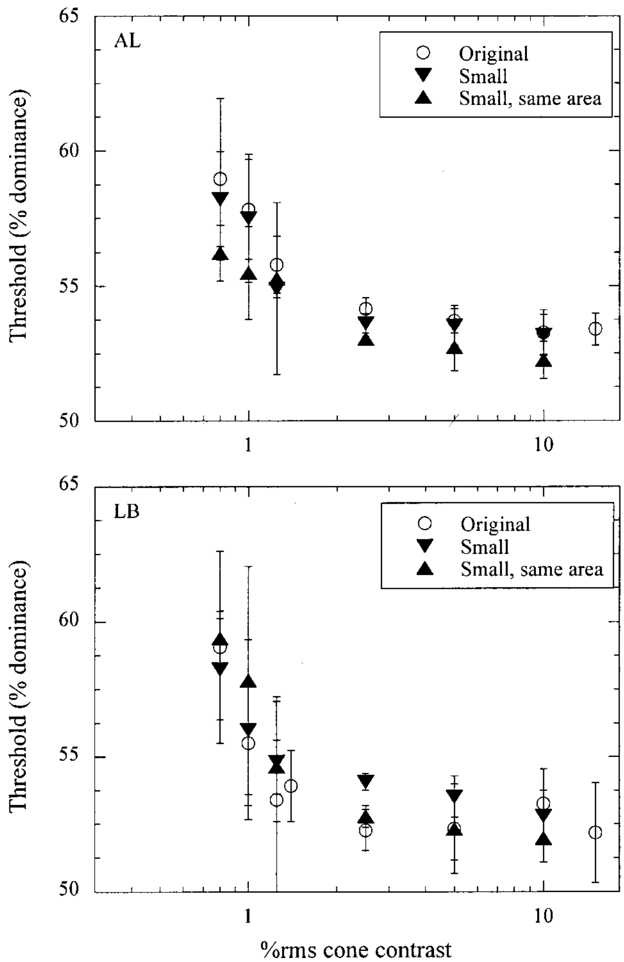

Fig. 8. Dominance required for the segmentation of arrays of chromatic elements, each one fourth of the area used in the main experiment. In the small condition (downward triangles), the number of elements was kept constant and the area of the array was changed; in the small/same area condition (upward triangles), the area of the array was kept constant and the number of elements was changed.

differences among elements does not stem from our having fortuitously chosen an advantageous element size, but the results do not demonstrate that performance actually depends on mechanisms that are sensitive to the properties of individual elements. It could, for example, depend on higher-order integration by mechanisms that compute the average chromaticity or luminance across each half of the array. To establish whether observers can use information on a scale smaller than each half of the array, we explored the segmentation of surfaces that were identical in their first-order statistics but differed in their second-order characteristics.

The two element classes from which surfaces were assembled were defined by probability distributions of element chromaticities that had equal mean (centered around white) but different shape: In one (uniform) distribution, all chromaticities within a limiting range were equally likely to occur; in the other (Gaussian) distribution, chromaticities close to white occurred more often. The term second-order thus describes the order of the statistics of the color histogram of each distribution rather than the joint statistics of pairs of elements.[23] At 100% dominance the two sides of the array appear to differ in overall contrast (Fig. 3, left): The side on which elements are chosen with a uniform probability distribution (left) appears to have higher contrast. To distinguish surfaces the observer must sample on a scale smaller than each half of the array. An observer performing optimally would sample at the scale of the individual elements.

We measured segmentation thresholds for arrays in which the distributions of element values lay along either the achromatic axis or the L–M axis. The Gaussian distribution had a standard deviation of 10× the visibility threshold and was clipped at ±3 standard deviations. The range of the uniform distribution was clipped at the same limits as the Gaussian distribution. The variance of the uniform distribution was thus slightly more than 3× that of the Gaussian distribution. Both distributions had zero mean; the average contrast of each was high enough that increases in contrast would not be expected to improve performance. Other details were the same as those for the experiments of Fig. 6. We tested observer AL and observer AB (instead of observer LB).

Figure 10 shows that both observers could distinguish surfaces that differed in their second-order statistics. To do this requires mechanisms that can distinguish the statistics of distributions of local variations in contrast. For

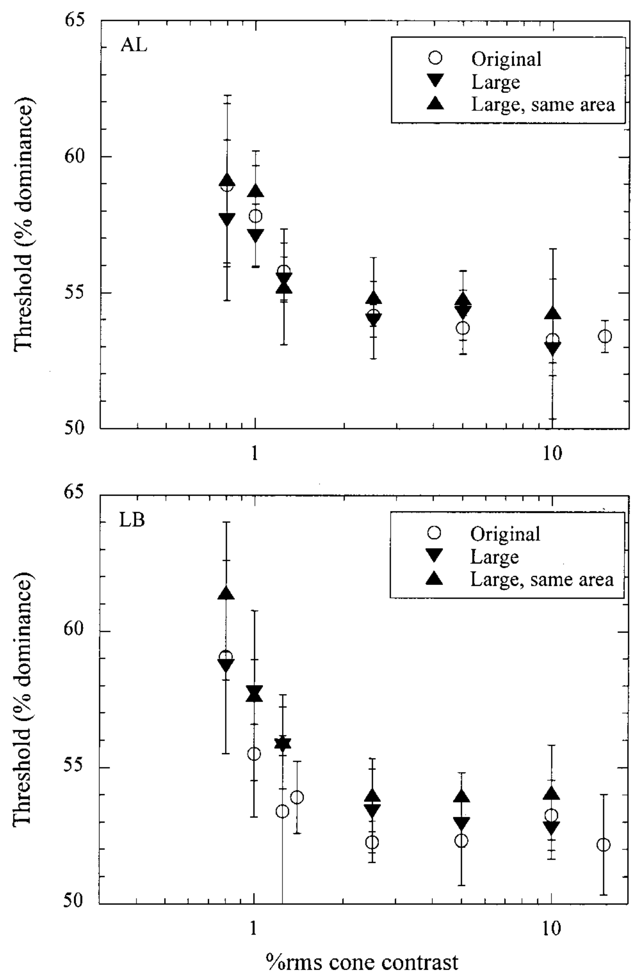

Fig. 9. Dominance required for the segmentation arrays of chromatic elements, each 4× the area used in the main experiment. Other details as for Fig. 8.

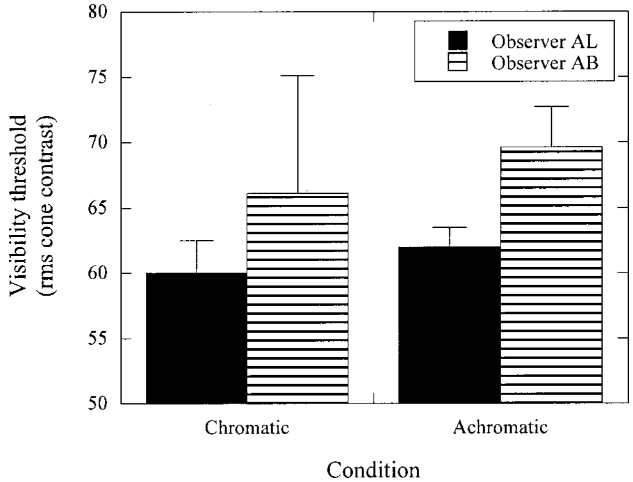

Fig. 10. Dominance required for segmentation of arrays in which the two classes of elements were defined by the different shapes of their distributions of chromaticity or luminance. One distribution was Gaussian, the other uniform (see text for details). Error bars represent 95% confidence intervals around the means, each based on three sets of measurements.

the particular distributions used here, surfaces defined by elements differing in color were no more distinguishable than those defined by elements differing in luminance. The results shown in Fig. 10 make clear that when observers are forced to analyze spatial variations in chromaticity on the scale of individual elements, surfaces can be distinguished satisfactorily; large-scale spatial summation of information is evidently not necessary.

5. EFFECT OF NOISE ON SEGMENTATION

Accidental variations in surface properties (particularly those caused by shadows and highlights) are common in natural scenes and constitute noise in which underlying differences in surfaces have to be discerned.[24] To understand the relative importance of brightness and color differences, one wants to know how robust they are as segmentation cues in the face of noise.

A. Segmentation by First-Order Color Differences

1. *Stimuli*

Using arrays of elements that differed in chromaticity at a contrast of 10× visibility threshold—high enough that in the absence of noise, performance is independent of contrast (see Fig. 6)—we measured the dominance required for the array to be segmented, as a function of the amplitude (i.e., standard deviation) of noise added to it.

Each element in the array had added to it an offset in color space chosen from a one-dimensional distribution of Gaussian noise, with zero mean and values clipped at ±3 standard deviations. Offsets displaced elements along either the L–M axis (chromatic noise) or the achromatic axis (achromatic noise). The bottom row of Fig. 2 shows examples of an array of elements with and without added achromatic noise.

2. *Results*

Figure 11 shows, for two observers, the relative dominance required for segmentation in the presence of different amounts of chromatic noise (circles) or achromatic noise (squares). Relative dominance is the ratio of the threshold dominance in noise to the threshold dominance in the absence of noise, with threshold here taken as the percentage by which dominance exceeds 50%. Vertical bars represent 95% confidence intervals around the mean of three sets of measurements. Noise amplitude is expressed in units of visibility threshold. This was established in a separate set of measurements in which we added noise to chromatic elements at 10× threshold configured at 50% dominance. We used the slope of the best-fitting straight line to characterize the effects of noise. For both observers chromatic noise substantially impairs segmentation: At 5× its visibility threshold, chromatic noise raises the threshold dominance by a factor of 2–2.7. On the other hand, achromatic noise had little effect, and the highest noise amplitudes we could produce (~50× threshold) raised threshold dominance only by ~10%. The slopes of the lines obtained with chromatic noise were 81× (observer AL) and 59× (observer LB) steeper than those obtained with achromatic noise.

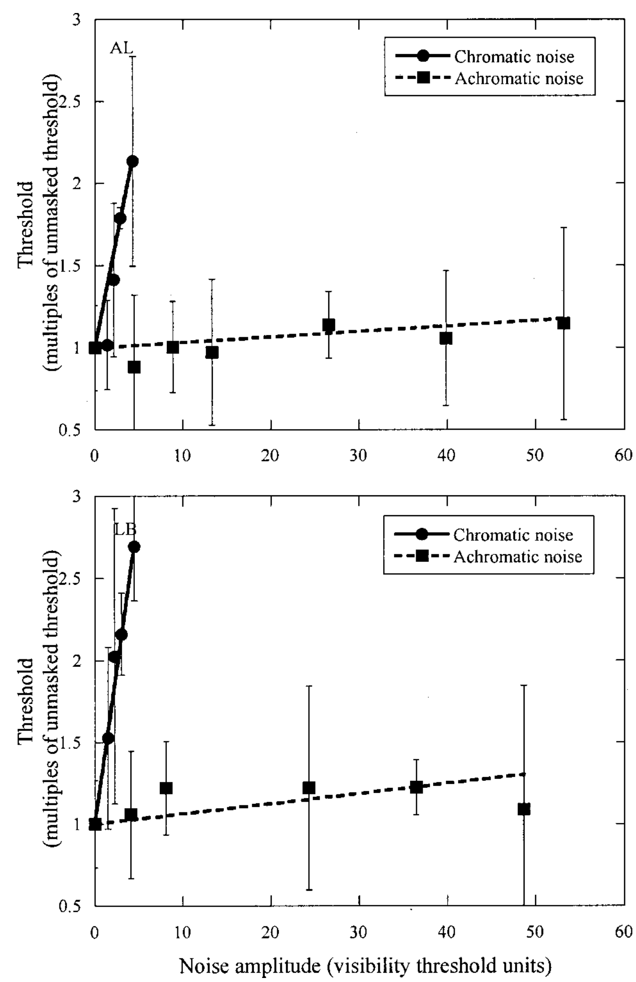

Fig. 11. Dominance required for segmentation of a chromatic array in the presence of chromatic noise (solid lines) and achromatic noise (dashed lines). Thresholds are expressed as multiples of the unmasked threshold. Noise amplitude is expressed as the standard deviation in visibility threshold units. Points have been fitted linearly with the zero-noise point fixed. The observers' sensitivities to chromatic noise are 81× (upper) and 59× (lower) greater than their sensitivities to achromatic noise.

B. Segmentation by First-Order Brightness Differences

1. Stimuli

We used arrays of elements differing in brightness, at a contrast set to be 10× visibility threshold, and measured the dominance required for the array to be segmented, as a function of the amplitude of chromatic or achromatic noise added to it. Display limitations restricted the maximum amplitude of chromatic noise to 6.5% rms cone contrast. The top row of Fig. 2 shows examples of the array with and without added chromatic noise.

2. Results

Figure 12 shows, in the same format as Fig. 11, how the relative dominance required for segmentation depends on the amplitude of added noise. Over the range of contrasts realizable on the monitor, segmentation thresholds are affected by both achromatic noise and chromatic noise. Achromatic noise at 5× threshold raises threshold dominance by a factor of 2 (thick solid lines)—about as much as chromatic noise raises the threshold dominance required for chromatic targets to be distinguished (Fig. 11). Chromatic noise (thick dashed lines) provides a weaker though still effective mask: The slopes of the lines obtained with achromatic noise were 4.5× (observer AL) and 2.4× (observer LB) steeper than those obtained in chromatic noise, a much smaller ratio than characterizes the lines in Fig. 11.

If the mechanisms underlying segmentation by color and segmentation by brightness were similarly affected by internal and external noise, we would expect the effect of chromatic noise on a chromatic target (Fig. 11, solid lines) to equal that of achromatic noise on an achromatic target (Fig. 12, solid lines). For direct comparison, the solid lines from Fig. 11 have been replotted in Fig. 12 as the thin solid lines. Although for both observers the effect of chromatic noise on a chromatic target was slightly greater than the effect of achromatic noise on an achromatic target (23% greater for observer AL and 36% for observer LB), the chromatic–chromatic masking lines generally fall within the 95% confidence intervals around the achromatic–achromatic masking lines. The lines representing the effect of achromatic noise on a chromatic target from Fig. 11 have also been replotted in Fig. 12 (thin dashed lines). The effect of achromatic noise on a chromatic target is significantly less than the effect of chromatic noise on an achromatic target (thick dashed lines).

Figure 13 shows, for both observers, slopes of the masking functions obtained in all the experimental conditions. It provides a summary of the effectiveness of the two kinds of masks on the two kinds of targets. The most potent mask is chromatic noise applied to a chromatic target; the least potent mask is achromatic noise applied to a chromatic target.

C. Why Is Segmentation by Color Resistant to Achromatic Noise?

The small influence of achromatic noise on an observer's ability to distinguish surfaces defined by first-order differences in chromaticity might reflect the different spatial sampling properties of the mechanisms that deal with the chromatic signal and achromatic noise. If the observer could engage mechanisms that integrated color signals over a large region, these would afford some protection against the influence of achromatic noise, which leaves the mean color of each distribution unchanged. Were this the explanation of the robust performance, we might expect that when the only information that could distinguish surfaces occurs on a scale smaller than half of the array, performance will be more susceptible to the influence of noise. To explore this idea we measured the dominance required to distinguish surfaces by their second-order chromatic statistics in the presence of achromatic noise and the dominance required to segment by second-order achromatic statistics in the presence of chromatic noise. The two classes of elements from which surfaces were assembled were drawn from a one-dimensional Gaussian distribution or a one-dimensional uniform distribution. The Gaussian distribution was clipped at ±3 standard deviations, where the standard deviation was set to 10× the visibility threshold; the uniform distribution was clipped at the same limits. As before, the variance of the uniform distribution was thus slightly more

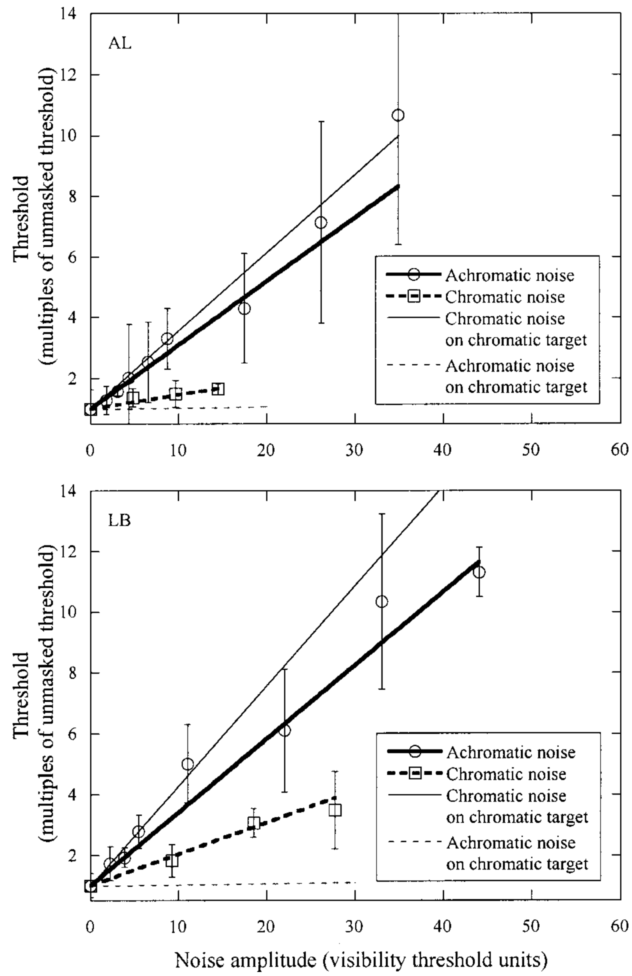

Fig. 12. Dominance required for segmentation of an achromatic array in the presence of achromatic noise (thick solid lines) and chromatic noise (thick dashed lines). The observers' sensitivities to achromatic noise are 4.5× (upper) and 2.4× (lower) greater than their sensitivities to chromatic noise. For direct comparison, the masking lines from Fig. 11 have been replotted as thin solid lines (chromatic noise on a chromatic target) and thin dashed lines (achromatic noise on a chromatic target).

than 3× that of the Gaussian distribution. As in the previous noise experiment, the standard deviation of the noise was expressed in visibility threshold units. We tested observer AL and observer AB (instead of observer LB).

Figure 14 shows, for both observers, the slopes of the masking functions that describe how the dominance threshold depends on noise amplitude. When achromatic surfaces differ in their second-order statistics, chromatic noise is an effective mask. This was also the case when surfaces were defined by differences in first-order statistics (Fig. 13). When chromatic surfaces differ in their second-order statistics, achromatic noise is an effective mask. This was not the case when surfaces were defined by differences in first-order statistics. For surfaces defined by first- and second-order statistics, the effects of chromatic noise on an achromatic target exceeded the effects of achromatic noise on a chromatic target. However, the relative difference was far more pronounced for surfaces defined by first-order statistics (13.8× for observer AL and 16× for observer LB) than for surfaces defined by second-order statistics (1.8× for observer AL and 1.2× for observer LB). These results suggest that chromatic textures are immune to masking by achromatic noise when (in the case of first-order statistics) observers can benefit from summing signals over a substantial region; when the task requires that observers sample on a scale smaller than half of the array, achromatic noise substantially impairs performance.

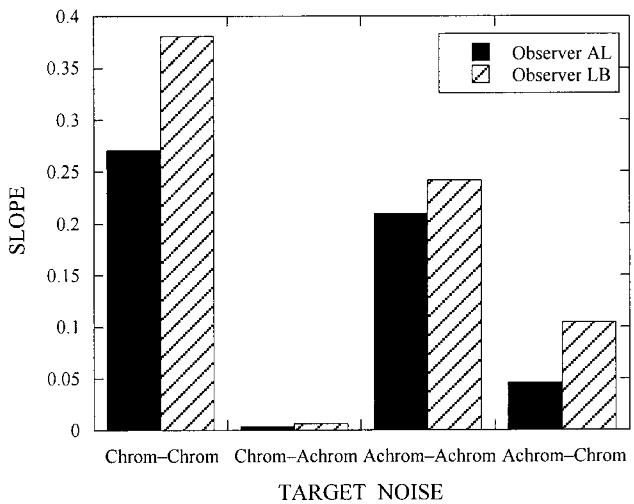

Fig. 13. Potency of achromatic and chromatic noise in masking the segmentation of arrays in which the two classes of elements differed in their first-order statistics (mean chromaticity or mean luminance). For each of two observers the bars show the slopes of masking functions for all combinations of element type (chromatic or achromatic) and noise type (chromatic or achromatic).

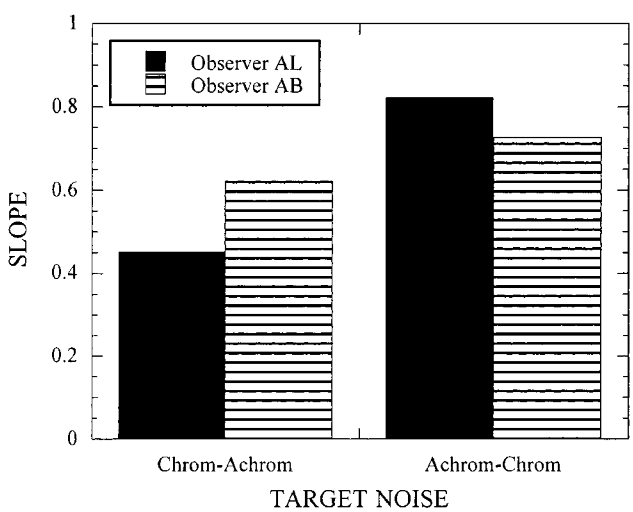

Fig. 14. Potency of achromatic and chromatic noise in masking the segmentation of arrays in which the two classes of elements were defined by different second-order statistics (the different shapes of their distributions of chromaticity or luminance). For each of two observers the bars show the slopes of masking functions for a chromatic target in achromatic noise and an achromatic target in chromatic noise.

6. DISCUSSION

We have demonstrated several circumstances under which differences between the chromatic structures of surfaces make the surfaces readily distinguishable.

First-order color variations provide not only effective cues to surface structure but, at low contrasts, better cues than do corresponding brightness variations (Fig. 7). This is reflected by the superior visibility of arrays defined by color variation: When the contrasts of elements are equated for visibility, observers perform equally well whether the elements are defined by differences in chromaticity or differences in luminance (Fig. 6). This finding suggests that the performance is limited only by the internal signal-to-noise ratio of the stimulus and that, at least in the chromatic domain, stimuli detected with equal reliability will provide equally effective segmentation cues. There is no need to postulate different types of mechanisms underlying segmentation by color and segmentation by brightness when mechanisms that sum signals over large regions are sufficient for the task.

The greater efficiency with which observers detect arrays whose elements vary in color than arrays whose elements vary in brightness (Fig. 5) is not easily explained. The arrays used to establish detection thresholds were at 50% dominance, so their mean chromaticity and luminance matched that of the background against which they were presented and would provide little information to mechanisms that integrated signals over a large region. (When stimuli to be detected are large regions of uniform illumination, the advantage of color over luminance is considerably greater than the factor of 2 found here. Greenstein *et al.*,[25] using pinwheel-like stimuli (opposite quadrants of a circular stimulus subtending 4 deg), found that threshold for a red–green stimulus was approximately one sixth that for an achromatic one. Li and Lennie,[20] using uniform square stimuli subtending 4 deg, found that thresholds for red and green stimuli were approximately one tenth those for light and dark stimuli.) Evidently chromatic information is more efficiently handled even in the middle range of spatial frequencies that were dominant in the arrays used in our experiments.

Our experiments to explore the influence of noise showed that segmentation by first-order color differences is robust to achromatic noise (Fig. 11)—indeed, pseudoisochromatic plates used in tests of color vision rely on this fact. On the other hand, segmentation by first-order brightness differences is clearly impaired by chromatic noise (Fig. 12), though not dramatically. Similar asymmetries have been found in other studies.[13,20,24,26] The asymmetry was accompanied by different subjective impressions of the two kinds of arrays (see Fig. 2). When achromatic noise was added to arrays of colored elements, the noise was assimilated as part of the surfaces defined by the colored elements. On the other hand, when chromatic noise was added to arrays of achromatic elements, the differences in color were more salient than the differences in brightness, making it difficult to discern two surfaces of different brightnesses. Color differences often signify differences in surface or object materials, especially in cases where objects are of similar shape and scale.[17] With this in mind, the asymmetry in the percepts is consistent with achromatic noise on a chromatic array being interpreted as chromatically textured surfaces under variable illumination (Fig. 2, bottom right); chromatic noise on an achromatic array is interpreted as elements of different material distributed throughout a surface (Fig. 2, top right).

The relative immunity of segmentation by first-order color differences to disruption by achromatic noise seems to depend on the visual system's capacity, when required, to engage mechanisms that sum signals over a large region. Surfaces defined by second-order chromatic variations cannot be discerned by mechanisms that simply sum signals over a large region. These surfaces must be distinguished by mechanisms sensitive to local variations in contrast (on a scale similar to that of the noise) and which then capture, on a larger scale, the statistical structure of the surface. For surfaces of this kind, the advantage of color is diminished and in some cases removed altogether. The integration that gives color its advantage seems to depend on the visual system's capacity to abstract an estimate of the mean from an array whose fine structure (on the scale of the individual elements) is very salient. The fact that the stimuli used here are spatially complex does not preclude the possibility that simple mechanisms can perform the segmentation task. For example, the existence of chromatic mechanisms that integrate over large areas can be well predicted from the low-pass chromatic contrast sensitivity function. The mechanisms of integration must arise at a relatively high level in the visual pathway. Neurons in the lateral geniculate nucleus and striate cortex have small receptive fields: In and within a few degrees of the fovea most of these integrate over regions no bigger than 1 deg and often much less.[27]

In discerning the structure of scenes, the visual system therefore seems to be better equipped to discount brightness variations than to discount color variations. This is useful in a world in which brightness variations are often accidental, resulting from shadows or highlights, and therefore need to be discounted when one is locating object boundaries by identifying differences in the reflectance of surfaces.[24]

ACKNOWLEDGMENTS

The authors thank Lawrence Bice and Alison Bice for serving as observers. This work was supported by National Institutes of Health grants EY 04440 and EY 01319 and was presented in part at the European Conference on Visual Perception (Tübingen, Germany, August 21–25, 1995) and at the annual meeting of the Association for Research in Vision and Ophthalmology (Fort Lauderdale, Florida, April 21–26, 1996).

*Present address and to whom correspondence should be sent: Department of Vision Sciences, SUNY State College of Optometry, 33 West 42nd Street, New York, New York 10036-8003; ali@sunyopt.edu.

†Present address: Center for Neural Science, New York University, 4 Washington Place, New York, New York 10003; pl@cns.nyu.edu.

REFERENCES

1. G. J. C. van der Horst, C. M. M. de Weert, and M. A. Bouman, "Transfer of spatial chromaticity contrast at threshold in the human eye," J. Opt. Soc. Am. **57**, 1260–1266 (1967).
2. R. Hilz and C. R. Cavonius, "Wavelength discrimination measured with square-wave gratings," J. Opt. Soc. Am. **60**, 273–277 (1970).
3. D. G. Green, "The contrast sensitivity of the colour mechanisms of the human eye," J. Physiol. **196**, 415–429 (1968).
4. O. Estévez and H. Spekreijse, "A spectral compensation method for determining the flicker characteristics of the human colour mechanisms," Vision Res. **14**, 823–830 (1974).
5. D. Varner, D. Jameson, and L. M. Hurvich, "Temporal sensitivities related to color theory," J. Opt. Soc. Am. A **1**, 474–481 (1984).
6. O. Estévez and C. R. Cavonius, "Flicker sensitivity of the human red and green color mechanisms," Vision Res. **15**, 879–881 (1975).
7. D. Regan and C. W. Tyler, "Temporal summation and its limit for wavelength changes: an analog of Bloch's law for color vision," J. Opt. Soc. Am. **61**, 1414–1421 (1971).
8. K. T. Mullen, "The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings," J. Physiol. (London) **359**, 381–400 (1985).
9. N. Sekiguchi, D. R. Williams, and D. H. Brainard, "Efficiency in detection of isoluminant and isochromatic interference fringes," J. Opt. Soc. Am. A **10**, 2118–2133 (1993).
10. P. Cavanagh, C. W. Tyler, and O. E. Favreau, "Perceived velocity of moving chromatic gratings," J. Opt. Soc. Am. A **1**, 893–899 (1984).
11. L. F. Van Sickle and W. S. Geisler, "Stereopsis in the absence of chromatic aberrations," Optics News **18**, 174 (1989).
12. J. R. Jordan III, W. S. Geisler, and A. C. Bovik, "Color as a source of information in the stereo correspondence process," Vision Res. **30**, 1955–1970 (1990).
13. E. Switkes, A. Bradley, and K. K. De Valois, "Contrast dependence and mechanisms of masking interactions among chromatic and luminance gratings," J. Opt. Soc. Am. A **5**, 1149–1162 (1988).
14. W. McIlhaga, T. Hine, G. R. Cole, and A. W. Snyder, "Texture segregation with luminance and chromatic contrast," Vision Res. **30**, 489–495 (1990).
15. F. A. A. Kingdom, B. Moulden, and S. Collyer, "A comparison between colour and luminance contrast in a spatial linking task," Vision Res. **32**, 709–717 (1992).
16. M. Gur and V. Akri, "Isoluminant stimuli may not expose the full contribution of color to visual functioning: Spatial contrast sensitivity measurements indicate interaction between color and luminance processing," Vision Res. **32**, 1253–1262 (1992).
17. J. D. Mollon, "'Tho' she kneel'd in that place where they grew,' The uses and origins of primate colour vision," J. Exp. Biol. **146**, 21–38 (1989).
18. J. Krauskopf, D. R. Williams, and D. W. Heeley, "Cardinal directions of color space," Vision Res. **22**, 1123–1131 (1982).
19. A. M. Derrington, J. Krauskopf, and P. Lennie, "Chromatic

mechanisms in lateral geniculate nucleus of macaque," J. Physiol. (London) **357**, 241–265 (1984).
20. A. Li and P. Lennie, "Mechanisms underlying segmentation of colored textures," Vision Res. **37**, 83–97 (1997).
21. R. T. Eskew, Jr., and R. M. Boynton, "Effects of field area and configuration on chromatic and border discriminations," Vision Res. **27**, 1835–1844 (1987).
22. D. H. Marimont and B. A. Wandell, "Matching color images: the effects of axial chromatic aberration," J. Opt. Soc. Am. A **11**, 3113–3122 (1994).
23. B. Julesz, H. L. Frisch, E. N. Gilbert, and L. A. Shepp, "Inability of humans to discriminate between visual textures that agree in second-order statistics—revisited," Perception **2**, 391–405 (1973).
24. E. Switkes and K. K. De Valois, "Luminance and chromaticity interactions in spatial vision," in *Color Vision*, J. D. Mollon and L. T. Sharpe, eds. (Academic, London, 1983), pp. 375–383.
25. V. C. Greenstein, D. Halevy, Q. Zaidi, K. L. Koenig, and R. H. Ritch, "Chromatic and luminance systems deficits in glaucoma," Vision Res. **36**, 621–629 (1996).
26. K. R. Gegenfurtner and D. C. Kiper, "Contrast detection in luminance and chromatic noise," J. Opt. Soc. Am. A **9**, 1880–1888 (1992).
27. K. H. Foster, J. P. Gaska, M. Nagler, and D. A. Pollen, "Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the macaque monkey," J. Physiol. (London) **365**, 331–363 (1985).