

Theory

Search asymmetries? What search asymmetries?

RUTH ROSENHOLTZ
Xerox PARC, Palo Alto, California

In order to establish a search asymmetry, one must run an experiment with a symmetric design and get asymmetric results. Given an asymmetric design, one expects asymmetric results, and such results do not imply an asymmetry in the search mechanisms. In this paper, I argue that a number of experiments purporting to show search asymmetries contain built-in design asymmetries. A saliency model of visual search predicts the results of these experiments, using only a simple measure of target-distractor similarity, without reliance on asymmetric search mechanisms. These results have implications for search mechanisms and for other experiments purporting to show search asymmetries.

Our lives are filled with activities that one may think of as *visual search*. We search our bookcases for a particular book, look for our keys, and try to find a friend's face in a crowd. In studying visual search, an experimental subject typically looks for a *target* item among some number of other items, referred to as *distractors*. In some cases, the number of distractors has little effect on the ease of finding the target (*efficient search*), whereas in other cases, each additional distractor makes the search task significantly more difficult (*inefficient search*; see Wolfe, 1996, for a review of the literature).

Many researchers have studied visual search as a means toward learning about mechanisms in the visual system. Researchers have taken efficient search as evidence of the existence, in the visual system, of a basic feature detector that responds to the features of the target, but not to the features of any of the distractors. Another paradigm uses visual search experiments as evidence that indicates whether or not the visual system codes two feature dimensions independently (Driver, McLeod, & Dienes, 1992; Treisman, 1988).

Within this framework, researchers have traditionally considered search asymmetries particularly interesting and fortuitous. Suppose that search for an item containing a particular feature, *t*, among distractors with a different feature, *d*, is efficient but that the reverse is inefficient. Researchers have taken such an asymmetry as evidence that a basic feature detector that responds to feature *t* exists in the visual system, but that a detector for feature *d* is absent. This seems to suggest a particular importance

for feature *t* in the visual system and a lack of importance for feature *d*.

However, there are a number of reasons to study search not as a means to uncover visual system mechanisms, but for its own sake, or to learn about visual attention. People interested in user-interface design want to know how to draw attention when it is desired and not to annoy or mislead observers when attention is not required. Researchers studying cockpit design want pilots to notice warning lights and easily scan their instruments. In image compression, regions of an image that do not draw attention perhaps can be more heavily compressed without a noticeable reduction in image quality.

Search asymmetries may be interesting to researchers hoping to reveal visual system mechanisms, but they make modeling visual search more complicated, since each new experiment potentially adds a new component to the model. Traditional accounts of visual search (e.g. Dick, 1989; Driver et al., 1992; Ivry & Cohen, 1992; Nagy & Cone, 1996; Treisman & Gelade, 1980; Treisman & Gormican, 1988) suggest that it is best modeled by a list of feature detectors and associated search asymmetries. These accounts suggest that a visual search model requires red-not-pink and moving-not-stationary detectors, but no pink-not-red or stationary-not-moving detectors, and so on. Such models are more complicated than models in which a single measure of similarity between the target and the distractors describes performance. Furthermore, if visual search is best described by a list of asymmetries, one has to run many experiments to ferret out all the asymmetries, in order to get a complete, rather complicated model.

Simpler models are desirable for applications, as well as, in certain cases, for the study of human vision. From this point of view, one would prefer models of visual search based on a measure of similarity between target and distractors, rather than models that enumerate search asym-

Thanks to Bart Anderson for useful discussions on asymmetries in feature domains other than motion, color, and orientation. Correspondence concerning this article should be addressed to R. Rosenholtz, Xerox PARC, 3333 Coyote Hill Rd., Palo Alto, CA 94304 (e-mail: rruith@parc.xerox.com).

metries. However, perhaps this is simply not the way the visual system works. Perhaps visual search performance is truly best described by a list of asymmetries, in which case one must make do with a complicated model.

The next section of this paper argues that many classic search asymmetries do not necessarily reveal the existence of asymmetric visual search mechanisms. To obtain evidence of an asymmetric mechanism, one must run an experiment with a symmetric design and get asymmetric results. However, many experiments purporting to show search asymmetry are asymmetrically designed. The succeeding section offers a model that can qualitatively predict many known results by considering target-distractor similarity, with no need for asymmetric search mechanisms.

This reinterpretation of these search results serves as a reminder of the following caution about studying visual system mechanisms through psychophysical experiments: One's conclusions about mechanisms are only as good as one's underlying model. With one model, the visual system appears to contain numerous asymmetric search mechanisms. Another model calls into question many of those asymmetries. What did we learn, then, about mechanisms?

The next section discusses the requirements for a symmetric experiment, reexamines a number of motion, color, and orientation search experiments, and shows that they may be viewed as containing built-in design asymmetries. The succeeding section presents a measure of target-distractor similarity and shows that it qualitatively predicts the results of these experiments without need for any asymmetric search mechanisms. Although this paper calls into question a number of search asymmetries involving simple features, such as motion, color, and orientation, there may yet be search asymmetries involving more complex features, such as shape, line length, and three-dimensional (3-D) appearance. The final section suggests ways in which the analysis in this paper might be extended to deal with search asymmetries involving more complex features.

Asymmetric Experimental Designs

An expression is said to be symmetric if parts of it may be interchanged without changing the whole (Daintith & Nelson, 1989). In this section, it is asked whether this condition was met by a number of experiments intended to demonstrate search asymmetries.

What does a symmetric experimental design look like? A pair of experiments has a symmetric design if it preserves relative differences between features and changes only absolute feature values. The simplest way to visualize this is to look at the experiments in some appropriate uniform feature space, such that the Euclidean distance between two points in the feature space determines the discriminability of the two features. A uniform feature space is one in which, if one plots contours of equal discriminability (e.g., what colors are equally discriminable from mid-gray), they form circles, and the circles are the same size throughout the space. Plotting all of the elements of the experimental stimuli in this feature space, a pair of experiments has a symmetric design if the feature-space representations differ only by a rotation, reflection, and/or translation. Figure 1 shows what a pair of experiments with a symmetric design might look like in feature space.

This paper enumerates three classes of asymmetric designs. Some search experiments have obvious design asymmetries. Other designs appear symmetric or asymmetric, depending on feature space in which one views the experiments. In such cases, one can make an argument for or against asymmetry on the basis of independent knowledge of the appropriate feature space. In the absence of such knowledge, one might choose the feature space that yields the more parsimonious explanation of the results. Given the same explanatory power, the more parsimonious explanation of the results is the one that is simpler—that is, its model has fewer parameters, where each asymmetry essentially counts as at least one parameter. Finally, a number of designs initially appear symmetric but become asymmetric once one considers the background or frame of reference for the experiment. A




Figure 1. A sample pair of experiments with symmetric design, represented in a uniform feature space. A design is *symmetric* if a distance-preserving transformation (reflection, rotation, and/or translation) takes the feature-space representation of one experiment into the other experiment of the pair. A disk (●) represents the target, and a cross (×) represents each type of distractor.

pair of experiments may have symmetrically switched the target and distractors but not properly switched the background or frame of reference, thus changing the relationship of the target to the background in an asymmetric way.

Case 1: A clearly asymmetric design. In one classic motion search asymmetry, it has been said that search for a moving target among stationary distractors is efficient (Dick, Ullman, & Sagi, 1987), whereas search for a stationary target among moving distractors is inefficient (Dick, 1989). This has been taken as evidence for a special status for motion detectors in the visual system.

However, consider how this effect has typically been tested and demonstrated (e.g., Dick, 1989): In the case of search for a stationary target, at any given instant in time, the observers see distractors moving in a number of different directions. This is hardly symmetric to a condition in which, at a given instant, the target moves in a single direction, whereas the distractors do not move. Figure 2 shows what the two experiments look like, plotted in two-dimensional (2-D) velocity space.

The obvious, more symmetric design would be to compare search for a moving target with search for a stationary target among distractors that coherently move in a single direction. Royden, Wolfe, and Klempen (2001) have run this experiment, and we will discuss it later in this paper.

Case 2: Symmetry depends on the proper feature space. Two other classic motion search results have also been claimed to have demonstrated search asymmetries (Driver et al., 1992; Ivry & Cohen, 1992). These experiments appear to have a symmetric design when viewed in one feature space but appear to have built-in design asymmetries when viewed in another. This paper argues that the most parsimonious feature space for these experiments is the one in which the experimental designs ap-

pear asymmetric and that, thus, the experiments should not be used to demonstrate search asymmetries.

Ivry and Cohen (1992) reported efficient search for a quickly oscillating target among slowly oscillating distractors but inefficient search for the more slowly oscillating target. We depict this situation in Figure 3. Viewed in Cartesian speed–direction space, the design of these two experiments appears symmetric. However, viewed in 2-D velocity space, the design appears asymmetric. No rotation, translation, or reflection transforms one experiment into the other.

Another set of experiments have been claimed to show the asymmetry that motion direction is coded independently of speed but speed is not coded independently of direction. In this experimental paradigm, observers search for a target item with a unique value along some *relevant* feature dimension, while experimenters add variability along an *irrelevant* feature dimension. If the added variability impairs search for the target, this is taken as evidence that the two feature dimensions are not coded independently in the visual system. If the added variability has no effect on the search task, researchers take this to signify independent coding of the two feature dimensions.

Driver et al. (1992) tested motion speed and direction as the relevant and irrelevant features. Adding variability in motion direction impaired search for a unique speed (the speed task). However, adding variability in motion speed did not impair search for a unique motion direction (the direction task). The authors concluded that speed is not coded independently of direction but that direction is coded independently of speed.

In this case, the relevant thing to check is whether the transformation, in going from a homogeneous search condition to a heterogeneous one, is the same in the speed task as in the direction task. Viewed in Cartesian speed–




Figure 2. Search for a moving (A) versus stationary (B) target in feature space. Note that this pair of experiments does not appear to have a symmetric design. The inset figures depict cartoon versions of the stimuli, with the targets and the distractors represented by disks (●) and crosses (×), respectively, and arrows indicating motion. In the actual displays, target and distractors had the same shape. No attempt has been made to replicate the number, size, or density of display elements in these cartoon representations.




Figure 3. Search for a slowly oscillating target (A, C, E) versus search for a quickly oscillating target (B, D, F). In speed–direction space (C, D), the design appears symmetric, but not in two-dimensional velocity space (E, F). Cartoon versions of the stimuli are shown in panels A and B. A disk (●) represents the target, and a cross (×) represents each type of distractor.

direction space, these two pairs of experiments seem to have a *nearly* symmetric design, as is shown in Figure 4. The transformation from Figure 4C to Figure 4D is similar to the transformation from Figure 4A to Figure 4B. The different numbers of distractor classes and difficulties in relating the scales of direction classes and speed detract from the symmetry of the design. Strictly speaking, experiments should have completely symmetric design, but one might argue that, viewed in speed–direction space, these experiments might be close enough to symmetrical to safely draw conclusions about search asymmetries.

However, viewed in 2-D velocity space (Figure 5), the experiments look quite different. In this space, adding variability in direction changes the experiment in a different way than adding variability in speed. Thus, it is not surprising that added variability impacted the speed task in a different way than it impacted the direction task. If 2-D velocity space is the correct space in which to view these experiments, that itself implies that speed is not encoded independently of direction—one of the conclusions drawn by Driver et al. (1992). The interpretation that their experimental design was asymmetric calls

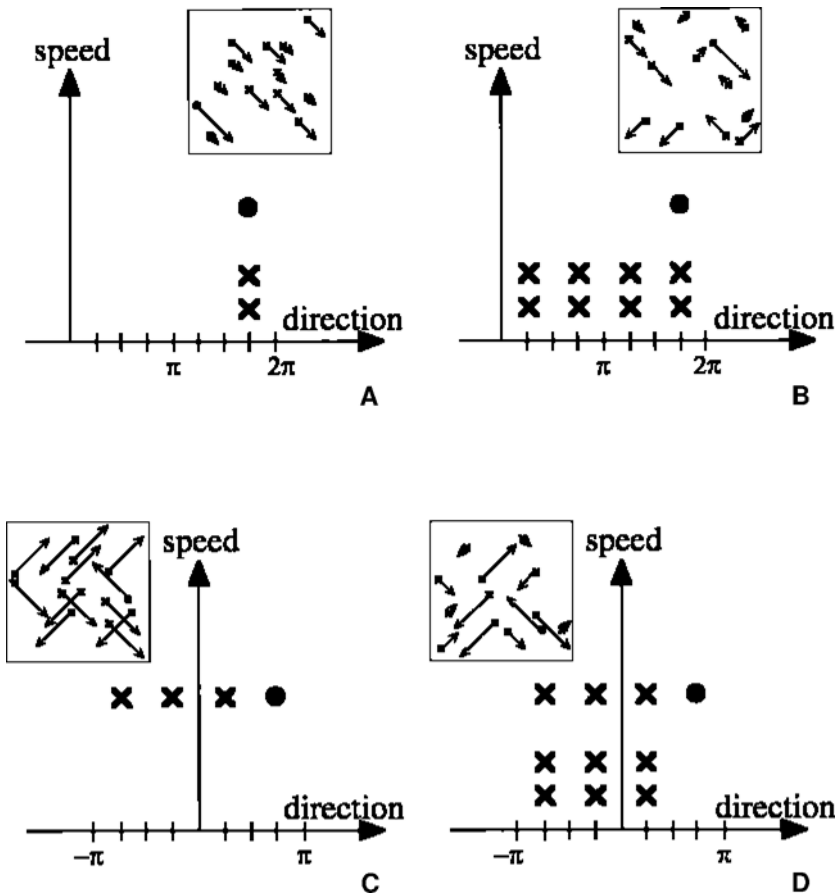


Figure 4. Search for a unique speed (A, B) or direction (C, D). In the homogeneous conditions (A, C), the distractors and the target have the same value along the irrelevant feature dimension (direction or speed, respectively). In the heterogeneous conditions (B, D), the distractors vary along the irrelevant dimension. In speed–direction space, shown here, the design of the two tasks, of search for a unique speed and for a unique direction, appears relatively symmetric. A disk (●) represents the target, and a cross (×) represents each type of distractor.

into question not this conclusion but, rather, the *asymmetry* that direction is encoded independently of speed, but not vice versa.

In both motion search experiments discussed in this section, representation in Cartesian speed–direction space leads to the interpretation of the experimental designs as (fairly) symmetric, thus supporting the conclusion that each set of experiments found a search asymmetry. Representation in 2-D velocity space leads one to interpret the designs as asymmetric and thus to draw no conclusions about asymmetric search mechanisms. Which is the best feature space to use?

In some cases, one has independent knowledge of the best feature space to choose. Such knowledge might be based on which space was a uniform feature space—that is, in which space discriminability contours were circles of uniform size throughout the space. Such independent knowledge is available for, for example, color, but not for motion.

In the absence of independent knowledge of the appropriate feature space, I argue that one should choose 2-D velocity space and the interpretation without search asymmetries. If a simple, symmetric model can explain all three of the classic motion search results, just by representing the motions in 2-D velocity space, that model explains the results in a simpler way than a model that chooses Cartesian speed–direction space and must contain asymmetric search mechanisms.

Case 3: Take into account the background when judging design symmetry. A number of experimental designs initially seem quite symmetric (Nagy & Cone, 1996; Winterbottom & Nagy, 1999). Take, for instance, search for a unique saturation, depicted in Figures 6A and 6B.¹ With only one target type and homogeneous distractors, this design appears completely symmetric—a simple reflection or rotation in feature space transforms one experiment into the other. Nagy and Cone reported that search for a more saturated target is easier than




Figure 5. Search for a unique speed (A, B) or direction (C, D). In the homogeneous conditions (A, C), the distractors and the target have the same value along the irrelevant feature dimension (direction or speed, respectively). In the heterogeneous conditions (B, D), the distractors vary along the irrelevant dimension. In two-dimensional velocity space, the design of the two tasks appears asymmetric. A disk (●) represents the target, and a cross (×) represents each type of distractor.

search for a less saturated target. They suggested that the apparent search asymmetry may be due to signals that are more saturated traveling faster through the visual system than do less saturated ones.

However, consider what the picture looks like once we consider the background color, marked with a **B** in Figure 6. The background was a dark gray for both experiments. Dark gray, being unsaturated, is closer to the less saturated colors used than to the more saturated colors. Thus, when we add the background to our representation of these experiments, the design no longer appears symmetric. No longer can a simple reflection or rotation in feature space transform one experiment into the other.

One cannot ignore the background in search displays. Consider an extreme example in which the target is nearly indistinguishable from the background; this search task is likely to be difficult, regardless of the color of the distractors. On the other hand, search for a target that is distinguishable from the background should be easy when the distractors are barely distinguishable from the background, regardless of the color of the distractors. At the very least, models and experimental designs should con-

sider contrast between the target/distractors and the background, as opposed to absolute feature values.

Figure 7 shows displays of search for a more or less saturated target against a background more saturated than the target or the distractors. This background, although clearly discriminable from both the target and the distractors, is closer in color to the more saturated elements. Against this background, it appears likely that search for a more saturated target would be more difficult than that for a less saturated target.

From Figure 8, we might expect no asymmetry when searching for a target differing from the distractors only in hue, against a dark gray background. In this case, the design of the two experiments is completely symmetric even considering the background, and Nagy and Cone (1996) find no search asymmetry. (Treisman & Gormican, 1988, found an asymmetry in search for a unique hue. However, these asymmetries were small, and D'Zmura, 1991, and Nagy & Cone, 1996, were unable to replicate this result.)

Similarly, Winterbottom and Nagy (1999) have presented the following color search experiment. In their first

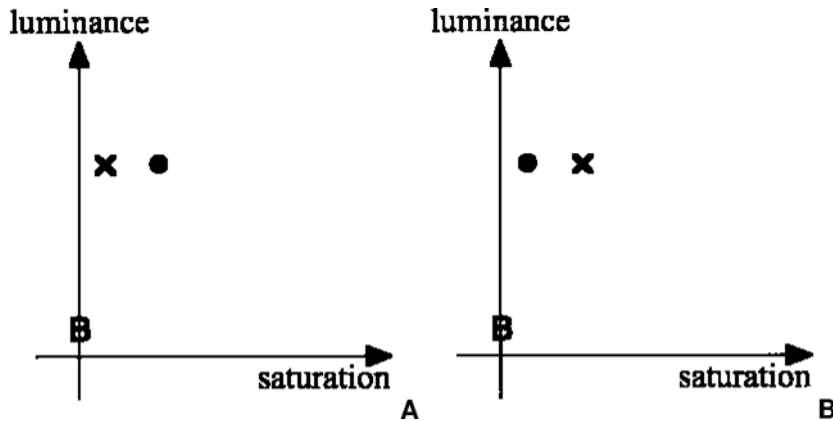


Figure 6. Search for a more saturated (e.g., red) target among less saturated (e.g., pink) distractors (A), and vice versa (B). Once we take into account the background color (B), the experimental design, which appears symmetric when one ignores the background color, appears asymmetric. A disk (●) represents the target, and a cross (×) represents each type of distractor.

experiment, they compared search for a light gray target among homogeneous mid-gray distractors with that among heterogeneous distractors of the same luminance but varying degrees of red or green. Figure 9 depicts this situation. They found no difference between the heterogeneous and the homogeneous cases.

Winterbottom and Nagy (1999) then performed a similar experiment, shifting the target and distractors—but, notably, *not the background*—a bit in the red direction. Thus, observers searched for a pinkish target among distractors that were more red than those in the previous experiment. Figure 10 depicts the situation. In this experiment, Winterbottom and Nagy found that search was significantly more difficult in the heterogeneous case. They concluded that observers can make use of an independent luminance channel to distinguish between target and distractors when the target is achromatic but cannot make use of such an achromatic mechanism when the

target is not achromatic. However, again they based their conclusions on an assumption of design symmetry of the two experiments. When one considers the background, the experimental design no longer appears fully symmetric. One cannot transform the feature-space plots of the first pair of experiments (Figure 9) into those of the second pair of experiments (Figure 10), using only a combination of rotation, reflection, and translation.

In orientation search, the “background” may also account for known search asymmetries. Treisman and Gormican (1988) showed that observers search efficiently for an oblique line among vertical lines but search inefficiently for a vertical line among obliques. They suggested that perhaps observers expect vertical lines more than oblique ones and, thus, oblique lines draw attention because they violate expectations. They demonstrated that the search asymmetry reverses when one presents displays surrounded by an oblique rectangular frame.




Figure 7. Search for a more saturated (red) target among less saturated (pink) distractors (A), and vice versa (B), against a saturated magenta background. It appears easier to find the less saturated target against this background.




Figure 8. Search for a color differing only in hue—for example, search for a yellow target among orange distractors (A), and vice versa (B). In these plots, angle represents hue, distance from the origin represents saturation, and depth into the page indicates luminance (the background lies at a lower luminance than do the target and the distractors). These experiments appear to have a completely symmetric design, and we would expect no asymmetry in search performance, in agreement with experimental results. A disk (●) represents the target, and a cross (×) represents each type of distractor.

Asymmetric expectations are one possible interpretation of these results; another is that the rectangular shape of the display (either the aperture or the border of the computer monitor) may contribute vertical and horizontal orientation estimates, whose existence makes the design asymmetric. Mori (1999) has not been able to replicate this result with the oblique frame, however, for reasons that remain unclear. He did find that the asymmetry reversed when he surrounded each display element by an oblique rectangular frame.

Now, let us revisit the asymmetry of efficient search for a moving object among stationary ones but inefficient search for a stationary object among objects moving in random directions. In the section on Case 1, it was argued

that a design in which search for a moving object was compared with search for a stationary object among distractors moving coherently in one direction would be more symmetric.

Royden et al. (2001) ran this more symmetrically designed experiment. They found efficient search for a stationary target under these conditions. However, search for a moving target was still somewhat more efficient than search for a stationary target. This might suggest a remaining asymmetric mechanism in motion search. However, they lit the room so that observers could see the border of the monitor and a number of other stationary objects. Arguably, the stationary background makes the two search tasks asymmetric, and one would expect search

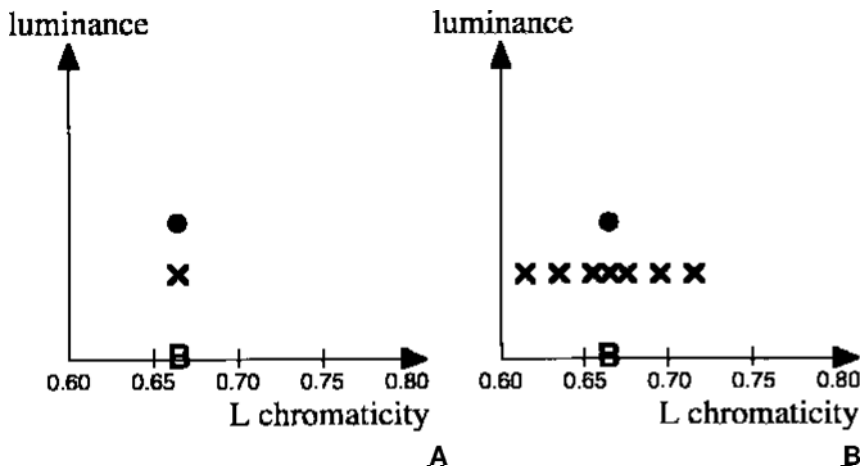


Figure 9. Search for a light gray target among (A) homogeneous mid-gray distractors and (B) distractors with added variability in the red–green direction. Displays used a dark-gray background. A disk (●) represents the target, and a cross (×) represents each type of distractor.

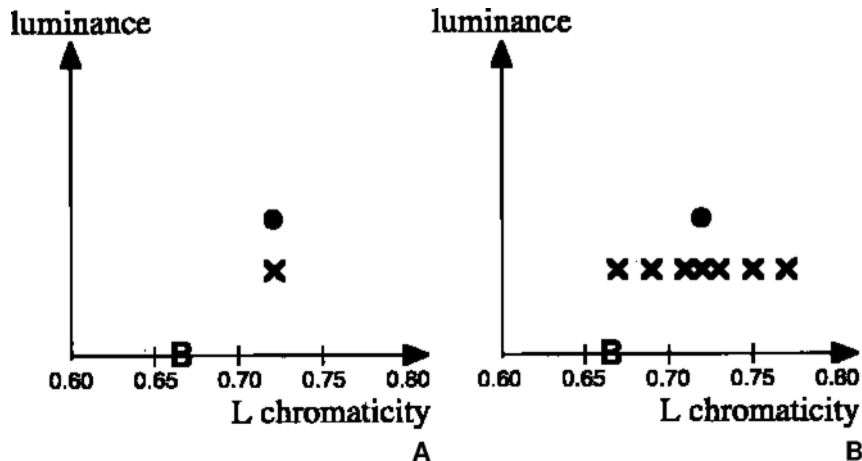


Figure 10. Search for a light pink target among (A) homogeneous darker pink distractors and (B) distractors with added variability in the red–green direction. The target and distractors have been shifted in the red direction, relative to the experiment shown in Figure 8. The background remains a dark gray. A disk (●) represents the target, and a cross (×) represents each type of distractor.

for a stationary target to be less efficient than search for a moving target, owing to similarity of the stationary target to the stationary background.

Explaining Search Results Without Asymmetric Mechanisms

In this paper, it is argued that a number of experiments have built-in design asymmetries and thus may not indicate the existence of asymmetric search mechanisms. However, many models of visual search predict little effect from these design asymmetries. Take, for instance, the asymmetry of Winterbottom and Nagy (1999; see Figures 9 and 10). A number of models of visual search (e.g., Palmer, Ames, & Lindsey, 1983; Treisman & Gelade, 1980) depend predominantly on the difference between the target and the type of distractor most easily confused with the target. Winterbottom and Nagy's target and the distractor type most similar to it lie always at the same distance. Thus, these models predict the same performance in the two sets of experiments, and deviation from equality might indicate an asymmetric search mechanism. Furthermore, the target and the distractors have, in each condition, the same degree of linear separability. Thus, linear separability models of visual search (Bauer, Jolicoeur, & Cowan, 1996; D'Zmura, 1991) also interpret the asymmetric performance as evidence of the existence of asymmetric mechanisms. In a sense, the symmetry of the design depends on the model of visual search.

Therefore, I must not only argue that the search experiments had asymmetric design, but also demonstrate the existence of a model that predicts the experimental results without reference to asymmetric mechanisms. This section offers a simple search model that predicts the results of the experiments discussed above, without refer-

ence to search asymmetries. This model is more parsimonious than previous explanations, since it relates search performance to a simple measure of target–distractor similarity, without asymmetric mechanisms.

The saliency model for visual search. The model that predicts these experimental results without asymmetric mechanisms is as follows (Rosenholtz, 1999). First, represent the features of each display element as a point, \mathbf{p}_i , in an appropriate uniform feature space (e.g., uniform color space). From the distribution of the features present in the display, compute the mean and covariance of the distractor features, μ and Σ , respectively. One can add internal noise at this stage, by adding the noise covariance to the covariance of the distractors. The model then defines target saliency as the Mahalanobis distance, Δ , between the target feature vector, \mathbf{T} , and the mean of the distractor distribution, where

$$\Delta^2 = (\mathbf{T} - \mu)' \Sigma^{-1} (\mathbf{T} - \mu).$$

In this equation, \mathbf{T} and μ are vectors, Σ is a matrix, and the prime indicates a vector transpose.

Essentially, the model uses, as the measure of target saliency, the number of standard deviations between the target feature vector and the mean distractor feature vector. In the simplest version of the model, the more salient the target, the easier the search.

One way to account for the effects of the background on the search task is to count the background as a distractor. What weight the background should have remains an open question. For color search, perhaps one should weight the background according to the proportion of the display area it accounts for—thus, target saliency depends on the average color of the display. Alternatively, the visual system may reduce the weight of the back-

ground, judging it an unlikely candidate for the target. The qualitative predictions in this paper do not require a concrete answer as to how to weight the background.

In many cases, one can immediately tell from the representation of the stimuli in feature space whether the model predicts efficient or inefficient search. The figures that follow represent the mean and covariance of the distractors by the 1σ covariance ellipse (and sometimes the 2σ , 3σ , etc.), centered at the mean distractor feature vector. When the target falls within this ellipse, the saliency model predicts inefficient search. The farther outside these ellipses the target lies, the easier the predicted search. For more subtle cases, calculating the saliency, Δ , clarifies the prediction.

Predicting motion search results. A previous paper (Rosenholtz, 1999) has already shown that this model predicts the results of the three classic motion search asymmetries. In the original moving versus stationary experiments shown in Figure 2, the stationary target has zero saliency, since its velocity lies right at the mean of the distractor velocities. Therefore, the saliency model correctly predicts inefficient search for a stationary target among randomly moving distractors. A moving target among stationary distractors has large saliency (the actual value depends on the target speed relative to the amount of internal noise), and the model correctly predicts efficient search.

Figure 11 depicts the situation for Royden et al.'s (2001) experiment, in which the moving distractors move coherently. Here the stationary background arbitrarily counts as $1/30$ of the distractors. The height of the ellipses indicates the amount of internal noise added. The presence of a stationary background causes the ellipse about the moving distractors to spread toward the origin. Thus, the stationary target lies closer to the 1σ ellipse than does the moving target and has lower saliency. Because of the visibility of

the stationary background, the model predicts somewhat less efficient search for a stationary target. This situation is akin to search for a vertically oriented line among oblique lines, with a visible vertical frame around the display.

Recall that Ivry and Cohen (1992) report efficient search for a quickly oscillating target among slowly oscillating distractors, but inefficient search for the more slowly oscillating target. Figure 12 depicts this situation. The saliency of the slowly oscillating target is less than 1, whereas the saliency of the quickly oscillating target is significantly larger than 1. Thus, the model correctly predicts the experimental results, without reference to asymmetric mechanisms. Consideration of visible stationary objects, such as the border of the monitor, causes the model to predict an even larger difference between search for a quickly oscillating target and that for a slowly oscillating target.

The motion search experiments depicted in Figure 5 involved adding variability in motion direction or speed, with the intent to probe independent coding of speed and direction. Using isotropic, normally distributed internal noise with a standard deviation of 0.2 deg/sec , we find target saliencies of roughly 4.3 and 3.4 for search for a unique speed among homogeneous and heterogeneous distractors, respectively (Figures 5A and 5B). For search for a unique direction among homogeneous and heterogeneous distractors (Figures 5C and 5D), we find average saliencies of 2.5 and 2.3, respectively. One can easily believe, from these numbers, that adding variability in direction significantly affects search for a unique speed but that adding variability in speed does not significantly affect search for a unique direction.

Thus, the saliency model qualitatively predicts the results of all three classic motion search asymmetries, demonstrating that one can model the results without relying on asymmetric search mechanisms.




Figure 11. (A) Search for a stationary target among coherently moving distractors; (B) search for a moving target among stationary distractors. In both cases, observers could see a stationary background (B), for example, the border of the monitor. Counting this background as a distractor causes the distribution of distractors to be pulled toward the stationary target (A). Thus, we correctly predict efficient search for a stationary target, but somewhat less efficient search for a moving target. A disk (●) represents the target, and a cross (×) represents each type of distractor.




Figure 12. Search for a slowly oscillating target among quickly oscillating distractors (A), and vice versa (B). The slowly oscillating target (A) lies within the 1σ covariance ellipse shown and has a saliency of less than one. The quickly oscillating target (B) lies outside the 1σ covariance ellipse and has a saliency of more than one. The saliency model correctly predicts easier search for a quickly oscillating target. A disk (●) represents the target, and a cross (×) represents each type of distractor.

Predicting color search results. Figure 13 depicts the experiments that purported to show asymmetry in search for a unique saturation. The plots show the 1 and 2σ covariance ellipses. Here, the width (minor axis) of the ellipse indicates the internal noise used, and the weight given the background is equal to the total weight of the distractors. The more saturated target lies farther outside the 2σ ellipse than does the less saturated target, leading to the prediction of more efficient search for the more saturated target, in agreement with experimental results.

Figure 14 shows the situation for Winterbottom and Nagy's (1999) experiments, with the same proportional weighting of the background as in the previous example. The internal noise is set so as to have a standard deviation of one half the difference between the most similar colors in the display. From Figures 14A and 14B, one can see that the gray target has the same saliency in the two

conditions, and the model correctly predicts no difference in performance for homogeneous versus heterogeneous distractors. On the other hand, the pink target in Figure 14C has higher saliency than in Figure 14D, and the model correctly predicts easier search in the homogeneous condition.

The saliency model qualitatively predicts these color search "asymmetries," again demonstrating that one can model the results without relying on asymmetric search mechanisms.

Conclusions

The analysis in this paper has demonstrated that a number of asymmetrical results in visual search experiments involving low-level features, such as motion and color, may not indicate the presence of asymmetric mechanisms in the visual system. The experiments used to




Figure 13. Search for a more saturated target among less saturated distractors (A), and vice versa (B). The model correctly predicts easier search for the more saturated target. A disk (●) represents the target, and a cross (×) represents each type of distractor.

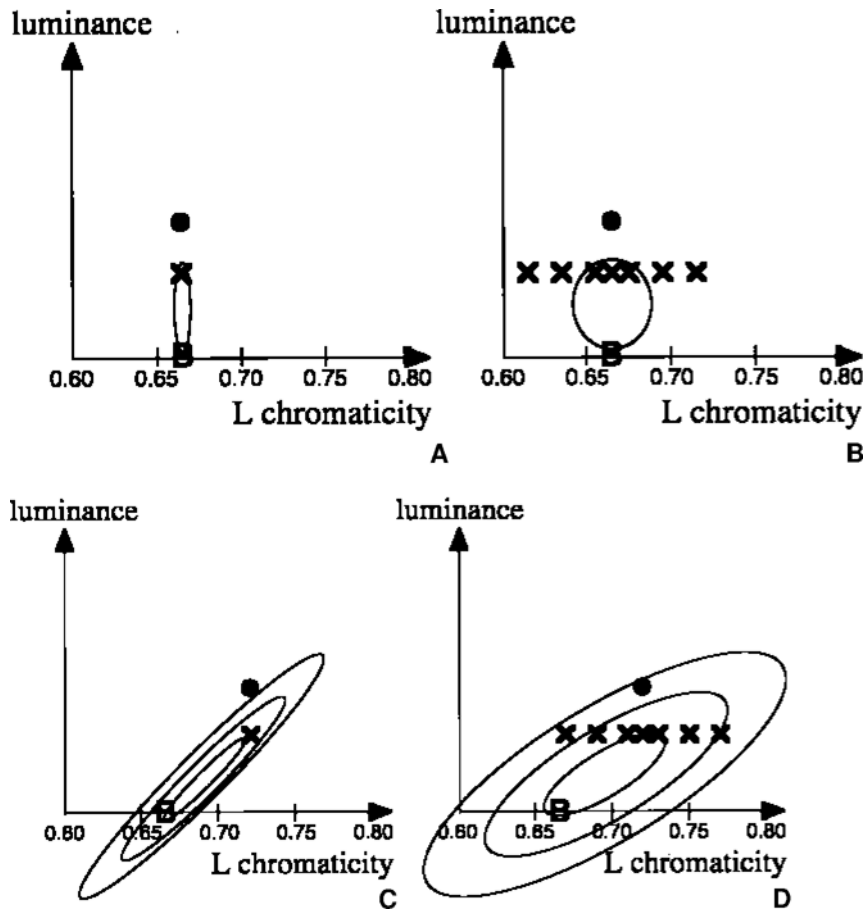


Figure 14. Search for a light gray target (A, B) or a light pink target (C, D) among homogeneous (A, C) or heterogeneous (B, D) distractors. The saliency model correctly predicts easier search for a pink target among homogeneous distractors than among heterogeneous distractors. A disk (●) represents the target, and a cross (×) represents each type of distractor.

demonstrate these asymmetries contain design asymmetries and thus are inconclusive when it comes to proving search asymmetries.

This paper has also shown that one can model the results of these experiments by using a simple saliency model, based on a measure of target–distractor similarity, with no need for asymmetric mechanisms. The fit of the model is somewhat sensitive to the choice of the weighting for the background, as compared with the distractors, as well as to the choice of internal noise. Further work remains to select these parameters. The point here is merely that a simple model exists that can predict the experimental results without the need of asymmetric mechanisms.

These results are good news for modeling of search and attention for application purposes, since they suggest that a simple model may suffice to predict the results of a broad range of search results. The results call into question a number of conclusions about search mechanisms that have been drawn on the basis of these experiments and earlier models of visual search.

The process of analyzing the design asymmetries of a number of motion, color, and orientation search experiments has suggested a number of new experiments, aimed at the question of whether search asymmetries will remain when tested with more symmetric designs. What happens when searching for a stationary target among moving distractors in a motion Ganzfeld? What happens when one changes the background color in color search displays? These suggested experiments bring up the point that the saliency model is fundamentally a model of pop-out—of rapid search for an oddball display element. Changes in the background in some of these suggested experiments may significantly change the *appearance* of the target and distractors (see Figure 7), which could complicate a task of search for a target with a particular appearance.

Although this paper has only discussed asymmetries in search for basic, low-level features, its analysis of the asymmetry of search experiments has implications for examining purported asymmetries in other feature domains. One should ask a number of questions of such ex-

periments. Do the experiments seem to be symmetric? If not, asymmetric results provide no evidence of asymmetric mechanisms. Do the experiments seem symmetric in one feature space but asymmetric in another, and might the latter feature space be the more appropriate one in which to view the experiments? Finally, in many cases, the background or area surrounding the stimulus contains features relevant to the search task. Do the experiments preserve the relationship between the background features and the target and distractors?

Motion, color, and orientation are probably among the best understood feature domains. Researchers have a good idea of the measured features and the appropriate feature space for each of these domains. The analysis here cannot yet explain asymmetries in feature domains for which it is unknown what features the visual system measures and how those features are quantified. For example, researchers have not yet determined the correct feature spaces for curvature, line length, proximity, shape (Treisman & Gormican, 1988), and 3-D appearance (Enns & Rensink, 1991; Sun & Perona, 1996). These representation problems need to be solved before one can determine whether asymmetries in these feature spaces are true asymmetries.

At first glance, it seems more difficult to explain away these asymmetries, and perhaps they are true asymmetries. However, it is possible that an account without asymmetries will appear, once we know the appropriate feature space.

Take the example of search for a curved line among straight lines being more efficient than vice versa (Treisman & Gormican, 1988). For this experiment, observers searched for a straight line among randomly oriented curved lines. In differential geometry, there is the notion of a *curvature vector* (see, e.g., Lipschutz, 1969). At a given point on a planar curve, this vector is perpendicular to the curve and points toward the center of the circle that best fits the curve at that point. The length of this vector is equal to the curvature at that point. Suppose that the curved and straight lines were represented by this curvature vector. For a straight line, the vector would have zero length (zero curvature), whereas for the curved lines, the vector would have nonzero length but random orientation owing to the random orientation of the curved lines. Thus, in this feature space, search for a curved line among straight lines, and vice versa, would look much like search for a moving element among stationary ones, and vice versa, as in Figure 2. In this feature space, the experiment appears to have an asymmetrical design.

As another example, consider the "asymmetry" that search for a C among Os is more efficient than vice versa. The standard account of this asymmetry is that a feature detector responds strongly to the Cs, but not as strongly to the Os, and that it is easier to detect a high response among low responses than vice versa.

Suppose instead that a feature detector responds throughout the image, rather than giving a single response to each letter. Let the average responses to a C and an O

be μ_C and μ_O , respectively. The average response μ_C need not be greater than μ_O . One might easily imagine that the detector response to different parts of a C might be more variable than the response to different parts of an O—an O is, after all, more homogeneous than a C. Let the standard deviations of the responses to a C and an O be σ_C and σ_O , where $\sigma_C > \sigma_O$. Then, the saliency of a target C among Os is $|\mu_C - \mu_O|/\sigma_O$, which is larger than the saliency of a target O among Cs, $|\mu_C - \mu_O|/\sigma_C$, owing to the larger variability in response to a C. Thus, the saliency model would correctly predict easier search for a C among Os than vice versa, without requiring an asymmetric search mechanism in the classic sense. The asymmetry lies not in the search mechanisms, but in the inhomogeneity of a C relative to an O.

The visual system may yet prove to have asymmetric mechanisms when it comes to visual search. However, by calling into question a number of asymmetries in simple features, such as motion, color, and orientation, the analysis here suggests that such asymmetries may be the exception, rather than the rule, as was previously thought.

REFERENCES

- BAUER, B., JOLICEUR, P., & COWAN, W. B. (1996). Visual search for colour targets that are or are not linearly-separable from distractors. *Vision Research*, **36**, 1439-1466.
- DAINTITH, J., & NELSON, R. D. (Eds.) (1989). *The Penguin dictionary of mathematics*. Harmondsworth, U.K.: Penguin Books.
- DICK, M. (1989). *Parallel and serial processes in motion detection*. Unpublished doctoral dissertation, Weizmann Institute, Rehovot, Israel.
- DICK, M., ULLMAN, S., & SAGI, D. (1987). Parallel and serial processes in motion detection. *Science*, **237**, 400-402.
- DRIVER, J., MCLEOD, P., & DIENES, Z. (1992). Are direction and speed coded independently by the visual system? Evidence from visual search. *Spatial Vision*, **6**, 133-147.
- D'ZMURA, M. (1991). Color in visual search. *Vision Research*, **31**, 951-966.
- ENNS, J. T., & RENSINK, R. A. (1991). Preattentive recovery of three-dimensional orientation from line drawings. *Psychological Review*, **98**, 335-351.
- IVRY, R., & COHEN, A. (1992). Asymmetry in visual search for targets defined by differences in movement speed. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 1045-1057.
- LIPSCHUTZ, M. (1969). *Schaum's outline of theory and problems of differential geometry*. New York: McGraw-Hill.
- MORI, S. (1999). Frame of reference for orientation in visual search. *Abstracts of the Psychonomic Society*, **4**, 37.
- NAGY, A., & CONE, S. M. (1996). Asymmetries in simple feature searches for color. *Vision Research*, **36**, 2837-2847.
- PALMER, J., AMES, C. T., & LINDSEY, D. T. (1993). Measuring the effect of attention on simple visual search. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 108-130.
- ROSENHOLTZ, R. (1999). A simple saliency model predicts a number of motion popout phenomena. *Vision Research*, **39**, 3157-3163.
- ROYDEN, C. S., WOLFE, J., & KLEMPEN, N. (2001). Visual search asymmetries in motion and optic flow fields. *Perception & Psychophysics*, **63**, 436-444.
- SUN, J., & PERONA, P. (1996). Early computation of shape and reflectance in the visual system. *Nature*, **379**, 165-168.
- TREISMAN, A. (1988). Features and objects: The 14th Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*, **40A**, 201-237.
- TREISMAN, A., & GELADE, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, **12**, 97-136.
- TREISMAN, A., & GORMICAN, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, **95**, 15-47.

- WINTERBOTTOM, M. D., & NAGY, A. (1999). The achromatic mechanism in visual search [Abstract]. *Investigative Ophthalmology & Visual Science*, **40**, S808.
- WOLFE, J. M. (1996). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13-73). Hove, U.K.: Psychology Press.

opt for a more intuitive space and/or the space used by the original experimenter. For judging the symmetry of these particular experiments, this use of a nonuniform space should have minimal effect, since their asymmetry is due solely to an asymmetric relationship between the background and the display elements.

NOTE

1. Although in principle, we should be representing color experiments in a uniform feature space, such as CIELAB, the figures in this paper

(Manuscript received March 1, 2000;
revision accepted for publication October 30, 2000.)