PERGAMON

Vision Research 42 (2002) 1421–1430

Vision Research

www.elsevier.com/locate/visres

Symmetry detection and visual attention: a “binary-map” hypothesis

Liqiang Huang [a], Harold Pashler [b,*]

[a] *Department of Optical Engineering, Zhejiang University, Hangzhou, 310027, China*
[b] *Department of Psychology 0109, University of California, San Diego, La Jolla, CA 92093, USA*

Received 19 December 2000; received in revised form 24 October 2001

Abstract

Recent research suggests that human symmetry-detection mechanisms cannot simultaneously compare different colors across the axis of symmetry (Nature 399 (1999) 115). In the present study, observers were required to judge symmetry in arrays composed of elements varying not only in color, but also in size, spatial frequency and orientation. In every case, response times increased with the number of different levels of a given feature. It is proposed that this increase reflects a sequential strategy whereby coarse “binary maps” are created by attentional filtering, and the symmetry of each map is then checked. Experiment 2 required observers to detect “pseudo-symmetry” (symmetry in feature values defined relative to an arbitrary featural boundary); the ease with which this task was accomplished supported the binary map hypothesis. The results suggest that (1) symmetry detection is spatially imprecise, and (2) attentional gating can operate prior to symmetry detection in the visual pathway.

Keywords: Symmetry; Binary map; Visual attention; Visual search

1. Introduction

Mirror symmetry is often highly salient to human observers (Barlow & Reeves, 1979), and the detection of symmetry may play an important role in mammalian vision in general. Some writers view symmetry as one of the most important aspects of early visual analysis (Wagemans, 1995). Recent research suggests, however, that human symmetry detection machinery may lack some capabilities one might naturally have attributed to it. Specifically, symmetry detection machinery appears “blind” to color in the sense that it cannot simultaneously compare the colors of various regions lying on opposite sides of the axis of symmetry (Morales & Pashler, 1999). [1]

Evidence for this somewhat counterintuitive claim came from tasks requiring people to judge symmetry in the arrangement of color in a regular grid pattern. Here the observer decided whether each and every square in the grid had the same color as the corresponding square located equidistant across the axis of symmetry, with asymmetric patterns differing in the color of only one or a few squares. The most obvious account of this kind of color-symmetry perception would be one that assumed that the colors of the pairs of opposing squares are checked simultaneously. This interpretation was challenged by several results. For one, response times (RTs) increased steeply as the number of different colors in the display was increased (holding constant the number of squares in the grid). This increase occurred even when the larger set of colors was selected so that every pairwise color difference within the set was greater than any difference within the smaller set.

This increase in RTs (several hundred milliseconds per additional color) was taken to favor an alternative account of color-symmetry detection. According to this alternative account, symmetry perception mechanisms per se are blind to color, and in order to perform the required task, observers must therefore use an indirect strategy. The strategy involves repeatedly “filtering” the display on the basis of color, extracting “subfigures” comprised of all the elements of a particular color and judging the symmetry of each subfigure. To make this

* Corresponding author. Fax: +1-619-534-7190.
E-mail address: hpashler@ucsd.edu (H. Pashler).

[1] This is rather different from the common claim that motion detection machinery is colorblind, meaning that it is unable to process purely chromatic (i.e., isoluminant) contours. Symmetry detection machinery is *not* color-blind in that sense (Troscianko, 1987).

0042-6989/02/$ - see front matter
PII: S0042-6989(02)00059-7

more concrete, consider how one might go about checking the color symmetry of a display consisting of red, green, blue and yellow squares. One could start by examining the red squares, determining whether the red subfigure is symmetric. Next one could check the green subfigure, and so on. If the subfigures are all symmetric, then the display as a whole must be color-symmetric. While seemingly laborious, this account explained a number of otherwise puzzling findings reported by Morales and Pashler (1999), several of which are discussed below.

In the present study, we asked whether a similarly laborious strategy might be necessary to detect not only color symmetry, but also symmetry in the spatial arrangement of other features. In the experiments reported below, observers assessed the symmetry of grids varying in a number of different ways. The displays were grids of regularly spaced items differing with respect to color, size, orientation or spatial frequency. When the elements varied in color, the task required judging color symmetry, as in Morales and Pashler (1999). When the elements differed in size or spatial frequency, however, the task could potentially be performed by assessing the fine-grained symmetry of the raw image within the luminance domain. To be effective, such an analysis would need to be very precise, however.

2. General method

2.1. Subjects

In Experiments 1–3, subjects were students from the Department of Optical Engineering in Zhejiang University, China. In Experiment 4, subjects were students from the University of California, San Diego. All had normal or corrected-to-normal vision.

2.2. Apparatus

Stimuli were presented on a high-resolution color CRT monitor driven by a PC. Responses were recorded from two adjacent keys on a standard keyboard. The observers viewed the displays from a distance of about 60 cm.

2.3. Stimuli

In symmetric displays, the entire display was symmetrical about a vertical axis bisecting the screen (the axis was not shown). Asymmetric displays were always created by starting with a symmetric display and altering it (the nature of the alteration differed between experiments, as described below).

2.4. Procedure

Displays were symmetric on one-third of trials, and asymmetric on two-thirds. Observers were informed about these frequencies. Each trial began with a small green fixation cross presented in the center of the screen. Observers were instructed to fixate the cross, which remained present for 700 ms. The cross was followed by a short blank interval (400 ms), which was in turn followed by the display. After the display appeared, observers responded by pressing one of two adjacent keys with fingers of right hand. A 'continue' message was presented after the observers had responded. The observer then pressed the space bar to start the next trial. They were told to respond as quickly and accurately as possible. A beep sounded whenever the observer made an error. In each experiment, observers were allowed to practice until they stated that they felt comfortable with the task (this took between 20 and 60 min depending on the subject).

3. Experiment 1

Experiment 1a–d examined the detection of symmetry in displays composed of elements that differed in color, size, orientation and spatial frequency, respectively. In each experiment, the displays were constructed out of elements that could assume either two or four levels of the feature that was varied. The values of the feature that varied were randomly assigned to items, with the constraint that each value was assigned to equal numbers of elements in the symmetrical display. The specific feature values employed were chosen such that every pair of choices in the four-choice display was more discriminable than the pair used in the two-choice display.

Suppose subjects perform the task by 'feature matching'—comparing the feature values of each grid element with the feature value of that item's reflection (the element at the corresponding position across the axis of symmetry). In that case, responses to the four-feature displays should be faster than to the two-feature displays, because the magnitudes of the discrepancies are greater in the former case. While the number of squares could obviously have an effect (depending on whether the comparison of different pairs is parallel or serial), the number of feature choices ought to make no difference. On the other hand, suppose subjects use the laborious strategy of sequentially filtering the display for one feature value after another, as Morales and Pashler (1999) proposed for color-symmetry judgments. In that case, the observer will first have to identify the subset of the display having one particular spatial frequency value (for example), then another spatial frequency value, and so on. The result should be a marked slowing when there are more feature choices.

In Experiment 1, the distance between the centers of adjacent items was 1.3 cm (1.24°), both vertically and horizontally. The display contained 32 items (four rows of eight). The overall extent of the display was approximately 5 cm (4.8°) high by 10 cm (9.6°) wide.

Experiment 1a was similar to the method used by Morales and Pashler (1999), who found that observers were slower to judge color symmetry in displays containing a greater number of colors (despite using larger color differences in this condition). However, the present experiment considered a factor neglected in the Morales and Pashler studies: that observers could potentially determine that a display was symmetric based upon noting that these boundaries in the two-color display were themselves symmetric, a strategy that would not work with the four-color display. Whereas Morales and Pashler attributed the number-of-colors effect to a sequential search through colors, it seems conceivable that the factor just mentioned could have produced the effect even if there were no serial search.

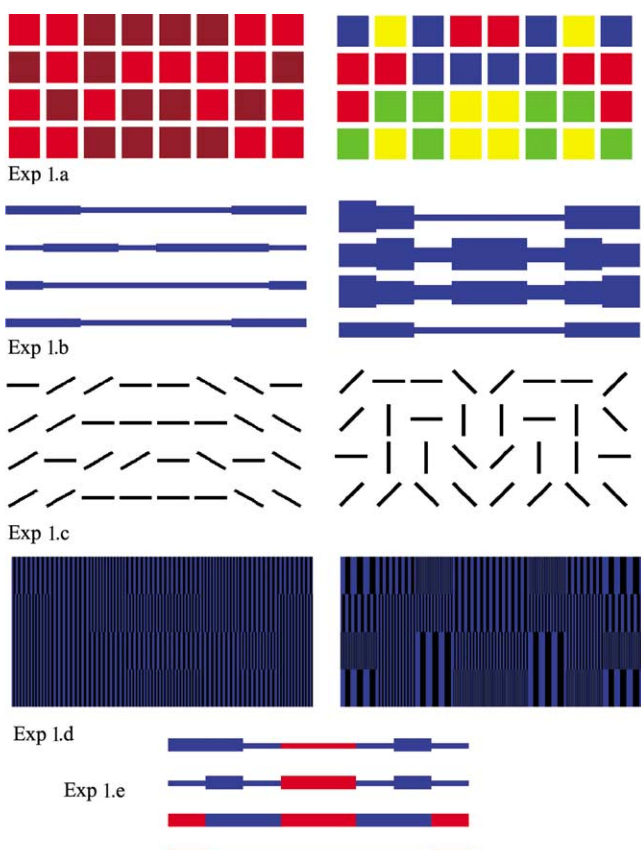

Fig. 1. Types of displays used in Experiments 1a–e. In Experiment 1, we measure the effect of increasing the number of feature values in several different featural dimensions: From top to bottom, they are: Experiment 1a (color experiment); Experiment 1b (size experiment); Experiment 1c (orientation experiment); Experiment 1d (spatial frequency experiment); and Experiment 1e (conjunction of color and size). In Experiment 1a–d, the left display illustrates a two-feature-value trial and the right display shows a four-feature-value trial. All displays shown in Fig. 1 are asymmetric.

To address these concerns, in Experiment 1a the colored regions were spatially separated from their neighbors as shown in Fig. 1. The gaps between regions measured 0.3 cm (0.3°). The size of the colored regions was 1 cm (1°), so detection of chromatic boundaries would not be useful (all displays contained boundaries throughout the grid).

As shown in Fig. 1, in the size experiment (Experiment 1b) all items were horizontal bars. The items were adjacent to their horizontal neighbors, but not to their vertical neighbors. In four-size displays, each item measured 1.07 cm, (1.02°) 0.78 cm, (0.74°) 0.49 cm, (0.47°) or 0.2 cm (0.19°). In two-size displays, each item was 0.29 cm (0.28°) or 0.18 cm (0.17°). In the orientation experiment (Experiment 1c), the stimulus items were lines pointing in varying directions. In four-orientation displays, each item was vertical, horizontal, 45° left-tilted, or 45° right-tilted. In two-orientation displays, each item was horizontal or tilted. In the spatial frequency experiment, (Experiment 1d) the items were vertical oriented square-wave gratings randomly chosen to be blue or black. (The color variation was introduced to reduce grouping). In four-spatial frequency displays, each item could have a spatial frequency of 19.2 line pair/cm, (20.1 line pair/°) 9.6 line pair/cm (10.1 line pair/°), 4.8 line pair/cm (5.0 line pair/°) or 2.4 line pair/cm (2.5 line pair/°). In two-spatial frequency displays, each item could have a spatial frequency of 9.6 line pair/cm (10.1 line pair/°) or 6.4 line pair/cm (6.7 line pair/°).

Unlike the other studies, Experiment 1e involved feature conjunctions. The items were conjunctions of size (0.18 cm (0.17°) or 0.44 cm (0.42°)) and color (red or green). The feature values in both dimensions were randomly assigned to items, with the constraint that each conjunction value was assigned to equal numbers (four) of elements in the symmetrical display. The value in one feature dimension of one item was altered to make asymmetric displays.

In each experiment, observers were given sufficient practice to become comfortable with the task, and then performed six experimental blocks of 40 trials.

3.1. Results

An ANOVA was used to assess the effect of number of feature choices. The results of Experiment 1a–d are shown in Table 1. In all four experiments, responses to the two-choice displays were substantially faster than responses to four-choice displays and the error rates were substantially lower, despite the fact that discriminability was always higher in the four-choice display.

If the dimensions are checked in parallel by the symmetry detection mechanism, we should expect that

Table 1
Mean reaction time and error rates for Experiments 1a–d

	Experiment 1a (color)	Experiment 1b (size)	Experiment 1c (orientation)	Experiment 1d (frequency)
RT (ms) (4 choices)	1528	1464	2549	1861
RT (ms) (2 choices)	970	889	1314	774
Difference of the RT	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
Error rate (4 choices)	0.12	0.17	0.21	0.11
Error rate (2 choices)	0.04	0.11	0.11	0.06
Difference of the error rate	$p < 0.01$	$p < 0.05$	$p < 0.01$	$p < 0.01$

the time to detect an asymmetry in one dimension (e.g. color dimension) to be unaffected by whether the display has the possibility of being asymmetric in another dimension or not. The discriminability of stimulus is higher in the conjunction experiment; thus, the detection should be more efficient. In the conjunction experiment (Experiment 1e), the mean RT was 1420 ms. Mean RT for color asymmetry in this experiment was 1293 ms, much slower than for asymmetric trials in two-color display in Experiment 1a (881 ms), $p < 0.005$. Mean RT for size asymmetric displays was 1188 ms, well in excess of the mean RT for asymmetric trials in two-size display in Experiment 1b (783 ms), $p < 0.005$. However, display elements in Experiment 1e were smaller than those in Experiment 1a. This might have slowed responses in Experiment 1e, and thus the difference cannot confidently be attributed to the multidimensional nature of the task. However, these large differences are consistent with the hypothesis that the symmetry detection system cannot process two dimensions at the same time, thus necessitating a shift from one dimension to another in the present experiment.

In these experiments, responses to asymmetric displays were always significantly faster than to symmetrical displays, and the error rate for symmetrical displays was significantly lower than for asymmetric displays. That is, observers sometimes missed an asymmetry whereas they were unlikely to report an asymmetry that was not present. These findings are in line with the typical results when asymmetries are created by a small perturbation of a symmetric display, as in this study.

3.2. Discussion

Over a variety of different kinds of feature variation, RTs were strongly affected by the number of possible values of a given feature. This effect swamped any beneficial effect provided by using more highly discrepant elements in the four-choice displays. One interpretation of these results is that in all cases, symmetry is assessed in what may be termed 'binary maps', i.e., internal representations that specify only the spatial distribution of a particular feature value, derived from an attentional filtering of the image. Each binary map divides the visual field into two categories: feature-positive and feature-negative. With four-choice displays, three (or possibly four) such maps would have to be constructed and checked before symmetry could be verified, whereas with two-choice displays, a single map would suffice.

The results also replicate the findings of Morales and Pashler (1999) for color symmetry detection, and strengthen those results by showing that number of colors still has a very large effect even when the individual elements do not abut each other, thereby eliminating the possible confound noted above.

What makes the present results more striking than those reported by Morales and Pashler (1999), however, is the fact that if the observer merely checked the fine-grained (e.g., pixel-level) symmetry of the gray-scale images, all the asymmetries presented in Experiments 1b, c and d should have been detectable. Why do people not carry out such a point-by-point analysis, saving themselves the trouble of constructing and assessing a series of binary maps? One possibility is that the symmetry detection machinery simply lacks the spatial acuity for such a detailed comparison. The sequential filtering strategy allows symmetry judgments to be computed at a very coarse level (display elements rather than pixels). The claim that symmetry detection machinery is coarse might seem to conflict with the claims of Barlow and Reeves (1979), who reported that the mechanism for achieving symmetry is versatile and efficient. However, our task required a more complete analysis of the display than did their task. In Barlow and Reeves' study, asymmetric displays were composed of randomly positioned dots, whereas in our study the asymmetric displays contained only a single pair of discrepant items. Thus, in their task the observer might analyze only a small proportion of the spatial information present in the display.

4. Experiment 2

While the observation that it takes longer to verify the symmetry of displays containing a larger number of feature values is consistent with the sequential filtering theory described above, other explanations for that effect might also be proposed. The binary map idea makes additional, more distinctive predictions. If each grid element is compared to a standard in the construction of the binary image (the "filtering" referred to above), then perhaps subfigures need not even be composed out of identical elements. For example, if a binary map can be created that includes all elements darker than some standard (say, mid-gray), then black and darker gray could both be categorized as feature-positive, while light gray and still lighter gray could both be categorized feature-negative.

This suggests a strong test of the theory by presenting what might be termed "pseudo-symmetric" patterns. If people are able to filter in the flexible way described, then they should have no difficulty perceiving the sort of "pseudo-symmetry" illustrated in Fig. 2 (Panels A, B, C, and D) despite the complete absence of true symmetry. On the other hand, if the visual system possesses mechanisms that determine whether precise symmetry is present in the gray-scale image, the task should be quite difficult.

In Experiment 2, the distance between the centers of adjacent items measured 1.3 cm (1.24°), both vertically and horizontally. The whole display contained 32 items (four rows of eight). The overall extent of the display was approximately 5 cm (4.77°) high by 10 cm (9.55°) wide.

Experiment 2a–e were designed to test the prediction described above. As shown in Fig. 2, in these experiments, displays were never truly symmetric; instead, they were either pseudo-symmetric (to be categorized "symmetric") or pseudo-symmetric with an alteration (categorized "asymmetric" for purposes of the task). Pseudo-symmetry was created as follows. First, each item in the display was randomly assigned to one of two feature classes (call them A and B), with the constraint that equal numbers of elements were assigned to these classes in the symmetrical display. The assignment of one item to a class was altered if the display was to be asymmetric. In Experiments 2a, b and e, four colors were used to construct the displays. Items in class A and B situated on the left side of the display were assigned to two of these colors (one color for A, the other for B), while items in the others were assigned the remaining two colors. The RGB values of the colors used in the Experiment 2a were, left side: bright-blue-green (128, 255, 255) and red; right side: rose (255, 0, 128) and black. The colors used in the Experiment 2b are: left part: dark gray and white, right part: black and white gray. The colors used in the left part and right part of Experiment 2e were reversed: left part: rose, dark-green-blue (0, 128, 128), right part: dark-green-blue, rose. Different from Experiment 1, the squares in Experiment 2a b and e are adjacent to each other. In Experiment 2c and d, we used a symmetrical stimulus and an asymmetric stimulus. In Experiment 2c, the symmetrical stimulus was a horizontal line, and the asymmetric stimulus

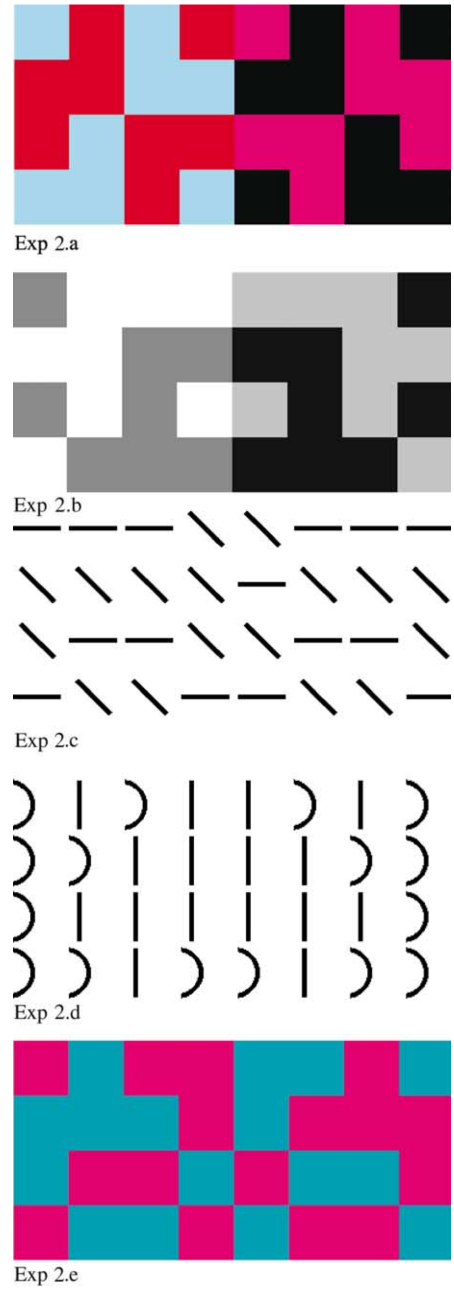

Fig. 2. Types of displays used in Experiments 2a–e. In Experiment 2 we measure the efficiency of "pseudo-symmetry" detection. From top to bottom, they are: Experiment 2a and b: items varied in color dimension and attentional filtering is possible. Experiment 2c items varied in orientation dimension and attentional filtering is possible. Experiment 2d items varied in curvature and attentional filtering is possible. Experiment 2e items varied in color dimension and attentional filtering is impossible.

was a tilted line. In Experiment 2d, the symmetrical stimulus was a vertical oriented bar, and the asymmetric stimulus was a right tilted curve.

Assuming the binary map account is correct, then in four of the experiments (2a–d), the observers should be able to find a standard with which they can derive a symmetrical binary map for all and only the pseudo-symmetric displays. In Experiment 2e, on the other hand, no appropriate standard is available because each element and its reflections were opposing elements in the same set. Here, it was predicted, the judgment of symmetry should be highly inefficient.

In each of the experiments the observers began with as many practice trials as they required to become comfortable with the task, followed by six blocks of 40 trials.

4.1. Results

In Experiment 2a, b and e, four colors were given in the display. In Experiment 2a and b, the pseudosymmetric displays can be converted to symmetric binary maps by selecting an appropriate standard within color space. Indeed, performance was generally excellent (Table 2). By contrast, mean RTs in Experiment 2e were more than twice as long as in any of the other studies. In Experiment 2c and d, an appropriate standard for orientation or curvature, respectively, was available; again performance was quite good.

4.2. Discussion

The results show that observers can readily judge "pseudo-symmetry" when comparison of element features to a standard would transform pseudo-symmetric displays into a symmetric binary map. Obviously a true point-by-point symmetry comparison would judge all these displays asymmetric. The fact that observers could assess pseudo-symmetry almost as fast as true symmetry in Experiment 1 further supports the conclusion that the proposed sequential filtering process is the means by which all these tasks are performed. One might suggest that pseudo-symmetry in Experiment 2 could be detected by checking the symmetry of the boundaries separating adjacent elements. This fails to account for the poor performance in Experiment 2c, however.

5. Experiment 3

We have hypothesized that the machinery of symmetry detection may be limited in its precision or processing capability (beyond the constraints imposed by front-end visual acuity). As noted earlier, this suggestion might seem to conflict with other researchers' conclusions that symmetry detection is versatile and highly efficient (Barlow & Reeves, 1979). As noted above, however, the tasks used here differ in important ways from those used by Barlow and Reeves (1979). In the studies described here, observers were required to make an exacting judgment of symmetry, spotting mismatches in the properties of individual elements. In Experiment 3, we explicitly examined the effect of the number of elements whose symmetry must be judged in a simpler symmetry-detection task.

In Experiment 3, as shown in Fig. 3, yellow squares were embedded in a dark green background. A variable number of squares (6, 10, 20, 30, 40, 50, or 60) appeared in randomly chosen positions within an array of 400 locations (20×20). Thus, the average eccentricity of a square did not vary with the number of squares in the display. Each square measured $0.16\ \text{cm} \times 0.16\ \text{cm}$ ($0.15°$) and the distance between neighbors was 0.16 cm ($0.15°$) vertically and horizontally. The overall extent of the display was approximately 6.3 cm ($6.0°$) high by 6.3 cm ($6.0°$) wide. The number of squares was chosen independently on each trial. The asymmetric trials were created by removing a square ($p = 0.5$) or adding a new square ($p = 0.5$). (As a result, the actual number of squares differed by one from the nominal display set size referred to above.)

The observers each began with sufficient practice trials, followed by six blocks of 40 trials.

5.1. Results and discussion

The results of Experiment 3 are shown in Fig. 3. Both RTs and errors rose rapidly with the number of squares, with errors rates approaching 50% with the larger numbers of squares. Plainly, in this task requiring detection of asymmetries involving perturbation of just a single item, symmetry detection machinery is readily overwhelmed by large numbers of stimuli. This fits with the suggestion raised earlier to explain why observers may choose to analyze (attentionally filtered) binary images in preference to trying to assess the symmetry of the point-by-point images.

The results of Experiment 3 suggest that in previous studies involving very large arrays of dots with asymmetric displays that were random (rather than symmetric with a perturbation), judgments of symmetry may have been based on coarse features occasionally and serendipitously generated by large clusters of dots,

Table 2
Mean reaction time and error rates for Experiments 2a–e

	Experiment 2a	Experiment 2b	Experiment 2c	Experiment 2d	Experiment 2e
Response time (ms)	806	873	1186	1394	2596
Error rate	0.06	0.04	0.11	0.13	0.19

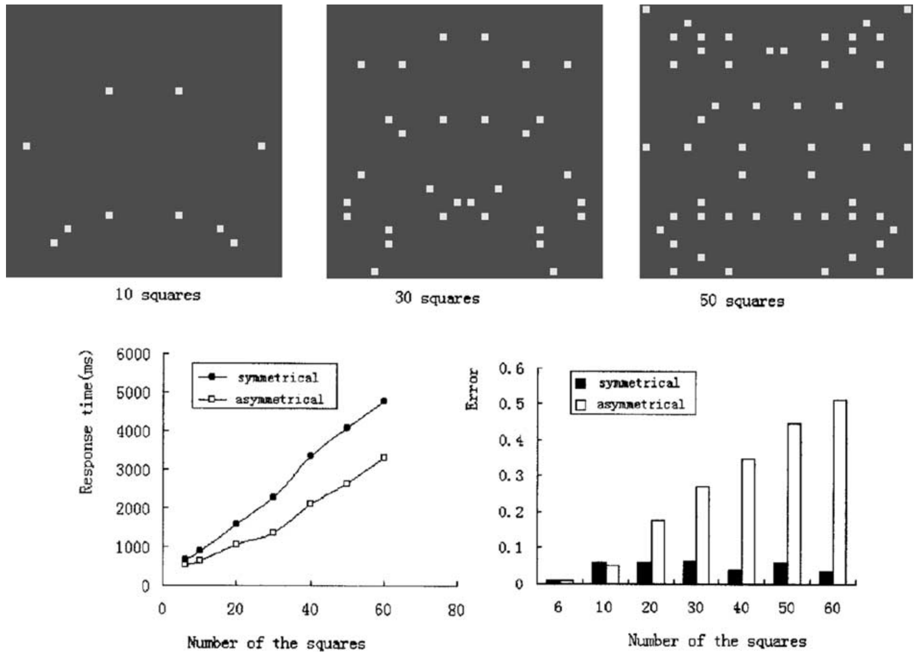

Fig. 3. Types of displays used in Experiment 3. In Experiment 3, we measure the set-size-effect of dots of a symmetry detection, along with mean RTs and error rates averaged over subjects.

rather than a comparison of individual points. It should be noted, however, that our results do not, strictly speaking, imply that symmetry detection is capacity limited (i.e., that the assessment of symmetry of any one dot pair is less efficient when the number of simultaneous comparisons is increased). If all dot pairs were analyzed with some fixed degree of accuracy or efficiency, increasing the number of dot pairs would be likely to increase the probability of an error. A strong test of capacity limitations in symmetry detection would require a different type of experiment. However, for present purposes the results are illuminating, suggesting that observers cannot make rapid and precise symmetry judgments on displays comprised of many elements.

6. Experiment 4

The result of Experiment 3 indicated that the performance was much worse when the number of dots was increased. This raised an alternative explanation for the result of Experiment 1: the elements may be perceptually grouped into chunks. Increasing the number of possible feature values would make these chunks more complex, and increase the number of chunks (because the probability of large groupings would be reduced). So if the symmetry is actually compared "chunk by chunk", the RTs will also increase with the number of possible features. Reasonably the grouping will play a role in these experiments.

Morales and Pashler (1999) offered converging evidence for the binary map hypothesis, in the form of an experiment using what they termed the "ABBA–ABCD" method. Here, they compared responses to two different types of four-color displays in which two squares mismatched. In the ABBA condition, the two mismatched pairs were restricted to two colors, whereas in the ABCD condition, the two mismatched pairs involved all four colors. Responses were faster in the ABCD condition, where all four colors were involved. This would be predicted by the hypothesis of sequential scanning through colored subfigures, because any mismatch would become evident in the first subfigure checked. The effect cannot be easily explained by the "chunk by chunk" account or any other account focusing on the role of complexity, because the spatial complexity (number of chunks) is the same in the two kinds of trials.

In Experiment 4, we attempt to change the "starting color" of the presumed sequential checking process by presenting a "prime"—a patch that matches one of the colors in the display that follows. It seemed to us likely that such a prime would determine the initial color checked in any sequential checking process. If so, responses should be faster when the color of prime is the same as one of the two "mismatching" colors in the display. By contrast, a mismatching prime should slow detection of an asymmetry. The "complexity" account would not predict any such effect.

In this experiment, as shown in Fig. 4, observers were required to judge if the whole display (both the prime

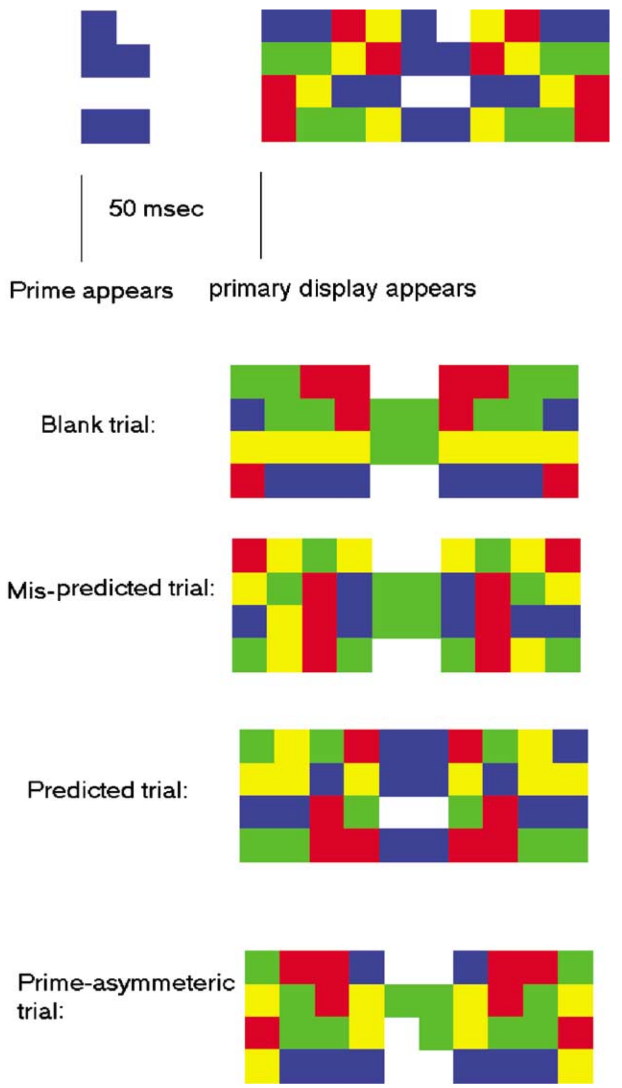

Fig. 4. Types of displays used in Experiment 4. Uninformative prime changes the efficiency of symmetry detection. From the top to bottom, they are: temporal sequence of prime and primary display. Samples illustrate blank trial, mis-predicted trial, predicted trial, and prime-asymmetric trial.

and the subsequent "primary display") was symmetric. The display was symmetric on one third of the trials; on another third of the trials there was an asymmetry in the prime (consisting of an unpaired colored square); on a final third of the trials, there was an asymmetry in the primary display, namely a pair of mismatched squares. In the latter case, the two colors of the pair of mismatched squares might include the color of the prime (1/6 of trials) or not (the remaining 1/6 of trials). The prime was red, green, blue or yellow (randomly chosen). If the prime was symmetric, it consisted of two or three pairs of squares; asymmetric primes were created by removing or adding one square without introducing any new colors. (The remaining squares in the prime region were the same color as the background.) The primary display (columns 1–4 and 7–10) was similar to four-color trials in Experiment 1a except in two respects. First, all squares were adjacent to their neighbor, and second, the two halves of the primary display were separated to allow a region two squares in width (2.5 cm), which was occupied by the prime. The prime appeared 50 ms before the primary display.

In Experiment 4, the distance between the centers of adjacent items measured 1.3 cm (1.24°), both vertically and horizontally, the whole display contained 40 items (four rows of ten). The two columns that are next to the axis could be blank (the same color as the background), as we will discuss in details in Experiment 4. The overall extent of the display was approximately 5 cm (4.77°) high by 12.5 cm (11.94°) wide.

Subjects were given enough practice to familiarize themselves with the task, followed by five blocks of 60 trials. The importance of accurate responding was emphasized in this experiment (without this emphasis, subjects tended to perform the symmetry judgment on the prime accurately, but to miss many asymmetries in the primary displays).

6.1. Results and discussion

RTs for asymmetric trials (where the asymmetry was in the primary display) were sorted according to whether the prime "predicted" the mismatch color (termed "predicted trials") or did not predict it (termed "mispredicted trials"). The mean RT for predicted trials was 2273 ms (error rate $= 0.09$) and for mispredicted trials, 2571 ms (error rate $= 0.13$), with significant differences for both RTs ($p < 0.01$) and error rates ($p < 0.05$).

The results are plainly consistent with the hypothesis that different colors are checked sequentially. Are they consistent with the alternative hypothesis that all of the colors are checked simultaneously? If so, one would have to suppose that the priming effect operates by speeding up processing of the portion of the display sharing the same color as the prime. There are at least two difficulties for such an account. First, the size of the color priming effect observed here (298 ms) seems much bigger than one could attribute to a speedup in perceptual processing studies using rapid serial visual presentations of colored displays suggest comprehension in less than 200 ms (Biederman, Mezzanotte, & Rabinowitz, 1982). Second, several studies question whether color priming produces any detectable facilitation in grouping and other early visual processing, even with primes that are—unlike those in the present study—informative about the properties of the relevant stimulus (Moore & Egeth, 1998; Posner, Snyder, & Davidson, 1980; Shih & Sperling, 1996).

Thus, the results of Experiment 4 provide converging evidence that color-symmetry detection is performed serially from color to color. Along with the ABBA–ABCD experiment described above (Morales & Pashler, 1999), the results speak against an alternative "chunk by chunk" account. Naturally, this does not imply that spatial complexity plays no role in the effect of number of colors reported in Experiment 1 above, and by Morales and Pashler (1999).

7. General discussion

The experiments described here suggest that perception of visual symmetry in grid patterns can be based on a sequential filtering of the image. This filtering is evidently flexible: it can "extract" the portions of the display containing an arbitrary feature value, or even different feature values situated on the same side of an arbitrarily chosen boundary. Observers seem to employ this strategy not only in detecting color symmetry (as previously suggested by Morales & Pashler, 1999; and confirmed in Experiment 1a above) but also when other local stimulus attributes such as spatial frequency and orientation are varied. Though these strategies seem peculiar at first blush, they make sense if one makes two key assumptions. The first is that human symmetry detection machinery is very limited in its precision, as suggested by the results of Experiment 3 (see also Huang and Pashler, under review). The second is that attentional mechanisms are flexible and capable of filtering sensory inputs at a point (or perhaps several points) along the visual pathway *prior* to the locus of symmetry detection (or at the very least that attentional filtering is capable of influencing the data upon which symmetry detection operates).

These conclusions are broadly similar to other recent findings demonstrating differences between symmetry perception and other basic visual mechanisms (Barrett, Whitaker, McGraw, & Herbert, 1999; Tyler, 1999). They are also specifically congenial to the proposals of Sperling and colleagues (e.g., Lu & Sperling, 1995a,b; Sperling & Lu, 1998) regarding a phenomenon they termed "third-order motion" perception. These authors postulated that motion perception can be computed directly on a "salience map" that is jointly determined by stimulus strength (bottom–up processes) and by selective attention (top–down processes). They term this kind of motion processing "third-order motion perception".

In one set of studies, Lu and Sperling (1995a) cleverly designed dynamic displays so that the observers' choice of which particular feature values to attend to (orientation and spatial frequency) determined the perceived direction of apparent motion. In essence, Sperling and Lu showed that motion perception can be based on the spatiotemporal distribution of attentionally selected features in the display. Analogously, the results described here indicate that symmetry perception can be based on the spatial distribution of attentionally selected features in the display. One can highlight the analogy between the two phenomena by constructing displays in which symmetry perception is biased in a way that is closely analogous to the motion biasing effect observed by Sperling and Lu. Consider Fig. 5. The display as a whole is not symmetric. However, if one chooses to attend to the red items in the figure, the symmetry about a horizontal axis is readily perceived, whereas if one attends to the green items, one perceives the symmetry about a vertical axis.

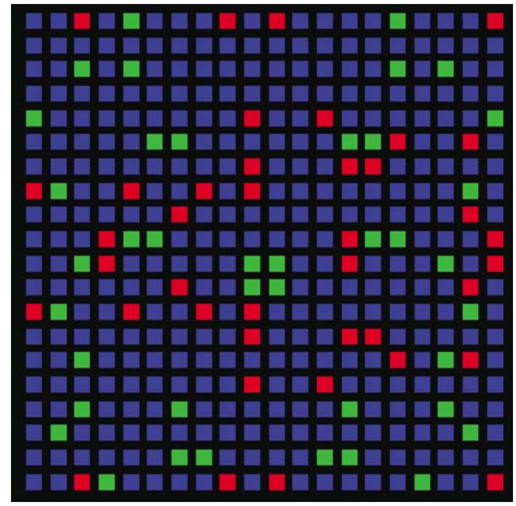

Fig. 5. A demo to show that "symmetry" can be different for the same stimulus according to different attentional setting. The observers will report a vertical symmetry if they are instructed to attend to the green pattern, and horizontal symmetry if they are instructed to attend to the red pattern.

To use the terminology of Sperling and Lu (1998), the findings reported here can be summarized by saying that when people try to judge the symmetry of regular grids composed of items varying along basic visual dimensions, they perform the task by decomposing it into a series of third-order symmetry judgments. Thus, in addition to shedding light on the nature of symmetry perception, these results reinforce the emerging consensus that the effects of selective attention can arise relatively early in the visual pathways.

The experiments are also relevant to what is sometimes termed the "binding problem" in perception. According to Treisman's "feature integration theory" (Treisman & Gelade, 1980), different features like color, orientation and motion are represented in separate maps, and these features can only be conjoined by limited-capacity attentional mechanisms. According to this theory, any task that requires binding information from more than one dimension will require focal attention. The great majority of research relating to this theory has used visual search tasks.

What has been less discussed is the fact that the theory would have implications for visual processing that go beyond visual search. At a minimum, the theory implies that any task that requires a judgment based on the spatial distribution of different features in a display (which requires keeping the different features straight) cannot be performed in parallel across both feature types and spatial locations.

The results described here (suggesting that the visual system cannot base a symmetry judgment on the simultaneous binding of different feature values to their respective locations) seem to fit with the theory. On the other hand, we would not suggest that the judgment of the symmetry of each of the colored subfigures requires a sequential checking of individual display elements in the subfigure (although the results described here do not rule that out). In the realm of visual search, too, the idea of sequential checking of all display elements has been challenged. Many investigators have concluded that detection of conjunction targets is more efficient than one would expect based on the simplest version of the theory, which proposed that display elements must be examined one by one (Wolfe, 1994, 1998). In summary, it may be the case in both conjunction search and multidimensional symmetry detection, that the underlying limitations on binding are roughly as described by Treisman and Gelade (1980), but observers have available various strategies for performing such tasks which sometimes tend to mask these binding limitations.

Acknowledgements

This work was supported by the National Science Foundation (SBR #9729778) and the National Institute of Mental Health (MH45584).

References

Barlow, H. B., & Reeves, B. C. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research, 19*, 783–793.

Barrett, B. T., Whitaker, D., McGraw, P. V., & Herbert, A. M. (1999). Discriminating mirror symmetry in foveal and extra-foveal vision. *Vision Research, 39*, 3737–3744.

Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and Judging objects undergoing relational violations. *Cognitive Psychology, 14*, 143–177.

Lu, Z.-L., & Sperling, G. (1995a). Attention-generated motion perception. *Nature, 377*, 237–239.

Lu, Z.-L., & Sperling, G. (1995b). The functional architecture of human visual motion perception. *Vision Research, 35*, 2697–2722.

Moore, C., & Egeth, H. (1998). How does feature-based attention affect visual processing? *Journal of Experimental Psychology: Human Perception & Performance, 24*, 1296–1310.

Morales, D., & Pashler, H. (1999). No role for colour in symmetry perception. *Nature, 399*, 115–116.

Posner, M. I., Snyder, C. R. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General, 109*, 160–174.

Shih, S., & Sperling, G. (1996). Is there feature-based attentional selection in visual search? *Journal of Experimental Psychology: Human Perception & Performance, 22*, 758–779.

Sperling, G., & Lu, Z.-L. (1998). A systems analysis of visual motion perception. In T. Watanabe (Ed.), *High-level motion processing: Computational, neurobiological, and psychophysical perspectives* (pp. 153–183). Cambridge, MA: MIT Press.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97–136.

Troscianko, T. (1987). Perception of random-dot symmetry and apparent movement at and near isoluminance. *Vision Research, 27*, 547–554.

Tyler, C. (1999). Human symmetry detection exhibits reverse eccentricity scaling. *Visual Neuroscience, 16*, 919–922.

Wagemans, J. (1995). Detection of visual symmetries. *Spatial Vision, 9*, 9–32.

Wolfe, J. M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review, 1*, 202–238.

Wolfe, J. M. (1998). Visual search. In H. Pashler (Ed.), *Attention*. London UK: University College London Press.