Psychonomic Bulletin & Review
1999, 6 (2), 183-203

Feature integration that routinely occurs without focal attention

MICHAEL KUBOVY
University of Virginia, Charlottesville, Virginia

DALE J. COHEN
University of North Carolina, Wilmington, North Carolina

and

JEFF HOLLIER
University of Virginia, Charlottesville, Virginia

To analyze visual scenes, the visual system decomposes the visual scene into features that are processed in parallel by separate subsystems. Certain theories (Treisman, Wolfe) propose that these subsystems function independently before focal attention integrates their output. We describe a new paradigm—the gestalt detection task—that directly assesses the degree of preattentive dependence between any two subsystems. We present five experiments that test whether the subsystems that process form and color function independently in processing brief (and, therefore, preattentively processed) stimuli. Our data show that these two subsystems interact during the preattentive processing of feature-dependent information. They are synergistic when the information they receive is consistent; they are antagonistic when the information they receive is inconsistent.

Herbert Simon (1969) noted that many complex systems are nearly decomposable into subsystems. In such systems, interactions within subsystems are orders of magnitude more common than interactions between subsystems. In the same spirit, the architect Christopher Alexander (1964) argued that difficult problems can be solved only if they are decomposed into separable structural components. It is, therefore, not surprising that the visual system—the complex system designed to solve the difficult problem of seeing—is thought to be decomposable into subsystems we call *modules*. The visual system first decomposes images by feature (or dimension); it then uses specialized modules to process those features in parallel; finally, it reintegrates the information produced by each module.

Many instantiations of this view share the *assumption of module independence*. There are two parts to this assumption: (1) The modules function independently, and (2) a qualitatively different process (often termed *focused attention*) reintegrates the outputs of the modules to produce a percept. The simplicity and elegance of this theory have ensured its wide acceptance.

The assumption of module independence has been tempered by increasing evidence—computational, neuroanatomical, and psychophysical—that the interaction among modules is not negligible. In this paper, we first review the evidence in favor of and against modular independence. We then present a new experimental paradigm, the *gestalt detection* paradigm, designed to explore module interaction, and report five experiments that show how the module for grouping by color interacts with the module for grouping by form.

This research was supported by DHHS Grant R01 MH47317. Parts of this work were previously presented by Kubovy and Cohen (1991, 1992). We thank Helen Kadlec for her superb review of this paper. Correspondence concerning this article should be addressed to M. Kubovy, Department of Psychology, Gilmer Hall, The University of Virginia, Charlottesville, VA 22903-2477 (e-mail: kubovy@ virginia.edu).

ARE MODULES INDEPENDENT?

Normative Considerations

Marr's (1976, 1982) principle of modularity states that "any large computation should be split up into a collection of small, nearly independent, specialized subprocesses" (Marr, 1982, p. 325)—that is, modules. A module is a processing unit that interprets a specific data type in the context of specific constraints (Jepson & Richards, 1992). A modular system has the advantage of reducing the computational capacity needed to solve a complex problem. Take, for example, a problem that faces neural networks that use back-propagation: They may take tens or hundreds of thousands of learning trials to learn to solve a problem. According to Rumelhart (1989) a divide-and-conquer strategy (i.e., modularity) can reduce the number of required trials by orders of magnitude.

A system that implements module independence is said to have a *blackboard architecture* (Erman, Hayes-Roth, Lesser, & Reddy, 1980). This architecture's name comes from the following image: Each module processes information to the best of its ability and writes its output on a reserved part of a blackboard. A so-called agent must

(1) resolve inconsistencies that may occur between the output of several modules and (2) combine their output, without having access to the data or to the nature of prior processing and without being able to influence the modules.

In contrast, a system that violates the hypothesis of module independence (Ullman, 1991) is said to have a *network architecture* (M. Green, 1991). For instance, Poggio, Gamble, and Little (1988) describe an algorithm for the parallel integration of feature dimensions. They assume each feature module to be retinotopic. Units in one module communicate with units in the other modules; the closer they are in retinotopic space, the stronger their interaction. The algorithm is based on the premise that an object's edge produces discontinuities in corresponding spatial locations in the different feature modules. Thus, consistent boundaries in the different feature modules are reinforced, whereas inconsistent boundaries are inhibited. Phillips, Kay, and Smyth (1995a, 1995b) have proposed a similar architecture in which local processors provide each other with contextual guidance.

Neuroanatomical Evidence

The neuroscience literature contains abundant evidence of segregated processing streams analogous to the blackboard architecture we have just discussed. Neuroanatomical and neurophysiological data suggest that three pathways serve different functions (Livingstone & Hubel, 1988). The pathways are summarized in Figure 1 (Livingstone & Hubel, 1987; Merigan & Maunsell, 1993). The first is the *magnocellular → thick stripe → MT* pathway. It is characterized by large, color-blind, fast cells that are most sensitive to motion and depth. The next is the *parvocellular → interblob →* interstripe (*pale stripe*) *→ V4* pathway. This pathway is characterized by small, color-sensitive, slower cells that are most sensitive to shape differences in a static field. It uses both luminance and color contrast to define shapes. Finally, the *parvocellular → blob → thin stripe → V4* pathway is also characterized by small, color-sensitive, slower cells, but these cells are most sensitive to color. These cells also have a lower acuity than the interblob pathway.

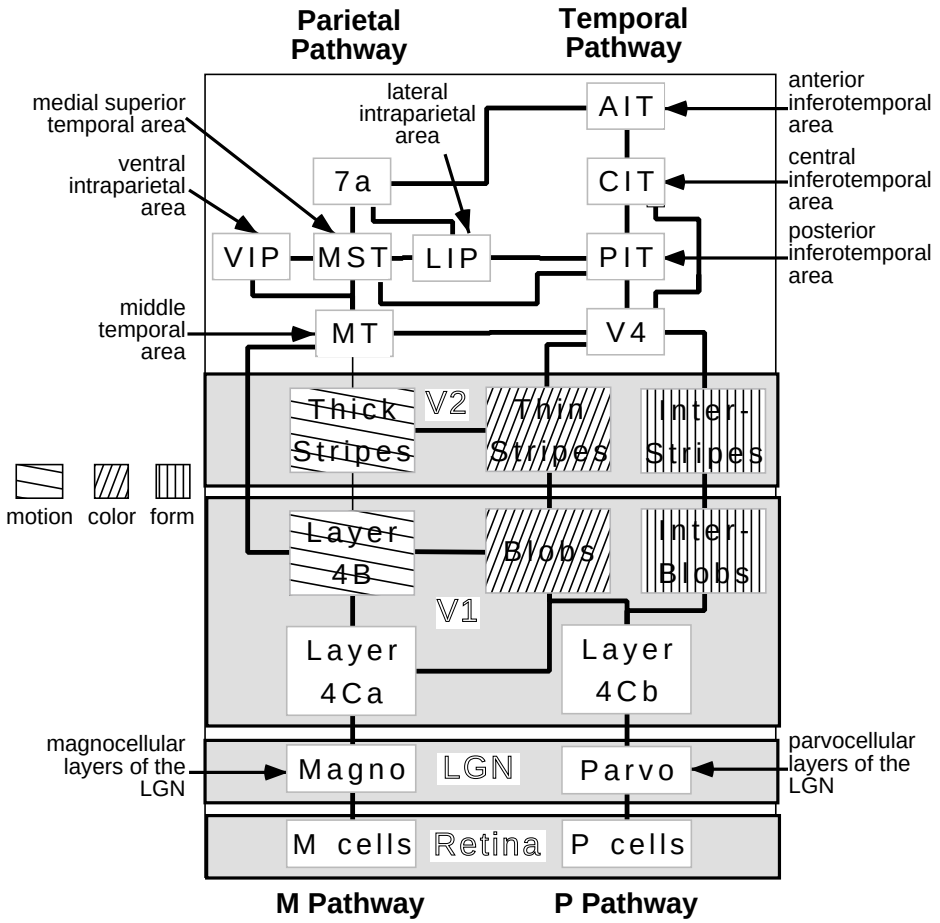

Figure 1. Pathways in the primate visual system. The horizontal connections represent possible pathways for module interaction. The legend on the left represents the major processing function of areas in *V*1 and *V*2.

Notwithstanding this evidence, it has become clear that this framework omits many pathways that suggest numerous forms of interstream communication (Felleman & Van Essen, 1991). These connections provide a neurological foundation for modular interaction and hint at a network architecture. In fact, several voices in the neurological literature have questioned the assumption of module independence (e.g., Merigan & Maunsell, 1993).

We have two reasons to believe that the processing of color is not independent of the processing of form. (1) The interstripes most sensitive to shape differences and the thin stripes most sensitive to color are interleaved. This anatomical arrangement suggests interaction; the appropriate connections have indeed been observed anatomically (Levitt, Yoshioka, & Lund, 1994). (2) Even the interstripes, which are most sensitive to form, contain cells that are sensitive to color.

Psychophysical Evidence

The psychophysical evidence most commonly cited in favor of module independence comes from three sources: studies of visual search, studies of object classification, and studies of texture segregation.

Visual search. In the visual search task, researchers ask observers to find a *target*, say a red triangle, among *distractors*, say red circles and green triangles. They present observers with distractors randomly scattered in the plane among which we may have embedded a target. We call a trial on which we presented a target *positive*, and one on which we did not present a target *negative*. We vary the number of distractors (N) in each display, ask the observers to determine as quickly as they could whether the target was present, and record their reaction times (RTs). The traditional interpretation of these data ran as follows. If the growth of RT as a function of N—RT = $f(N)$—was linear, with a slope of a, investigators inferred that the observers scanned the shapes one by one (i.e., they used a serial search strategy). If, in addition, for the negative trials, the slope of RT = $f(N)$ was linear with a slope of $2a$, investigators inferred that the observers stopped scanning as soon as they found the target (i.e., their serial search strategy was self-terminating; Sternberg, 1970). If a approached 0, they claimed either that the search was parallel (Treisman & Gelade, 1980) or, more conservatively, that the search was efficient (Wolfe, 1994).

In general, observers perform an efficient search when they can distinguish a target from a distractor by one feature alone—*single-feature* displays (e.g., if they search for a red triangle among green triangles or among green circles). Under these conditions, observers report that the target *pops out* and that the task is experienced as one of detection rather than of search. Conversely, observers perform an inefficient search that has features of a serial self-terminating search when they cannot distinguish a target from a distractor by one feature alone—*conjunctive* displays (e.g., if they search for a red triangle among red circles and green triangles). (See Beck & Ambler, 1973; Bergen & Julesz, 1983a, 1983b; Cavanagh, Arguin, & Treisman, 1990; Chmiel, 1989; Kleiss & Lane, 1986; Treisman, 1977, 1988, 1990a, 1990b; Treisman & Gelade, 1980; Treisman & Gormican, 1988; Treisman & Schmidt, 1982; Treisman & Souther, 1985; Wolfe & Cave, 1990; Wolfe, Cave, & Franzel, 1989.)

The most influential theories developed to account for these data, Treisman's feature integration theory (FIT; see references in the preceding paragraph) and Wolfe's (1994) Guided Search 2.0 (GS2) assume module independence. According to both FIT and GS2, early processing of displays is performed by specialized modules designed to perform grouping and segregation. Each module performs this task with respect to just one stimulus dimension—for example, color or form. If the input varies little with respect to the feature it processes, the corresponding module produces little or no output. In contrast, if the stimulus can be partitioned into uniform regions that differ in color or form, the appropriate module produces a tentative partition of the stimulus. Although, beyond this point, FIT and GS2 disagree on some matters, both propose that the processing ends with a serial attentional process that conjoins the information produced by the independent modules.

If the observer is asked to search for a single target that differs from the background with respect to a single dimension, the task is called a *single-feature search* (e.g., a black circle among white circles; Figure 2, top). According to both theories, the partition proposed by the color module isolates the target (the black circle), which

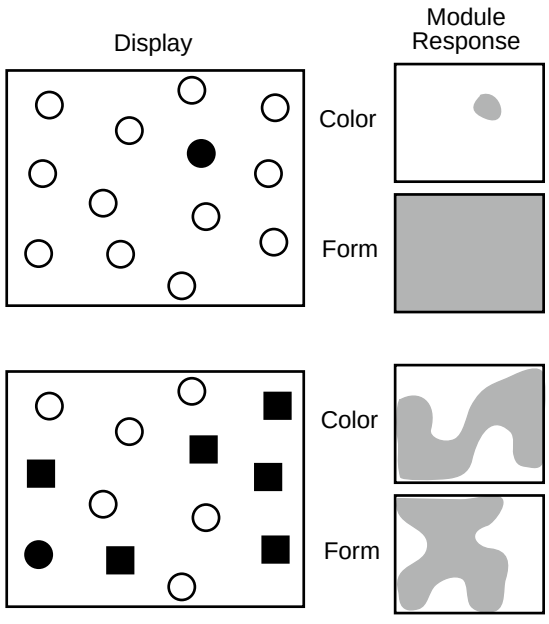

Figure 2. Top: Hypothetical responses of the color and form modules to a target that differs from distractors with respect to a single dimension. Bottom: Hypothetical responses of the modules to a stimulus with two kinds of distractors that differ from each other in two ways and a target that differs from each distractor in only one way.

pops out. The form module is silent, because the stimulus does not vary with respect to form. At this point, focal attention may integrate the target's two features—its color and its form. Without this integration, the observer would not know that the target is a black circle. If the observer is asked to search for a single target that differs from the background with respect to two dimensions, the task is called a *conjunctive-feature search* (e.g., a black circle among white circles and black squares). Here, too, the two theories agree that the modules produce inconsistent partitions of the stimulus (Figure 2, bottom). Hence, the output of the modules cannot draw the observer's attention directly to the odd man out. Therefore, an observer can only find the target by inefficiently scanning the stimulus element by element and integrating the features of each, using *focal attention*. The process called focal attention corresponds to the so-called agent in the blackboard architecture.

Sorting, separability, and module independence. Whenever FIT, GS2, or related theories are presented, two assumptions appear prominently: (1) module independence and (2) the requirement of focal attention for feature integration. Both assumptions originate in Garner's theory of *stimulus separability* and the idea that separability is equivalent to the operation of independent modules (Garner, 1974; Garner & Morton, 1969; Treisman, 1986). To illustrate this idea, consider a fragment of an experiment by Gottwald & Garner (1972). Using cards on which they had drawn one of three shapes (a circle, a square, or an equilateral triangle) in one of three colors (red, blue, or yellow), they asked the observers to sort 30- or 32-card decks into two classes. As Figure 3 (Gottwald & Garner, 1972) shows, when the task required the sorting of colored shapes by color (left column), the observers took about 383 msec per card (157 cards/min), regardless of the number of forms (two, four, or six). The same pattern held when the roles of form and color were interchanged. However, when the task required the sorting of six colored shapes into classes defined by a conjunction of color and form (right column), the observers took 134 msec longer per card (507 msec/card; 118 cards/min). Here is how Garner (1974) explains this type of result: "With [separable dimensions,] dimensional structure is primary, and cannot be ignored. Thus the highly cognitive processing [i.e., focal attention] must be carried out, a processing in which the observer actually deals with each dimension separately [i.e., the output of the independent modules], making a decision about each and then about the two in conjunction" (p. 148).

Since Garner's seminal work, Ashby and Townsend (1986) have provided a theoretical model that clarifies the vague and overlapping notions of independence previously discussed in the literature. Their general recognition theory (GRT) separates the recognition process into two perceptual stages and one cognitive stage. In the first perceptual stage, stimulus information is encoded. In the second, information is processed. The cognitive stage consists of the decision process required to make a response.

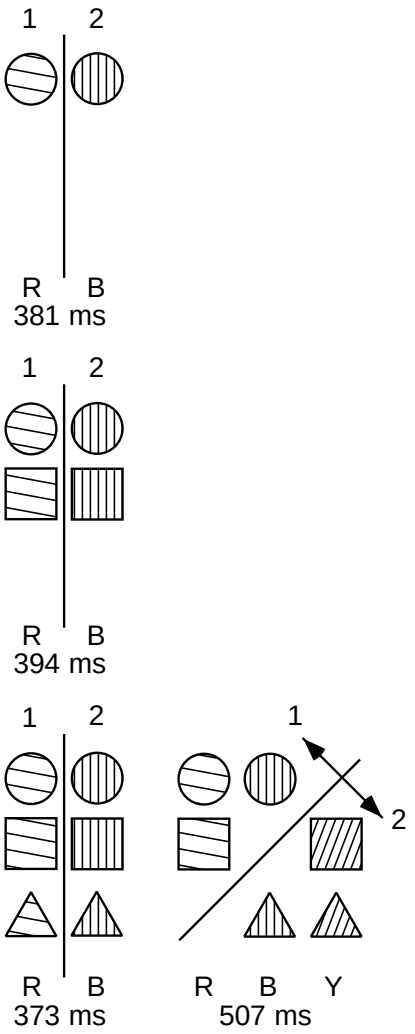

Figure 3. A summary of the Gottwald and Garner (1972) experiment showing the cost of sorting colored shapes by a conjunction of color and form.

Figure 4 summarizes the model. Suppose that the stimuli in a set vary with respect to two dimensions, X and Y; each of the stimuli is characterized by its level on these two dimensions. Let us assume that the modules that encode the stimuli are differentially sensitive to each dimension. If each module is sensitive only to one of these dimensions, according to Ashby and Townsend, such a system conforms to Garner's (1974) notion of perceptual separability. The degree to which the two modules' sensitivities overlap corresponds to the degree of perceptual integrality between the dimensions. In the second stage of the model, stimulus information is processed. The two modules may process their information independently, or they may exchange information during processing. According to Ashby and Townsend, modules that do not share information are perceptually independent, and those that do share information are perceptually depen-

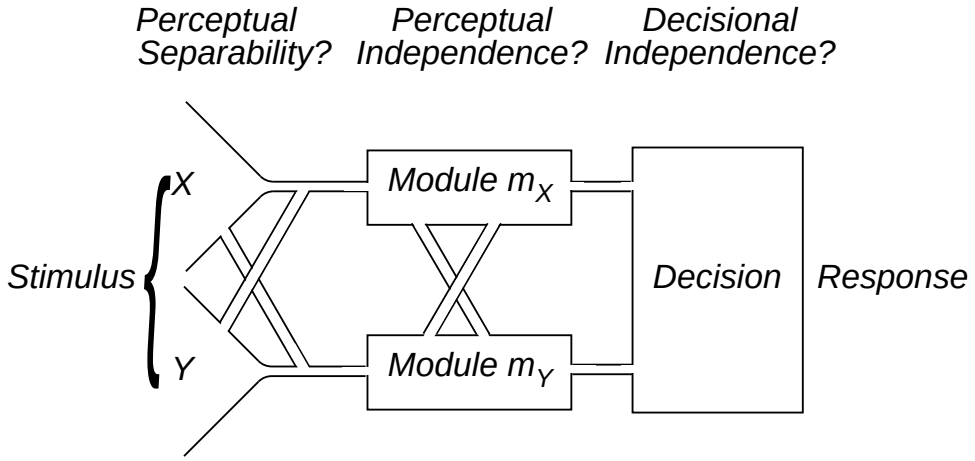

Figure 4. Process independence and three kinds of process correlation for two stimulus dimensions and their relevant response systems.

dent. Finally, during the decision stage, the outputs of each module are combined. If the subjects' decisions about the output of one module are not influenced by their decisions about the output of the other module, we will say that decisional independence holds. If however, the subjects' decisions about the output of one module influence their decisions about the output of the other module, we will say that decisional dependence holds. Thus, separability is equivalent to the principle of modularity, not to module independence.

Texture segregation and popout. Texture segregation experiments test module independence by assessing interference and redundancy. In such experiments, one shows observers arrays of elements that vary with respect to two or more dimensions. Consider, for example, Callaghan's (1989) experiments on texture segregation, which followed upon her earlier experiments in a similar vein (Callaghan, 1984; Callaghan, Lasaga, & Garner, 1986). A typical set of her stimuli is shown in Figure 5. She asked observers to identify the orientation of a boundary that bisected each stimulus. She found that observers located boundaries between homogenous regions (Figures 5a and 5b) more rapidly than they located boundaries between nonhomogenous regions (Figures 5d and 5e). This demonstration of interference suggests that when the color and form modules receive inconsistent information, they inhibit each other. She also found that observers located boundaries faster when they were redundant (Figure 5c) than when they were not (Figures 5a and 5b), but only when hue differences were small. This result is not conclusive, because her analysis did not allow her to infer how much gain would be expected merely from the activation of two concurrent processes. Using more refined analyses of RT distributions, Mordkoff and Yantis (1993) were able to do just that. They showed that RTs obtained with redundant targets were shorter than would be expected on the basis of the independent processing of single features.

One of the present authors (Cohen, 1997) reported experiments (conducted after the present studies were completed) that used the five types of stimuli shown in Figure 6 to study the rapid segregation (i.e., *popout*) of single elements (or pairs of elements) from a lattice of uniform distractors. Cohen applied signal detection, as we did in Experiments 1 and 2 of the present paper, and showed that (1) coincident targets are detected better than disparate targets (true at both 33- and 83-msec display time), (2) coincident targets are detected better than would be predicted assuming module independence (true for 33 and 83 msec), and (3) disparate targets are detected worse than would be predicted assuming module independence (true for 83 msec). From this result, he concluded that color and form modules act in synergy to detect targets that share a location in space and inhibit each other when the information each receives is not supported by information received by the other.[1]

GESTALT DETECTION PARADIGM

The gestalt detection paradigm uses texture segregation to assess interference and redundancy effects in early vision. Four desirable features distinguish this paradigm from earlier work. First, it allows us to manipulate the consistency of the information presented to the modules. We will observe different types of interaction between the modules when the information they process varies in consistency. Thus, we will be in a position to demonstrate interference effects such as those Callaghan (1989) obtained and coactivation effects such as those shown by Mordkoff and Yantis (1993). Second, it goes beyond the work of Cohen (1997), who demonstrated synergy and inhibition in popout, but not in texture segregation. Third, it allows us to manipulate the amount of time in which the modules can process the information. We will observe temporal variations in the interaction between

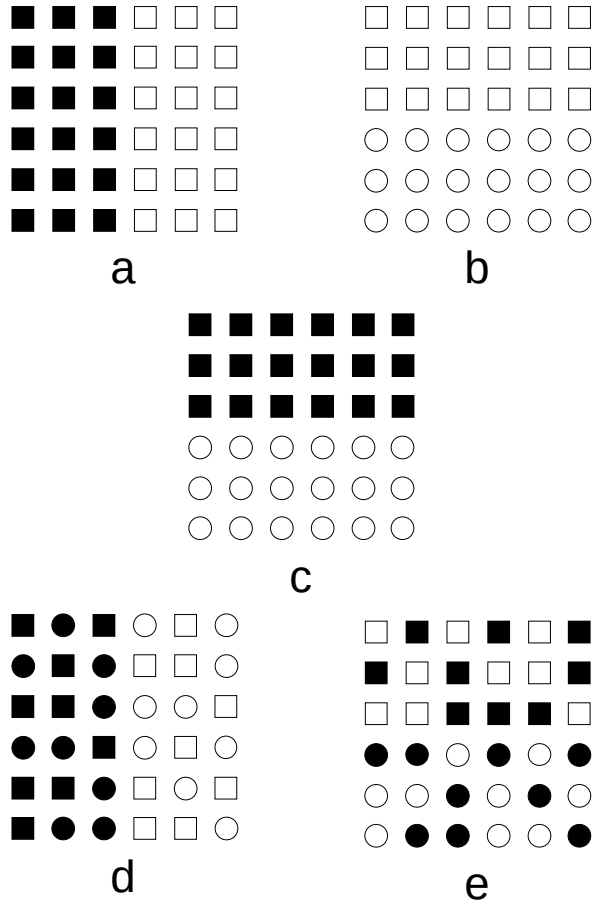

Figure 5. Typical lattices used by Callaghan (1989). Filled and unfilled forms represent different hues. Panels a and d have color boundaries. Panels b and e have form boundaries. Panel c has a form-and-color boundary. In Panels a–c, the boundaries partition the lattice into homogenous regions. In Panels d and e, the boundary partitions the lattice into non-homogenous regions.

modules. Fourth, it allows us to assess independence, using signal detection methodology.

Each of our stimuli consists of a sequence of frames,[2] one of which may be a target; the others are distractors. Thus the gestalt detection paradigm is a combination of rapid serial visual presentation (RSVP; Just, Carpenter, & Woolley, 1982; Mitchel, 1979; Potter, 1986; Young, 1984) and texture segregation (Bergen & Julesz, 1983a, 1983b; Enns, 1986; Gurnsey & Browse, 1987; Julesz, 1981, 1984; Julesz & Bergen, 1983).

Each frame contains a 25-element square lattice of colored forms. The elements can take on two levels of color, denoted $c_k(k = 1, 2)$, and two levels of form, $f_l(l = 1, 2)$. We say that a row is *color homogenous* if all the elements it contains are of uniform color. Parallel definitions can be given of color-homogenous column, form-homogenous row, and form-homogenous column. A row (or a column) can concurrently be color homogenous and form homogenous.

A *color-homogenous* region consists of two or three contiguous color-homogenous rows or columns. A *form-homogenous* region consists of two or three contiguous form-homogenous rows or columns. A frame is a *color target* if its rows (or columns) can be partitioned into exactly two (nonempty) color-homogenous regions; it is a *form target* if its rows (or columns) can be partitioned into exactly two (nonempty) form-homogenous regions. It is a *color-and-form target* if it is both a color and a form target.

Between the two color-homogenous regions of a color target, there is a *color boundary*; between the two form-homogenous regions of a form target, there is a *form boundary*. Between rows (or columns) in homogenous regions, there are *null boundaries*. As Figure 7 shows, in a 5×5 square array, two horizontal boundaries, $h_m(m = 1, 2)$, and two vertical boundaries, $v_n(n = 1, 2)$, are possible.

If a frame has one or two boundaries, it is a target; otherwise, it is a *distractor*. If it has two boundaries, one of them must be a color boundary, and the other must be a form boundary. The structure of a frame can be summarized by a pair $\langle b_C, b_F \rangle$, where $b_C, b_F \in \{h_1, h_2, v_1, v_2, \emptyset\}$, and Ø means the null boundary (Figure 8). If the frame has two boundaries, they are either *coincident* (i.e., $b_C = b_F$) or distinct (i.e., $b_C \neq b_F$). If the boundaries are distinct, they are either *parallel* or *orthogonal.* Table 1 gives the notation for different boundary types.

When given a demanding task, observers may consistently attend to particular locations. In a procedure reminiscent of the Averbach and Coriell (1961) experiments, Monheit and Johnston (1994) briefly showed their observers six colored letters placed at the vertices of a regular hexagon (whose sides subtended 0.86º) centered on a fixation cross. They replaced the six letters with masks

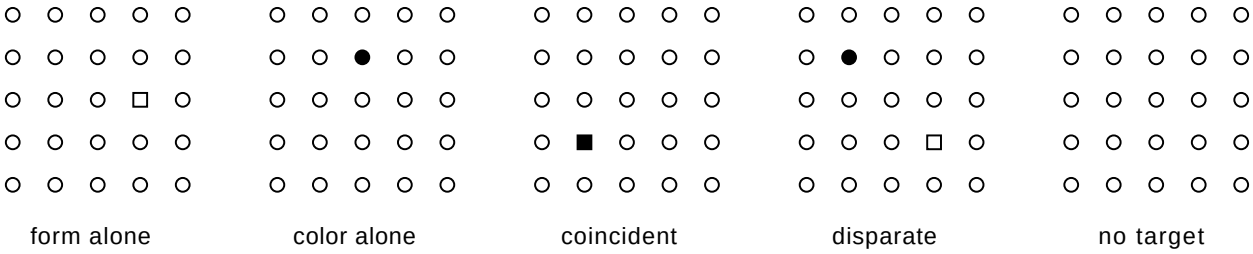

Figure 6. The five display types used in Cohen's (1997) visual detection tasks. Unfilled circles are distractors. A display can contain two targets (the *disparate* display), one target (the *form-alone*, the *color-alone*, or the *coincident* display), or none (the *no target* display). A target can differ from the distractor by one feature (form or color, as in the *form-alone*, the *color-alone*, and the *disparate* displays) or by two features (as in the *coincident* display).

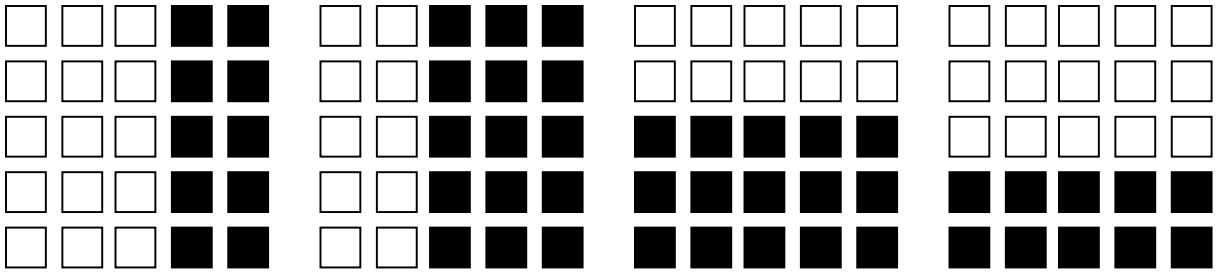

Figure 7. The four locations of boundaries, illustrated with a color boundary. From left to right, these arrays are $\langle v_2, \emptyset \rangle$, $\langle v_1, \emptyset \rangle$, $\langle h_1, \emptyset \rangle$, $\langle h_2, \emptyset \rangle$.

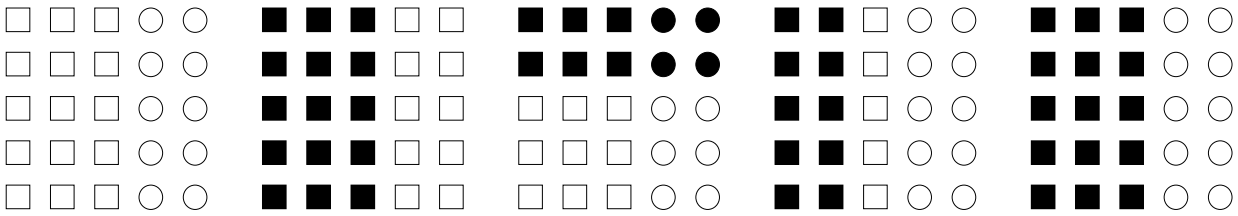

Figure 8. From left to right, examples of $\mathcal{F} = \langle \emptyset, v_2 \rangle$, $C = \langle v_2, \emptyset \rangle$, $\perp = \langle h_1, v_2 \rangle$, $\| = \langle v_1, v_2 \rangle$ and $\sim = \langle v_2, v_2 \rangle$ boundaries (see Table 1).

and the fixation cross with an arrow pointing at one of the vertices. They instructed the observer to report the color and identity of the corresponding letter. They found that the accuracy with which observers reported color was correlated across trials with the accuracy with which they reported form. Furthermore, the observers' performance (average percentage of correct color and form reports) was highest for the letter on the right (about 70% correct) and lowest for the letters on the bottom (about 36% correct). These data suggest that the observers' focused their attention on a region less than 0.8º in diameter.

The use of RSVP in the gestalt detection task requires observers to perform a segregation in time (i.e., to segregate the target from the distractors). They could not, however, perform this segregation if they had not detected grouping in space (i.e., by detecting or localizing a boundary). In other words, we are looking for popout in space and time. To force observers to see global boundaries rather than local differences, we made the distractors (1) similar to targets and (2) uncorrelated with the target within each trial. To decorrelate distractors from the target, we used the following procedure. For each of the 96 targets and for each boundary in a target, we randomly sampled three elements in the large homogenous region and two elements in the small homogenous region (Figure 9). We then interchanged the features of those elements, producing one distractor (out of the 6×10^{12} distractors possible). We repeated this process 90 times for each target, producing a pool of 8,640 distractors. To generate a trial, we sampled (with replacement) the requisite number of distractors, without regard for the targets from which they were derived. Changing only five elements of a target to obtain a distractor makes pseudoboundaries likely: If observers focused on the eight elements on a gray background in Figure 10, they would mistakenly believe that a target had been presented.

Table 1
Symbols for Boundary Types, Actual and Predicted

Symbol	Definition	Boundary Type
0	$\langle \emptyset, \emptyset \rangle$	none
C	$\langle b_C, \emptyset \rangle$	single, color
$\mathcal{F}$	$\langle \emptyset, b_F \rangle$	single, form
1	$\in \{C, \mathcal{F}\}$	single
2	$\in \{\sim, \|, \perp\}$	dual
$\sim$	$\langle b_C, b_F \rangle$ such that $b_C = b_F$	coincident
$\|$	$\langle v_C, v_F \rangle$ or $\langle h_C, h_F \rangle$ such that $b_C \neq b_F$	parallel
$\perp$	$\langle v_C, h_F \rangle$ or $\langle h_C, v_F \rangle$	orthogonal
$\otimes$	independent modules	dual
$\equiv$	identical modules	dual

We used two variants of the gestalt detection paradigm: *boundary detection* and *boundary localization*. In the boundary detection task, the observers were asked to detect the *presence* of a boundary in one of the frames. In the boundary localization task, a target was present on every trial, and the observers were asked for the *location* of *one* boundary, even if two were present.

EXPERIMENTS 1 AND 2
Boundary Detection

Experiments 1 and 2 used the boundary detection task. Models of independence allow us to predict the detectability of a two-boundary target from the detectability of a single-color boundary, d'_C, and the detectability of a single-form boundary, d'_F. We will denote the detectability of a two-boundary frame, predicted on the assumption of independent modules, $d'_\otimes$. For a pair of feature modules to be independent, two criteria may be formu-

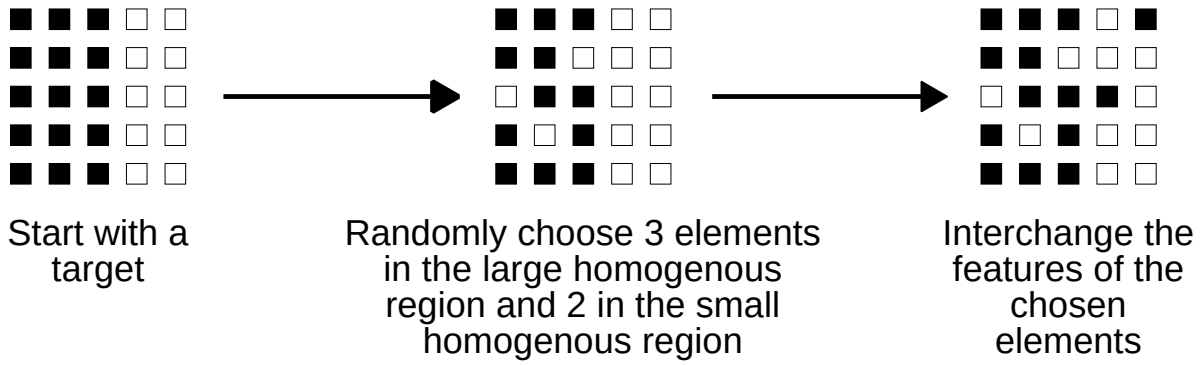

Figure 9. How to build a distractor.

lated: (1) strong independence ($d'_{\otimes} = d'_2$ for all dual-boundary conditions) and (2) weak independence ($d'_{\sim} = d'_{\parallel} = d'_{\perp}$).

Violations of these criteria of independence imply that the modules interact. However, the persuasiveness of such an inference depends on the interpretability of the pattern of violation. We will say that a network is synergistic if a dual-boundary stimulus is detected better than would be predicted by independence ($d'_2 > d'_{\otimes}$). The network is antagonistic if a dual-boundary stimulus is detected worse than would be predicted by independence ($d'_2 > d'_{\otimes}$).

We used a two-dimensional signal detection task (D. M. Green & Swets, 1966; Kadlec & Townsend, 1992; Macmillan & Creelman, 1991; Tanner, 1956). When all the frames on a trial are distractors of type **0** (see Table 1), some of them could be mistaken for a target. Our null hypothesis is dictated by the notion of module independence that we presented at the beginning of this article. Therefore, we assume that, on trials for which all the frames are **0** distractors, the evidence in favor of a color boundary is independent of the evidence in favor of a form boundary. Hence, it is natural to represent these noise trials by a bivariate normal distribution in which the correlation between the two variables is 0 (Figure 11).

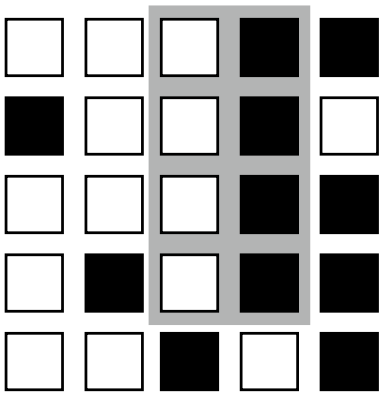

Figure 10. A distractor with a pseudoboundary.

On trials in which a $\mathcal{C}$ or an $\mathcal{F}$ target is presented, a signal of intensity, $d'_{\mathcal{C}}$ or $d'_{\mathcal{F}}$, respectively, is added. Therefore they must be represented by bivariate normal distributions with means $\overline{\mathcal{C}}$ and $\overline{\mathcal{F}}$. We place the origin of the decision space at the mean of the **0** distribution, $\overline{\mathbf{0}}$. The angle between $\overline{\mathbf{0}}\overline{\mathcal{C}}$ and $\overline{\mathbf{0}}\overline{\mathcal{F}}$, θ, expresses the correlation between the output of the two modules. The detectability of a **2** boundary is the distance between the means $\overline{\mathbf{2}}$ and $\overline{\mathbf{0}}$—that is,

$$d'_{\mathbf{2}} = \sqrt{d'^2_{\mathcal{C}} + d'^2_{\mathcal{F}} + 2d'_{\mathcal{C}}d'_{\mathcal{F}}\cos\theta}. \qquad (1)$$

The more synergistic interaction there is between the two modules, the smaller the θ. Since $d'_{\mathcal{C}}$ and $d'_{\mathcal{F}}$ do not change, the detectability of a **2** frame decreases as θ increases (because $\cos\theta$ decreases). To describe the relationship between the modules, we will need two new symbols (Table 1), which we will define in a moment: $\otimes$, which stands for independent modules, and $\equiv$, which stands for identical modules. As θ grows, the relationship between the modules goes through four conditions: (1) *identity*, $\theta = 0$; the modules are indistinguishable ($d'_{\mathbf{2}} = d'_{\equiv} = d'_{\mathcal{C}} + d'_{\mathcal{F}}$); (2) *synergy*, $0 < \theta < \pi/2$; the modules are cooperative ($d'_{\otimes} < d'_{\mathbf{2}} < d'_{\equiv}$); (3) *independence*, $\theta = \pi/2$; the modules are independent ($d'_{\mathbf{2}} = d'_{\otimes} = \sqrt{d'^2_{\mathcal{C}} + d'^2_{\mathcal{F}}}$); and (4) *antagonism*, $\theta > \pi/2$; the modules inhibit each other ($d'_{\mathbf{2}} < d'_{\otimes}$). A plausible lower bound on antagonism is given by $\min(d'_{\mathbf{2}}) = \max(d'_{\mathcal{C}}, d'_{\mathcal{F}})$, because any more antagonism would destroy information.

If we observe independence, we will have supported the current view, represented by FIT and GS2. If, on the other hand, we observe nonindependence, the refutation of the current view is not immediate. Our model of independence may be violated because: (1) the two modules have overlapping sensitivities (i.e., a failure of perceptual separability), or (2) the two modules interact during processing (i.e., a failure of perceptual independence; Figure 4). The key to resolving this indeterminacy lies in the pattern of nonindependence.

Perceptual separability fails (case 1) when a module that is primarily activated by one dimension is also activated by the other dimension (e.g., the form module is

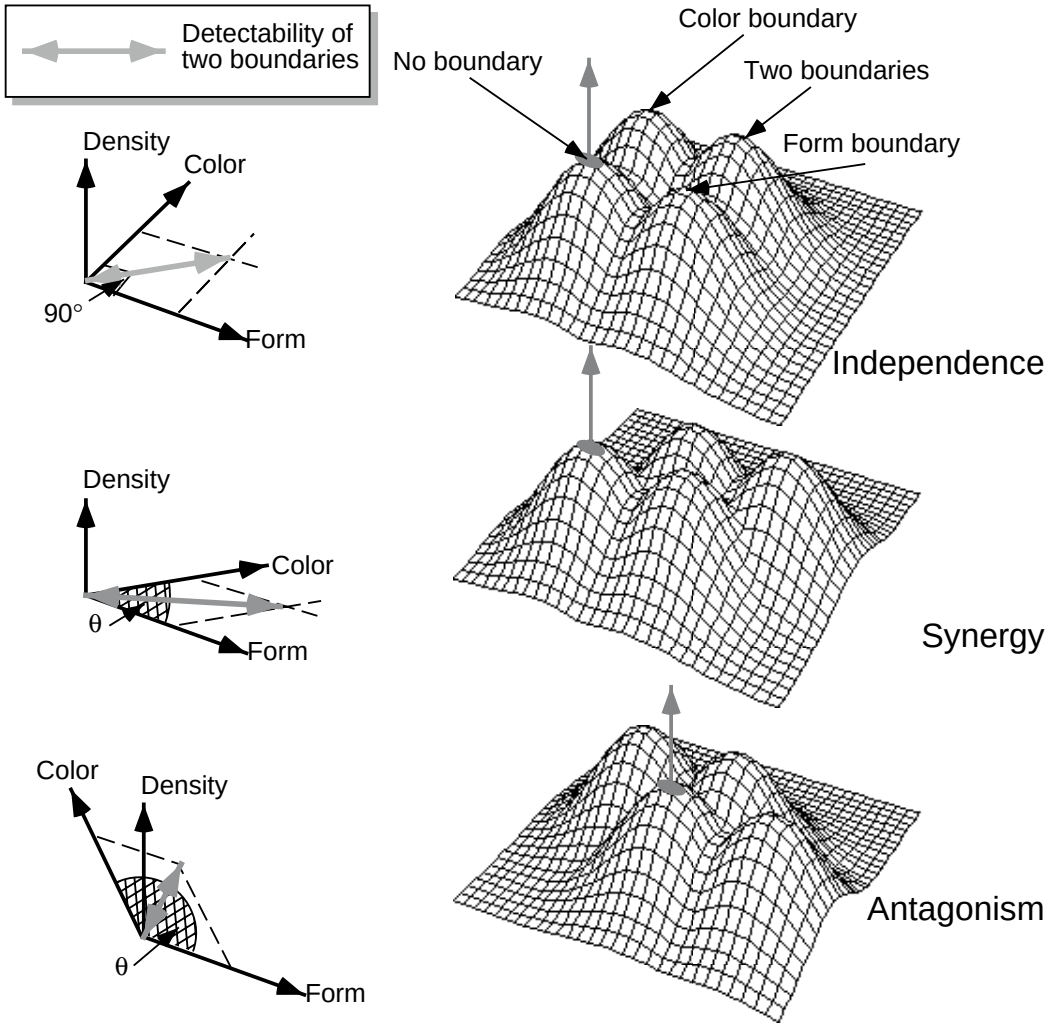

Figure 11. The perceptual effects of no-boundary ⟨0, 0⟩; color-boundary ⟨$\mathcal{C}$, 0⟩; form-boundary ⟨0, $\mathcal{F}$⟩; and two-boundary ⟨$\mathcal{C}$, $\mathcal{F}$⟩ targets when the form and color modules are independent, synergistic, and antagonistic.

activated primarily by form, but also, to some extent, by color). Once the module has been activated, the information stemming from the module's primary feature cannot be discriminated from information stemming from the other feature. Therefore, a failure of separability cannot account for different forms of dimensional interaction when the displays differ. The secondary dimension may inhibit or reinforce the module's activation by the primary dimension, but it must do so uniformly, regardless of display type.

Perceptual independence fails (case 2) when two modules communicate during processing. It is meaningful to speak of a failure of perceptual independence only if the contents of each module are different. Therefore, the distinction between perceptual independence and perceptual dependence is meaningful only when the two dimensions are, by and large, perceptually separable. Thus, we assume that each module begins processing information independently of the other. If the modules communicate, it would be in order to compare the information generated by their early independent processing and to subsequently process the information in a way that depends on the results of this comparison. So if we can infer that the interaction between the modules depends on the relation between the inputs to them—that is, if the display type determines whether antagonism or synergy is observed—we can conclude that we are dealing with a case of perceptual nonindependence. Otherwise, we cannot distinguish between a failure of perceptual separability and a failure of perceptual independence.

Method

Subjects

Five paid University of Virginia graduate or upper level undergraduate students were tested in Experiment 1, and 8 were tested in Experiment 2.

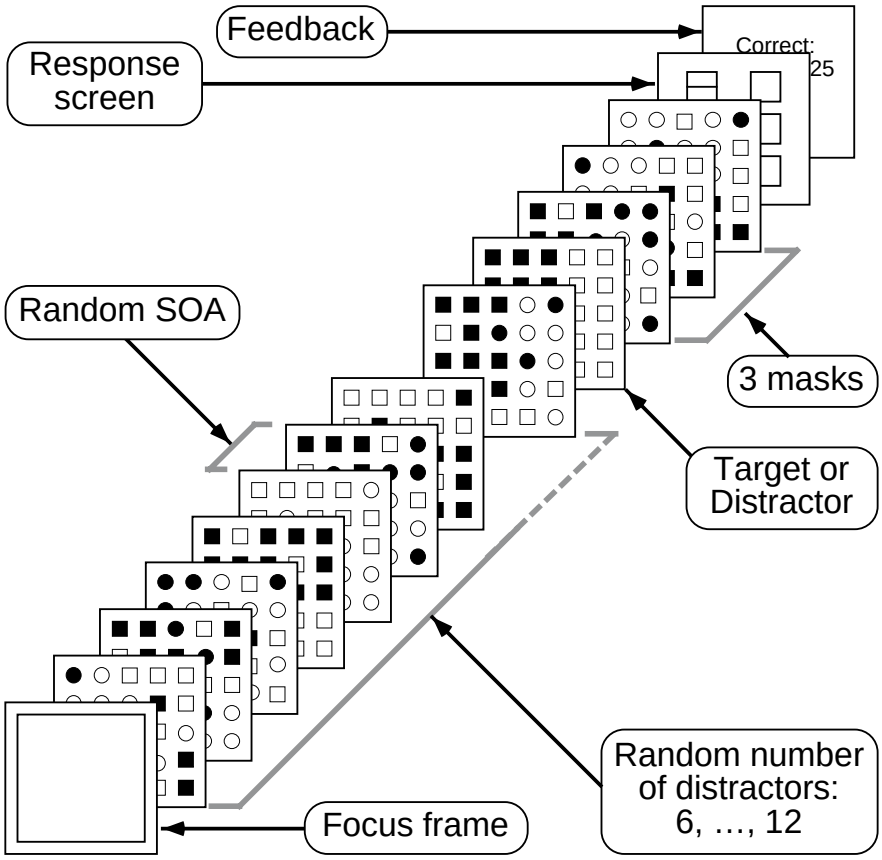

Figure 12. A trial of the gestalt detection paradigm.

Apparatus

Experiment 1 was conducted on a Masscomp 5000 with a 19-in. Sony Trinitron screen. Experiment 2 was conducted on a Silicon Graphics Indigo with a 16-in. Sony Trinitron screen. A viewing hood covered the computer screen. The hood kept the observers 61 cm from the screen and shielded them from ambient light.

Stimuli

Targets and distractors were square lattices of 25 colored forms. The two values of form were an equilateral octagon and an octagon with four long horizontal and vertical sides of equal length and four short sides of equal length at an angle of 45º. The two values of color were purple and light purple. During an observers' training sessions, we manipulated the differences between the values of form and color in order to achieve parity of performance between these two factors. We varied the ratio of the large side to the small side of the irregular octagon and the shade of the light purple. Each element subtended a visual angle of 0.6º. Each matrix subtended a visual angle of 6º.

We can generate 16 C (one color boundary) target types—2 color values $\times$ 2 form values $\times$ boundary positions. For instance, there are four ways to create a $\langle v_2, \varnothing \rangle$ target: color 1 on the right or on the left, combined with the choice of form 1 or form 2. If we add the 16 $\mathcal{F}$ targets, there are 32 single-boundary target types. When both form and color boundaries are present in a target, the boundaries can either be orthogonal ($\perp$), parallel ($\parallel$), or coincident ($\sim$). There are 64 dual-boundary target types—(2 versions of form $\times$ 4 locations for form boundary) $\times$ (2 versions of color $\times$ 4 locations for color boundary).

We described our procedure to generate distractors in Figure 9.

Procedure

On every trial, we presented a sequence of frames in rapid succession (Figure 12). This sequence contained several distractors, among which a target may be embedded. To prevent the observers from predicting when the target would appear (if it appeared), the number of distractors preceding it was a uniform integer-valued random variable on the interval [6,12]. To ensure that the observers did not continue to process the target after it disappeared, we presented three posttarget distractors that served as *masks*.

We measure the duration of each frame in 60ths of a second, called *ticks* (the refresh time of the computer screen). The stimulus onset asynchrony (SOA), in ticks, of every frame in the trial (both distractors and the target) was, in Experiment 1, a uniform integer-valued random variable on the interval [10,15] (167–250 msec in 16.7-msec steps). In Experiment 2, the interval was [2,7] (33–117 msec in 16.7-msec steps). The random variation was designed to prevent observers from developing a strategy for a particular SOA. Each trial began with a focus frame and ended with response and feedback screens.

On each trial, the observers had to decide whether or not a target had been presented on that trial. They indicated their response by moving a computer mouse so that the cursor was on one of the six icons shown in Figure 13 and by clicking on the mouse button to choose it. The six icons represented a 6-category rating scale, whose categories ranged from *I'm very sure a target appeared* to *I'm very sure a target did not appear*.

Each observer participated in approximately 30 sessions. Each session took an hour and consisted of 500 trials. The observers ran once a day, five days a week, until their obligation was complete. To better motivate them, we paid them only for correct responses. Ob-

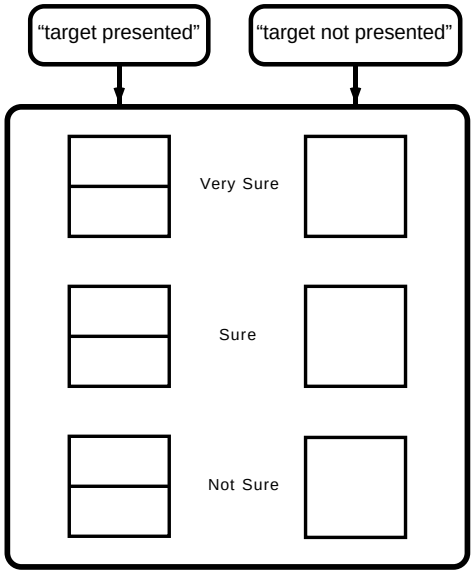

Figure 13. The response screen in the gestalt detection experiment.

servers earned approximately $5.50 each session. Observers practiced until their detection rates stabilized, and until we could equate their detection rates for C and $\mathcal{F}$ trials. This usually took between 1 and 2 weeks.

Results

We obtained a value of A_z—the area under the receiver-operating characteristic (ROC; Macmillan & Creelman, 1991)—for each observer, each duration, and each boundary type. For each observer, the estimate of A_z for each duration and each boundary type was based on about 500 observations. To obtain these estimates, we used ROCFIT, a program that calculates a maximum likelihood estimate of A_z from rating data (Metz, 1989). We used A_z, rather than d', because it is a nonparametric and, hence, more robust measure of detectability. As we explained above, we used the C and the $\mathcal{F}$ trials to obtain the region in which synergy holds for two-boundary targets. Because Equation 1 is expressed in terms of d', we used the following procedure: We (1) transformed the A_z data to d_a (Macmillan & Creelman, 1991, Equation 3.6), (2) applied Equation 1, with $\theta = \pi/2$ for the lower bound of the region of synergy (when the modules are independent) and $\theta = 0$ for the its upper bound (when the modules are identical), and (3) transformed the predicted d_a back to A_z. The predictions are based on the data of individual subjects and then averaged. The results appear in the bottom panels of Figure 14.

Single-Boundary Data

From the single-boundary graphs, we can see that our efforts to equate the detectability of single-form boundaries and single-color boundaries were not entirely successful. Most observers detected color boundaries somewhat better than they did form boundaries. We computed separate linear models for the short ($t_0 = 3$) and the long ($t_0 = 10$) experiments. In each, the dependent variable was A_z, and the independent variables were target SOA (treated as a continuous fixed variable), boundary type (treated as a discrete fixed variable), and observers (treated as a random discrete variable). For each experiment, we obtained a linear model (Table 2; standard errors are provided to the right of the coefficients of the model; divide these figures by 10^3).[3] Our linear models for the two experiments differed somewhat: Only with long durations (Table 2, columns 4 and 5) did we observe an effect of boundary type (the CF factor was coded +1 for color and −1 for form). However, the presence of an effect of boundary type does not affect our subsequent analyses.

The growth of $\widehat{A_z}$ as a function of SOA is virtually identical in the two experiments (slopes of 0.029 and 0.030, respectively). Regarding the intercepts, however, we cannot make strong comparative claims, since we adjusted the difficulty of the tasks to keep performance away from floor or ceiling. Any apparent continuity of the data between the two experiments should therefore be treated as happenstance. It is, however, true that it was harder to keep performance above floor for short durations than for long ones.

Dual-Boundary Data

In the top panels of Figure 14, we compare the detectability of orthogonal ($\perp$), parallel ($\parallel$), and coincident ($\sim$) boundaries to the region of synergy. Table 2 shows that the intercepts of the boundaries are ordered as expected: $A_z(\perp) < A_z(\parallel) < A_z(\otimes) < A_z(\sim) < A_z(\equiv)$, with the exception of a reversal for $t_0 = 10$, where the intercept for $A_z(\sim)$ is greater than the intercept for $A_z(\equiv)$, which we attribute to variability in the data.

Strong Modular Synergy

We first explore the strong version of the hypothesis of modular synergy, which has two components: (1) Modules enhance each other's performance when their input is consistent; $A_z(\sim) > A_z(\otimes)$, and (2) Modules inhibit each other's performance when their input is inconsistent, $A_z(\otimes) > A_z(\parallel)$ and $A_z(\otimes) > A_z(\perp)$.

To determine whether the performance of the color and form modules is synergistic when boundaries coincide, we computed, for the analysis of variance mentioned in the preceding paragraph, the contrast $A_z(\sim)$ versus $A_z(\otimes)$. We find that $A_z(\sim) > A_z(\otimes)$ for the long durations (Table 3, row 2), but not for the short ones (Table 3, row 1).

To determine whether the performance of the color and form modules is mutually inhibitory when boundaries do not coincide, we computed the contrast $A_z(\parallel)$ versus $A_z(\otimes)$. We find that $A_z(\parallel) < A_z(\otimes)$ for the long durations (Table 3, row 4), but not for the short ones (Table 3, row 3). However, the contrast $A_z(\perp)$ versus $A_z(\otimes)$ showed that when the boundaries are orthogonal, the modules are mutually inhibitory for both experiments (Table 3, rows 5 and 6).

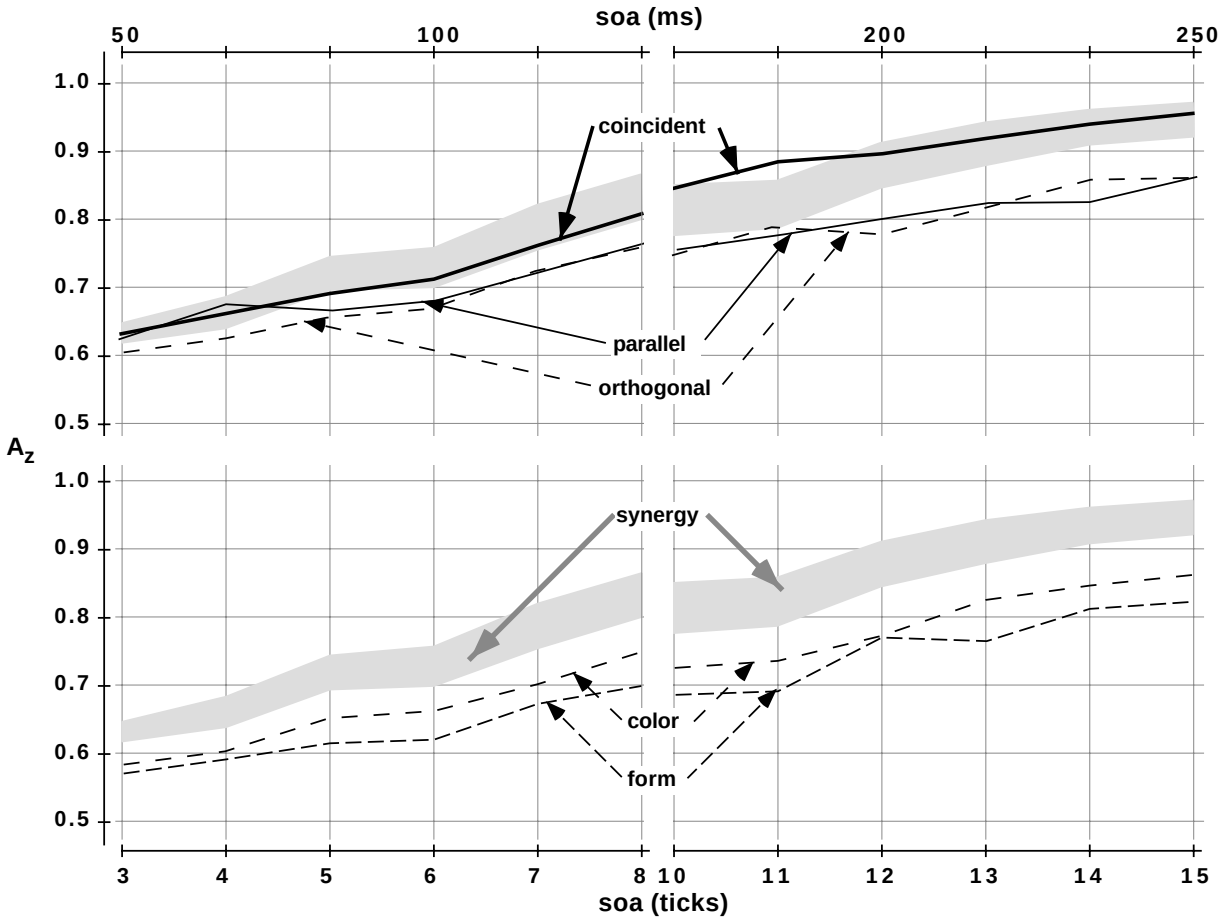

Figure 14. Observed detectability (A_z) of targets as a function of target stimulus onset asynchrony (SOA; in ticks, 1/60th sec; the equivalent value in milliseconds is given at the top of the graph). Data averaged over observers. Left panels: target durations 3 to 8 (t_0 = 3). Right panels: target durations 10 to 15 (t_0 = 10). Bottom: Single-boundary targets ($\mathcal{C}$ and $\mathcal{F}$) and predicted detectability of two-boundary targets under module synergy, bounded below by module independence and above by module identity (gray area). Top: Dual-boundary targets (coincident, parallel, and orthogonal) and predicted region of module synergy of two-boundary targets (same as in bottom panels). Estimates of A_z below the region of module synergy represent antagonism.

Weak Modular Synergy

Although the evidence just presented did not show that the modules interact at short durations when the boundaries are parallel, a weaker version of the hypothesis of modular synergy for parallel boundaries in the short-duration experiment can be supported—namely, that coincident boundaries are better detected than parallel ones, $A_z(\sim) > A_z(\|)$. This hypothesis does not distinguish between effects due to synergy, antagonism, or both. The contrast $A_z(\sim)$ versus $A_z(\|)$ (Table 3, rows 7 and 8) shows that coincident boundaries are more detectable than parallel boundaries for short and (of course) for long durations. (We did not find a statistically reliable difference between performance for parallel and orthogonal boundaries.)

Discussion

The data show that the form and color modules interact and, thus, cast doubt on a basic tenet of current theories of feature integration, the assumption of module independence.

For the longer durations, the data show that coincident boundaries are processed synergistically and that noncoincident boundaries trigger antagonism between the modules. For the shorter durations, the evidence in favor of module interaction is somewhat weaker. We were unable to show that coincident boundaries are detected better—or that parallel boundaries are detected worse—than would be expected if the modules were independent. However, two findings show that there is interaction even at short durations: (1) Orthogonal boundaries are detected worse than would be expected if the modules were independent, implying mutual antagonism between them, and (2) parallel boundaries are detected worse than coincident boundaries, implying either synergistic processing of coincident boundaries or antagonism between parallel boundaries.

These data imply a failure of perceptual independence, not of perceptual separability. If perceptual separability were imperfect (i.e., the dimensions were partially integral) and perceptual independence held, once information

Table 2
Linear Models for Detection Experiment

Coefficient	Experiment 1($t_0 = 3$) $\widehat{A_z}$	Experiment 1 $SE \times 10^3$	Experiment 2 ($t_0 = 10$) $\widehat{A_z}$	Experiment 2 $SE \times 10^3$
	Single Boundaries			
Constant	0.572	9.25	0.702	6.95
SOA $- t_0$	0.029	3.47	0.030	1.39
CF			0.019	3.57
	Dual Boundaries			
Constant	0.618	3.89	0.797	3.36
SOA $- t_0$	0.035	5.03	0.025	4.58
Boundary at t_0 (intercepts)				
⊥	−0.022	6.91	−0.046	12.8
∥	−0.0072	8.73	−0.040	12.8
⊗	−0.00630	6.91	−0.024	12.8
~	0.00844	7.11	0.059	12.8
≡	0.027	6.91	0.051	12.8
Boundary × (SOA $- t_0$)				
⊥	−0.00402	2.19	−0.002	2.51
∥	−0.00562	2.55	−0.044	2.51
⊗	0.00130	2.19	0.00719	2.51
~	−0.000324	2.22	−0.0035	2.51
≡	0.000867	2.19	0.00235	2.51

entered a module, information about the relationship between the two dimensions would be unavailable. Thus, in our experiments, we would have observed only synergy between modules, never antagonism. For example, consider *separable* dimensions *X* and *Y*—that is, variations in the stimulus with respect to *X* only affected module m_X, whereas variations with respect to *Y* only affected module m_Y. Under these assumptions, each module has the chance to detect a single boundary. Compare with this the case of imperfectly separable *X* and *Y*—that is, variations in the stimulus with respect to dimension *X* affected module m_X and also leaked to module m_Y and affected it weakly (*mutandi mutandis* for dimension *Y*). Suppose, furthermore, that we presented the observer a frame containing parallel boundaries: an *X* boundary on the left and a *Y* boundary on the right. Module m_X would see a strong boundary on the left and a weak boundary on the right (*mutandi mutandis* for module m_Y). Therefore, each module would have information regarding *two* boundaries, thereby producing a result that looked like synergy and precluding a result that looked like antagonism.

In contrast, if perceptual separability were complete while perceptual independence of the modules failed, some information about the relationship between the two dimensions would be available, since the failure of perceptual independence would presumably be in the service of intermodule communication. Because our data showed differential interaction between modules for different display types, we have demonstrated a failure of perceptual independence, but not of perceptual separability.

In order to confidently attribute the effects we have observed to a failure of perceptual independence, we must rule out the possibility that the effects are due to decision or response processes (for a discussion, see Cohen, 1997).

Table 3
Tests of Modular Synergy (Experiments 1 and 2)

Test	Experiment	Difference	*SE*	*F*	*df*	MS_{denom}	*p*(One-Tailed)
			$A_z(\sim) - A_z(\otimes)$				
1	1	0.0106	0.00645	0.681	1,4	0.0024	.228
2	2	0.0564	0.00902	9.816	1,3	0.0039	.0259
			$A_z(\otimes) - A_z(\|\|)$				
3	1	0.0183	0.0102	1.595	1,4	0.0029	.138
4	2	0.0446	0.00945	11.094	1,3	0.0021	.0224
			$A_z(\otimes) - A_z(\perp)$				
5	1	0.0288	0.00790	6.687	1,4	0.0019	.0305
6	2	0.0442	0.0115	7.405	1,3	0.0032	.0363
			$A_z(\sim) - A_z(\|\|)$				
7	1	0.0288	0.00633	10.378	1,4	0.0011	.0161
8	2	0.1010	0.0218	10.702	1,3	0.011	.0234

It is unlikely that our results are due to postperceptual processes. Decision and response processes cannot be responsible for synergistic effects, because they cannot increase the amount of perceptual information available. If, however, the observers did not do their best in some conditions and did as well as they could in others, the data may mimic synergy. For example, if the observers closed their eyes on a certain proportion of the single-boundary trials, we would underestimate the level of performance for dual-boundary trials predicted under the assumption of independence. Under these conditions, the quality of the observers' performance for dual-boundary trials would seem to be higher than that predicted by independence, even if the processes were actually independent. Therefore, to account for our results while claiming that perceptual independence holds, one must assume that the observers responded suboptimally in some, but not all, conditions. We do not see how our procedure could have permitted the observers to do so.

EXPERIMENTS 3, 4, AND 5
Boundary Localization

In Experiments 3–5, we used a boundary localization procedure. (We list the experiments in increasing order of the shortest SOA used; they were run in the reverse order.) The stimuli and procedure were identical to the boundary detections task described above, with two exceptions: (1) All trials contained a target, and (2) the observers were asked to identify exactly one of the boundaries present.

Method

Subjects

Five paid University of Virginia graduate or upper level undergraduate students were tested in Experiment 3, 4 were tested in Experiment 4, and 4 were tested in Experiment 5. None of these had participated in Experiments 1 or 2.

Apparatus

Experiments 3 and 4 were conducted on a Silicon Graphics Indigo with a 16-in. Sony Trinitron screen. Experiment 5 was conducted on a Masscomp 5000 with a 19-in. Sony Trinitron screen. A viewing hood covered the computer screen. The hood kept the observers 61 cm from the screen and shielded them from ambient light.

Stimuli

Targets and distractors were identical to those of the boundary detection task.

Procedure

The trial structure for the boundary localization task was identical to that for the boundary detection task, with the following exceptions. First, every trial contained a target. Second, the subjects' responses consisted of identifying the position of exactly one boundary in the target. The response screen was a four-alternative forced-choice display containing the four boundary positions (Figure 15). Although dual-boundary targets contained both a form and a color boundary, the observers' responses selected only a single boundary position, regardless of whether the target had one or two boundaries. The observers were not asked whether form or color defined the chosen boundary.

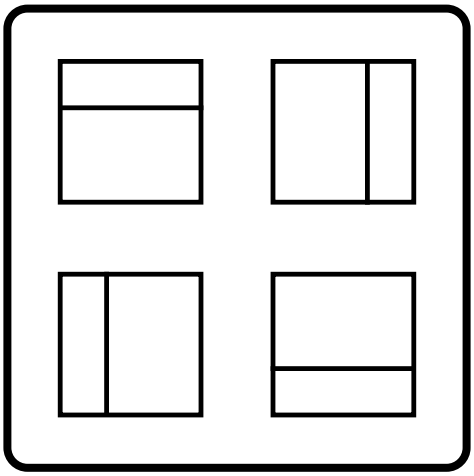

Figure 15. The response screen for the boundary localization task.

The SOA, in ticks, of every frame in the trial (both distractors and the target) was, in Experiment 3, a uniform integer-valued random variable on the interval [3,13] in steps of two; in Experiment 4, the interval was [6,16] in steps of two; and in Experiment 5, the interval was [10,15].

Each observer participated in approximately 30 sessions. Each session took an hour and consisted of 500 trials. The observers ran once a day, five days a week, until their obligation was complete. In an attempt to motivate the observers, we paid them only for correct responses. The subjects earned approximately $5.50 each session. They practiced until their localization rates stabilized. This usually took between 1 and 2 weeks. In the practice sessions, we adjusted the salience of the form and color boundaries until the probability of correctly localizing a form boundary equaled the probability of correctly localizing a color boundary, $p(\surd|\mathcal{F}) = p(\surd|C)$. To equate salience, we adjusted the similarity of the elements in each dimension. We made slight corrections after the practice sessions to ensure the stability of the observers' performance levels.

Results

We can estimate the predicted probability of correctly localizing a boundary in a two-boundary target, $p(\surd|\mathbf{2})$, if we assume that $p(\surd|\mathcal{F})$ and $p(\surd|C)$ are independent (as the blackboard models assume). The data provide a direct estimate of $p(\surd|\mathcal{F})$ and $p(\surd|C)$. To analyze the data, we compared the observer's $p(\surd|\mathbf{2})$ to the prediction based on the models of independence $p_{\otimes}(\surd|\mathbf{2})$. In the boundary localization task, there are four possible responses. The observer's response is correct if it corresponds to the location of one of the boundaries present in the target. We converted percent correct for each target type into a d'. To do this, we had to model the observer's decision space.

Decision Space Assumption

In each target, there are four potential boundary locations, $l \in \{1, 2, 3, 4\}$ (Figure 7), and two boundary types, $m \in \{c, f\}$. We denote the eight ensuing potential boundaries by $m_l \in \{c_1, \ldots, c_4, f_1, \ldots, f_4\}$, where c_1 refers to a color boundary in location 1. Every trial produces a level of activation, $R(m_l)$ (Figure 16), at each of the eight

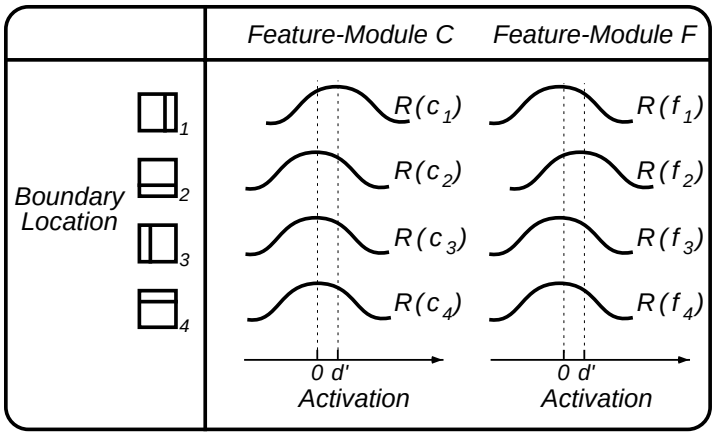

Figure 16. Graphical representation of the observer's internal space. The means of the distribution have been placed to represent the modules' response to a vertical right-hand color boundary orthogonal to a horizontal lower form boundary. Therefore, the means of the activation distributions $R(c_1)$ and $R(f_2)$ have been shifted to the right by d'.

potential boundaries, m_l. For locations in which there is no boundary, the corresponding activation is the noise distribution, $R(m_l) \sim \mathcal{N}(0, 1)$, a Gaussian with a mean of 0 and a standard deviation of 1, whereas the presence of a boundary in the target shifts the mean of the activations in the corresponding locations by d', so that $R(m_l) \sim \mathcal{N}(d', 1)$. This is the observer's decision space. In the single-boundary condition, the observer produces a correct response either (1) because the correct boundary location for the *correct* feature module had the highest activation or (2) because the correct boundary location for the *incorrect* feature module had the highest activation.

We propose a model of independence that estimates $p(\surd \mid \mathbf{2})$ on the basis of this assumption. We used this model as the basis for Monte Carlo simulations.

Model of Independence

Our model of independence assumes perceptual separability, perceptual independence, and decisional independence. According to this model, when two boundaries are present, the boundary chosen corresponds to the $R(m_l)$ with the highest activation. This model assumes that every $R(m_l)$ can be compared with every other $R(m_l)$, within and between feature modules (Figure 16).

To estimate $p(\surd \mid \mathbf{2})$, we first converted the probability correct for the single-boundary targets, $p(\surd \mid \mathbf{1})$, into d'. To do this, we ran a Monte Carlo simulation to find the function $d' = f[p(\surd \mid \mathbf{1})]$. Without loss of generality, we assumed that boundary $c1$ was present and, therefore, let $R(c_1) \sim \mathcal{N}(d', 1)$. For all $l \neq 1$, $R(m_l) \sim \mathcal{N}(0, 1)$. The simulation program randomly chose a value from the distribution of each representation. If $R(c_1) = \max[R(m_l)]$, the program deemed the answer to be correct. To construct a lookup table, in which one could enter a value of d' and read off the corresponding value of $p(\surd \mid \mathbf{1})$, we repeated this procedure 1,000 times for each $d' \in \{0.0, 0.05, 0.10, \ldots, 3.0\}$. To obtain the function $d' = f[p(\surd \mid \mathbf{1})]$, we fitted a function to the data in the table and obtained

$$d' = 4.03922\left(\sqrt[4]{p(\surd \mid \mathbf{1})} - \sqrt[4]{1 - p(\surd \mid \mathbf{1})}\right) + 1.07883, \quad (2)$$

which reproduced the values in our table to four decimal places.

To estimate $p(\surd \mid \mathbf{2})$, we used Equation 2 to convert the observer's proportion correct for each single-boundary target to a d'. For example, a target with coincident boundaries at location 1 (i.e., c_1 and f_1), $R(m_l) \sim \mathcal{N}(d', 1)$. For all $l \neq 1$, $R(m_l) \sim \mathcal{N}(0, 1)$. The simulation randomly sampled from each $R(m_l)$ distribution. If $R(m_l) = \max[R(m_l)]$, the answer was deemed correct. This procedure was repeated N times, where N represents the number of trials per condition. The outcome is a proportion correct for those trials.

This entire procedure was repeated 1,000 times. The mean proportion correct is the predicted outcome, given independence. The standard deviation of the simulated proportion correct is the standard error of the model, given N trials.

The Data

We obtained a value of

$$\lambda(\surd) = \ln \frac{1}{2} \frac{p(\surd)}{1 - p(\surd)}$$

—Tukey's (1977) folded logarithm transformation (fLog, a version of the *logit* transformation)—of the probability of correctly identifying a boundary, for each observer, duration, boundary type, and experiment. As we explained above, we used the single-boundary trials—$\mathcal{C}$ and $\mathcal{F}$ — under the assumption of our model of independence, to

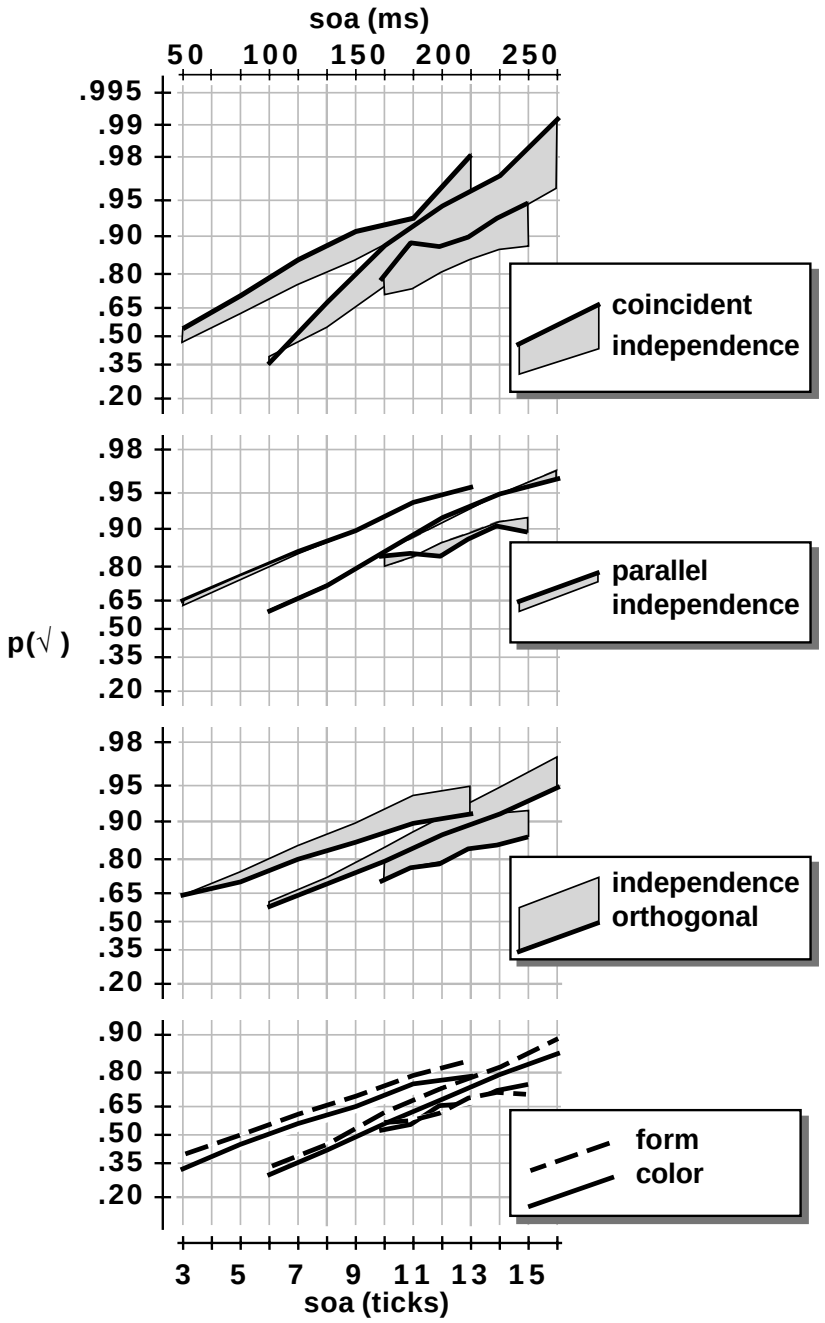

Figure 17. Bottom panel: Identification performance for single-boundary color (C) and form ($\mathcal{F}$) targets. In the three remaining panels, the data from each experiment are represented by a thick line; the predictions of our model of independence are represented by a thin line. Lower middle panel: Dual orthogonal boundaries ($\perp$). Performance is worse than that predicted by module independence, implying module antagonism. Upper middle panel: Dual parallel boundaries ($\|$). Performance is well predicted by module independence. Top panel: Dual coincident boundaries (~). Performance is better than predicted by the model, implying module synergy.

Table 4
Linear Models for Boundary Localization Experiments

		Experiment 3 ($t_0 = 3, t_6 = 13$)		Experiment 4 ($t_0 = 6, t_6 = 16$)		Experiment 5 $t_0 = 10, t_6 = 15$)	
Test	Coefficient	$\hat{\lambda}(\surd)$	*SE* $\times 10^2$	$\hat{\lambda}(\surd)$	*SE* $\times 10^2$	$\hat{\lambda}(\surd)$	*SE* $\times 10^2$
				Single Boundaries			
1	Constant	−0.402	2.69	−0.587	3.90	−0.538	6.13
2	SOA	0.136	2.06	0.103	2.08	0.089	0.80
3	*CF*	0.043	5.97	0.058	9.27	−0.005	5.82
				Noncoincident Dual Boundaries			
4	Constant	0.171	1.75	−0.062	3.33	0.176	5.55
5	SOA	0.149	2.07	0.117	2.43	0.093	1.05
	Boundary at t_0 (intercept)						
6	⊥	−0.146	4.06	−0.155	7.75	−0.197	4.60
7	⊗	0.040	4.06	0.394	7.75	0.066	4.60
8	||	0.044	4.06	0.456	7.75	0.028	4.60
				Coincident Dual Boundaries			
9	Constant	−0.208	2.40	−0.547	4.01	−0.454	7.08
10	SOA	0.244	3.27	0.184	3.72	0.156	1.50
	Boundary at t_0 (intercept)						
11	⊗	−0.032	2.26	0.059	6.14	0.010	7.62
12	~	0.051	2.27	0.082	6.14	0.193	7.62
	Boundary at × SOA						
13	⊗	−0.061	1.23	−0.045	1.68	−0.035	0.62
14	~	−0.003	1.27	−0.009	1.68	−0.016	0.62

obtain estimates of the performance that would be expected on dual-boundary trials if the color and form modules were independent.

Single-Boundary Performance

Before we examine the data that go to the heart of our argument, we examine the single-boundary data. Figure 17 combines the data of the three experiments. The bottom panel shows that we succeeded in equating the levels of detection performance for single-boundary color ($\mathcal{C}$) and form ($\mathcal{F}$) targets. Table 4 summarizes the linear models we fitted to the data. The model for single boundaries for Experiment 3 is $\lambda(\surd) = -0.402 + 0.136SOA + 0.043\mathcal{CF}$, where $\mathcal{C}$ targets are coded −1 and $\mathcal{F}$ targets are coded +1. Line 3 shows that the effect of boundary type (color vs. form) is 0.043 ± 0.0597, which means that it is indistinguishable from 0, which is also true of the two other experiments. We also note (line 2) that the growth of performance with SOA is on the order of 6 fLog units per second and that there is no statistical evidence that the three slopes differ.

Dual Boundary Performance

Orthogonal boundaries. To determine whether the color and form modules are antagonistic when boundaries are orthogonal (as Figure 17, lower middle panel, suggests) we computed, for the linear model mentioned in the preceding paragraph, the contrast $\lambda(\surd|\perp) - \lambda(\surd|\otimes)$. We find that $\lambda(\surd|\perp) < \lambda(\surd|\otimes)$ for the three experiments (Table 5, rows 1–3), but the support of Experiment 4 is weak. When we combine the evidence from the three experiments (Rosenthal & Rosnow, 1991, pp. 501–502), we

Table 5
Tests of Modular Synergy (Experiments 3–5)

Test	Experiment	Difference	*SE*	*F*	*df*	MS_{denom}	*p*(One-Tailed)
			$\lambda(\surd\|\perp) - \lambda(\surd\|\otimes)$				
1	3	−0.1856	0.05275	6.197	1,4	0.0834	.0338
2	4	−0.1940	0.08796	2.431	1,3	0.1860	.1085
3	5	−0.2634	0.05062	13.51	1,3	0.0616	.0175
			$\lambda(\surd\|\|) - \lambda(\surd\|\otimes)$				
4	3	0.3656	0.00955	732.84	1,4	0.0027	.0001
5	4	0.3102	0.08344	6.906	1,3	0.1672	.0393
6	5	0.3648	0.07891	10.672	1,3	0.1496	.0235

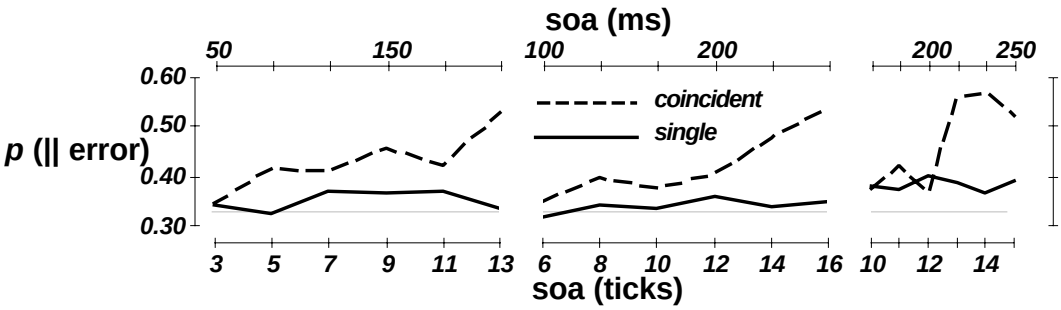

Figure 18. Comparison of the probability of parallel errors made on single-boundary trials to the probability of parallel errors made on coincident-boundary trials.

obtain $z = 2.985$, $p = .0014$, confirming that the color and form modules are antagonistic when the boundaries are orthogonal.

Parallel boundaries. Figure 17 (upper middle panel) suggests that the color and form modules are neither antagonistic nor synergistic when boundaries are parallel. We computed the contrasts $\lambda(\surd\,|||) - \lambda(\surd\,|\otimes)$ and found that all three F ratios were below 1, confirming the null effect we observed in the figure.

Coincident boundaries. The data for coincident boundaries allow us to draw firm conclusions regarding the synergy of modules. As Figure 17 (top panel) shows and rows 4–6 of Table 5 confirm, the data fall unequivocally above the levels of performance predicted by the model of independence. One further observation about coincident boundaries supports our contention that the two modules interact.

Parallel Errors

In Figure 18, we compare errors of localization that occurred on two types of trials: single-boundary trials and coincident-boundary trials. Since, in both types of trials, only one choice of four is correct, the observer can make one of three possible errors: a parallel error, in which he or she chooses a boundary parallel to the correct one, and two types of orthogonal errors. We consider erroneous parallel responses to be "better" than erroneous orthogonal responses, because they contain partial information. If errors are committed because the observer has no information about the nature of the target, the three types of errors should be equiprobable, and $p(\|\,\text{error}) = 1/3$. As SOA increases, the probability of parallel errors remains close to 1/3 for single boundaries but rises for coincident boundaries. As SOA is increased for coincident-boundary targets, there is a concomitant increase in the likelihood of a response that is closer to the truth than are the other responses. In other words, as the probability of a correct response increases, the quality of errors also improves. The fact that this rise does not occur when only one module is activated (i.e., in the single-boundary trials) implies that the effect is due to synergy between the modules.

Discussion

The data from the localization experiments provide three types of evidence that the color and form modules are not independent: (1) Orthogonal targets are processed less efficiently than our model of independence suggests; (2) coincident targets are processed more efficiently that our model of independence suggests; and (3) as the target lengthens, observers not only make fewer errors, but the errors they commit become concomitantly "better."

GENERAL DISCUSSION

The data from both the boundary detection and the boundary localization experiments show that the modules that process color and form are not independent. In both the detection and the localization experiments, we show that (i) coincident targets are processed synergistically by the color and form modules and (2) orthogonal targets are processed antagonistically by the color and form modules. These findings are consistent with the research of Cohen (1997), discussed in the introduction.

Psychometric Functions

The gestalt detection paradigm is a rigorous test of perceptual grouping because it requires subjects to segregate boundary information in both time and space. We obtained psychometric functions describing the growth of this grouping process over time. Even though different observers participated in these experiments, in the boundary detection task (Experiments 1 and 2), the slopes of these functions were similar for long and short target durations. In the boundary localization task, these functions were also similar, if not identical, across experiments. This suggests that we are assessing a common, low level perceptual process.

The psychometric functions from the boundary detection task support the principle of modularity—that is, color and form are processed in different modules. The subjects were less sensitive to the coincident target than one would predict if color and form were processed in the same module. As we remarked in the introduction, a

modular organization is advantageous for a system that must process complex information.

The psychometric functions from both the boundary detection and the boundary localization tasks refute the principle of modular independence. The subjects were more sensitive to the information in the coincident target than would be predicted by independence, and they were more sensitive to this information than to the information in the parallel and orthogonal targets. The subjects' sensitivity to the coincident target above that predicted given independent modules cannot be the result of either decision or response processes (for a detailed argument, see Cohen, 1997). This synergy must be a result of an interaction between the color and the form modules. The interaction between the color and the form modules is also evidenced in the subjects' reduced sensitivity to the parallel and orthogonal targets. The data show that when the boundaries of different features diverge, the color and form modules inhibit the activation of those boundaries. Thus, the subjects' sensitivity to both boundaries is reduced.

Although the violation of the principle of modular independence may cost the system some of the computational efficiency achievable by independent modules, it has an important benefit: There is no need for an agent to routinely combine the information processed by each module. If the agent is slow, its elimination in favor of direct interaction may increase the overall speed of processing.

Our account of the antagonism of modules could be wrong, in the following way. Consider the orthogonal boundary trials in the boundary detection experiment. We have assumed that each module "sees" one boundary and that the low performance in these trials is due to module antagonism. Suppose, however, that the observer does not see one color boundary and one form boundary but, rather, four rectangular patches, each defined by a different combination of the two colors and the two forms. Suppose, furthermore, that the observer was confused by this percept, because the frame would resemble a distractor more than it would a target.

Note, however, that this confusing percept could occur only if the color and form features had been integrated preattentively. So a critic who wished to adopt this view could argue that we have not observed antagonism between the modules but would have to concede that feature integration routinely occurs preattentively, which is our central contention.

The Role of Attention in Feature Integration

Our data show that the integration of feature dimensions occurred without the help of focused attention. The gestalt detection paradigm was designed so as to prevent subjects from using focused attention to detect the boundaries. The task required perceptual popout to occur across both space and time. Because the task required popout only of boundaries that extended over the entire breadth of the display and the displays were too large to be foveated, focused attention could only aid in the task if the subjects could scan the display. Because the majority of the displays were presented too quickly for subjects to shift their eyes even once, it is unlikely that focused attention could have influenced the experiment. Yet, the data show effects of modular interaction based on spatial information even at the shortest durations. This shows that focused attention is not necessary for feature integration.

Many early-selection theories of feature processing suggest that focused attention is necessary for feature integration (Treisman, 1977; Treisman & Gelade, 1980; Treisman & Schmidt, 1982; Wolfe & Cave, 1990; Wolfe et al., 1989). This claim is often supported by the results of the visual search task. The results of this task show that observers generally shift their focused attention to individual elements, to detect the the presence of a target in a conjunctive display.

Our results present a conundrum: Given the results of the visual search task, what is the role of focused attention in feature integration? We conjecture that focused attention is not necessary for feature integration under normal circumstances but that, when an inconsistency exists in the environment, focused attention may be called on to resolve that inconsistency. Specifically, under most circumstances, discontinuities in different feature dimensions produce coincident boundaries around objects. The visual system processes this information without the influence of focused attention. However, when discontinuities in different feature dimensions produce inconsistent or divergent boundaries, the visual system may inhibit those boundaries. Such a situation may trigger the activation of focused attention. The role of focused attention in such a situation is unknown. It may range anywhere from simply providing more precise information (because of the high density of photoreceptors in the retina) to actively conjoining features (as was suggested by Triesman and Wolfe).

Some Theory

Our data show that the visual system actively amplifies the activation of coincident boundaries created by different features and inhibits the activation of disparate boundaries created by different features. The effects of this interaction between modules is evident at stimulus presentations as short as 50 msec (Experiment 1) and increases as stimulus presentation times increase (parallel errors in Experiments 3–5). Cohen (1997) has called this pattern of data *active signal modulation*: The spatial location of discontinuites processed by each module is shared between them.

What is the goal of this computation? (This is Marr's 1982 computational theory question). We do not think that its goal is feature integration, but rather the early stages of the detection of objects. We conjecture that the function of each of the retinotopic maps found in the cortex is to offer a putative parsing of the visual field by performing grouping and segregation. In other words, they produce boundaries between regions in the visual field. Objects in the visual world produce coincident boundaries

in the visual field in many feature dimensions; for example, a tree trunk will produce coincident color and texture boundaries at its edge. Thus to maximize the probability that the visual system will correctly identify the boundary of an object, it can use a divide-and-conquer strategy. In such a system, a synergistic exchange of information among modules can insure that they only output consistent boundaries without requiring that they output identical boundaries. In a system that worked as a blackboard system, each module would inscribe its output on the cortical equivalent of a blackboard, and an agent would resolve whatever differences it found in the outputs of the modules. A side effect of this process would be feature integration.

In regions of the visual field where some modules detected boundaries and others did not, a synergistic exchange of information could inhibit the detection of nonexistent objects. The gradual relaxation toward a solution explains the positive correlation between SOA and module dependence as shown in the analysis of parallel errors in Experiments 3–5.

CONCLUSION

In summary, we presented a new task, the gestalt detection paradigm, that directly assesses the perceptual independence of features. We conducted five experiments, using variations of this paradigm to assess the perceptual independence of color and form. The data show that color and form are not processed independently and, further, that the interaction between modules is based on the spatial relationship of discontinuities in the feature modules.

REFERENCES

Alexander, C. (1964). *Notes on the synthesis of form*. Cambridge, MA: Harvard University Press.

Ashby, F., & Townsend, F. T. (1986). Varieties of perceptual independence. *Psychological Review*, **93**, 154-179.

Averbach, E., & Coriell, A. S. (1961). Short-term memory in vision. *Bell System Technical Journal*, **40**, 309-328.

Beck, J., & Ambler, B. (1973). The effects of concentrated and distributed attention on peripheral acuity. *Perception & Psychophysics*, **14**, 225-230.

Bergen, J. R., & Julesz, B. (1983a). Parallel versus serial processing in rapid pattern discrimination. *Nature*, **303**, 696-698.

Bergen, J. R., & Julesz, B. (1983b). Rapid discrimination of visual patterns. *IEEE Transactions on Systems, Man, & Cybernetics*, **13**, 857-863.

Callaghan, T. C. (1984). Dimensional interaction of hue and brightness in preattentive field segregation. *Perception & Psychophysics*, **36**, 25-34.

Callaghan, T. C. (1989). Interference and dominance in texture segregation: Hue, geometric form, and line orientation. *Perception & Psychophysics*, **46**, 299-311.

Callaghan, T. C., Lasaga, M. I., & Garner, W. R. (1986). Visual texture segregation based on orientation and hue. *Perception & Psychophysics*, **39**, 32-38.

Cavanagh, P., Arguin, M., & Treisman, A. (1990). Effect of surface medium on visual search for orientation and size features. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 479-491.

Chmiel, N. (1989). Response effects in the perception of conjunctions of color and form. *Psychological Research*, **51**, 117-122.

Cohen, D. J. (1997). Visual detection and perceptual independence: Assessing color and form. *Perception & Psychophysics*, **59**, 623-635.

Enns, J. (1986). Seeing textons in context. *Perception & Psychophysics*, **39**, 143-147.

Erman, L. D., Hayes-Roth, F., Lesser, V. R., & Reddy, D. R. (1980). The hearsay-ii speech–understanding system: Integrating knowledge to resolve uncertainty. *ACM Computing Surveys*, **12**, 213-253.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, **1**, 1-47.

Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.

Garner, W. R., & Morton, J. (1969). Perceptual independence: Definitions, models, and experimental paradigms. *Psychological Bulletin*, **72**, 233-259.

Gottwald, R. L., & Garner, W. R. (1972). Effects of focusing strategy on speeded classification with grouping, filtering, and condensation tasks. *Perception & Psychophysics*, **11**, 179-182.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Green, M. (1991). Visual search, visual streams, and visual architectures. *Perception & Psychophysics*, **50**, 388-403.

Gurnsey, R., & Browse, R. A. (1987). Micropattern properties and presentation conditions influencing visual texture discrimination. *Perception & Psychophysics*, **41**, 239-252.

Jepson, A., & Richards, W. (1992). A lattice framework for integrating vision modules. *IEEE Transactions on Systems, Man, & Cybernetics*, **22**, 1087-1096.

Julesz, B. (1981). Textons, the elements of texture perception, and their interactions. *Nature*, **290**, 91-97.

Julesz, B. (1984). A brief outline of the texton theory of human vision. *Theories in Neuroscience*, **7**, 41-45.

Julesz, B., & Bergen, J. R. (1983). Textons, the fundamental elements in preattentive vision and perception of textures. *Bell System Technical Journal*, **62**, 1619-1645.

Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, **111**, 228-238.

Kadlec, H., & Townsend, J. T. (1992). Signal detection analyses of dimensional interactions. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 181-231). Hillsdale, NJ: Erlbaum.

Kleiss, J. A., & Lane, D. M. (1986). Locus and persistence of capacity limitations in visual information processing. *Journal of Experimental Psychology: Human Perception & Performance*, **12**, 200-210.

Kubovy, M., & Cohen, D. J. (1991, November). *Are texture segregation by form and color independent of each other?* Paper presented at the 32nd Annual Meeting of the Psychonomic Society, San Francisco.

Kubovy, M., & Cohen, D. J. (1992, November). *Further tests of feature-module architectures for color and form*. Paper presented at the 33rd Annual Meeting of the Psychonomic Society, St. Louis.

Levitt, J. R., Yoshioka, T., & Lund, J. S. (1994). Intrinsic cortical connections in macaque visual area V2: Evidence for interaction between different functional streams. *Journal of Comparative Neurology*, **342**, 551-570.

Livingstone, M. S., & Hubel, D. H. (1987). Connections between layer 4B of area 17 and the thick cytochrome oxidase stripes of area 18 in the squirrel monkey. *Journal of Neuroscience*, **7**, 3371-3377.

Livingstone, M. S., & Hubel, D. H. (1988). Segregation of form, color, movement, and depth: Anatomy, physiology, and perception. *Science*, **240**, 740-749.

Macmillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide*. Cambridge: Cambridge University Press.

Marr, D. (1976). Early processing of visual information. *Philosophical Transactions of the Royal Society of London, Series B*, **275**, 483-524.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.

Merigan, W. H., & Maunsell, J. H. R. (1993). How parallel are the primate visual pathways? *Annual Review of Neuroscience*, **16**, 369-402.
Metz, C. E. (1989). *ROCFIT, IBM-PC version* (Tech. Rep.). Chicago: University of Chicago, Department of Radiology and the Franklin McLean Memorial Research Institute.
Mitchel, D. C. (1979). The locus of the experimental effects in the rapid serial visual presentation (RSVP) task. *Perception & Psychophysics*, **25**, 143-149.
Monheit, M. A., & Johnston, J. C. (1994). Spatial attention to arrays of multidimensional objects. *Journal of Experimental Psychology: Human Perception & Performance*, **20**, 691-708.
Mordkoff, J. T., & Yantis, S. (1993). Dividing attention between color and shape: Evidence of coactivation. *Perception & Psychophysics*, **53**, 357-366.
Phillips, W. A., Kay, J., & Smyth, D. (1995a). The discovery structure by multi-stream networks of local processors with contextual guidance. *Network: Computation in neural systems*, **6**, 225-246.
Phillips, W. A., Kay, J., & Smyth, D. M. (1995b). How local cortical processors that maximize coherent variation could lay foundations for representation proper. In L. S. Smith & P. J. B. Hancock (Eds.), *Neural computation and psychology: Proceedings of the 3rd Neural Computation and Psychology Workshop (NCPW3), Stirling, Scotland, 31 August–2 September 1994* (pp. 117-136). London: Springer-Verlag.
Poggio, T., Gamble, E. B., & Little, J. J. (1988). Parallel integration of vision modules. *Science*, **242**, 436-440.
Potter, M. C. (1986). Pictures in sentences: Understanding without words. *Journal of Experimental Psychology: General*, **115**, 281-294.
Rosenthal, R., & Rosnow, R. L. (1991). *Essentials of behavioral research: Methods and data analysis* (2nd ed.). New York: McGraw-Hill.
Rumelhart, D. E. (1989). The architecture of mind: A connectionist approach. In M. I. Posner (Ed.), *Foundations of cognitive science* (pp. 133-160). Cambridge, MA: MIT Press.
Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
Sternberg, S. (1970). Memory scanning: Mental processes revealed by reaction-time experiments. In J. S. Antrobus (Ed.), *Cognition and affect* (pp. 13-58). Boston: Little, Brown.
Tanner, W. P., Jr. (1956). Theory of recognition. *Journal of the Acoustical Society of America*, **28**, 882-888.
Treisman, A. (1977). Focused attention in the perception and retrieval of multidimensional stimuli. *Perception & Psychophysics*, **22**, 1-11.
Treisman, A. (1986). Properties, parts, and objects. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance: Vol. II. Cognitive processes and performance* (pp. 35-1–35-70). New York: Wiley.
Treisman, A. (1988). Features and objects: The fourteenth Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*, **40A**, 201-237.
Treisman, A. (1990a). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception & Performance*, **114**, 285-310.
Treisman, A. (1990b). Variations on the theme of feature integration: Reply to Navon. *Psychological Review*, **97**, 460-463.
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**, 15-48.
Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, **14**, 107-141.
Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, **16**, 459-478.
Tukey, J. W. (1977). *Exploratory data analysis*. Reading, MA: Addison-Wesley.
Ullman, S. (1991). Tacit assumptions in the computational study of vision. In A. Gorea (Ed.), *Representations of vision: Trends and tacit assumptions in vision research* (pp. 305-317). Cambridge: Cambridge University Press.
Wolfe, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, **1**, 202-238.
Wolfe, J. M., & Cave, K. R. (1990). Deploying visual attention: The guided search model. In A. Blake & T. Troscianko (Ed.), *AI and the eye* (pp. 79-103). New York: Wiley.
Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model of visual search. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 419-433.
Young, S. R. (1984). RSVP: A task, reading aid, and research tool. *Behavior Research Methods, Instruments, & Computers*, **16**, 121-124a.

NOTES

1. Monheit and Johnston (1994) conducted an experiment that appears to address the issue of independence of feature processing, but, in fact, does not. They observed positive correlations between the detection of color and form when they resided in a single object. In response to a query from M.K., J. C. Johnston (personal communication, January 6, 1995) wrote: "If it is granted that spatial attention induces strong positive correlations (because the two features in an object are in the same location/object) then there is no further reason to invoke any dependence in the feature processing itself. Nor is there any reason, of course, to disbelieve that there are correlations at the feature level. I do not know of any way to distinguish between the claims that the feature processing per se at an early extraction level are correlated, and the claim that the correlation is introduced during the transfer of already extracted features to post-attentional systems."

2. Defined by *The American Heritage Dictionary* as "a single picture on a roll of movie film."

3. For example, the predicted value of A_z for form and SOA = 12 is $A_z = 0.702 + 0.030 \cdot (12 - 10) - 0.019 = 0.743$. To compute the standard error of this value, we combine the standard errors of the terms as follows: $SE = \sqrt{0.00695^2 + 0.00139^2 + 0.00357^2} = 0.00794$. Thus, an approximate 95% confidence interval on the value of A_z we just calculated is $A_z = 0.702 \pm 2 \cdot 0.00794 = [0.694, 0.710]$.

(Manuscript received November 10, 1997;
revision accepted for publication July 13, 1998.)