

Theory and Evaluative Reviews

Visual search, visual streams, and visual architectures

MARC GREEN

Trent University, Peterborough, Ontario, Canada

Most psychological, physiological, and computational models of early vision suggest that retinal information is divided into a parallel set of feature modules. The dominant theories of visual search assume that these modules form a "blackboard" architecture: a set of independent representations that communicate only through a central processor. A review of research shows that blackboard-based theories, such as feature-integration theory, cannot easily explain the existing data. The experimental evidence is more consistent with a "network" architecture, which stresses that: (1) feature modules are directly connected to one another, (2) features and their locations are represented together, (3) feature detection and integration are not distinct processing stages, and (4) no executive control process, such as focal attention, is needed to integrate features. Attention is not a spotlight that synthesizes objects from raw features. Instead, it is better to conceptualize attention as an aperture which masks irrelevant visual information.

Two factors make early vision a difficult computational problem. First, the solution space is combinatorially explosive because images contain a large number of feature dimensions. Second, computation must be rapid, so processing time is limited. Many theories suggest that the visual system solves these problems by means of a divide-and-conquer strategy; the retinal image is decomposed into an array of separate representations that are processed in parallel. Authors have called these representations by such names as "feature spaces," "feature maps," "transform spaces," "intrinsic images," and "parameter spaces." Because of the minor theoretical distinctions among these terms, I will use the neutral term *feature modules*.

Although they differ in details, feature-module theories agree that early processing consists of two stages. The retinal image is initially coded by an array of detectors, each simultaneously tuned to many features (color, orientation, binocular disparity, etc.). Further, this *retinotopic map* is topographic; neighboring detectors represent the activity at neighboring image points. The retinotopic map then projects in parallel to a set of feature modules. Each module registers only one (or at most, a small number) of the image features coded in the retinotopic map. Lastly, it has also been suggested that much (Ballard, 1985, 1986; Barlow, 1981), or even possibly all (Treisman & Gelade, 1980), topographic information is lost during the mapping.

Feature-module theories have become popular in physiology, computational vision, and psychology (both psychophysics and cognition). This is not surprising, because division into "nearly decomposable" subproblems (Simon, 1969) is the best domain-independent method for solving complex problems. Although most researchers have adopted the divide-and-conquer strategy, the theories of different disciplines have been presented in different languages and rationales. This has prevented any serious attempt to create a single, unified framework for analyzing early vision. The existing literature might be termed "pick and choose." An author with expertise and knowledge of one discipline gratuitously chooses a few popular (and usually secondary-source) articles from another. The author then misunderstands and overinterprets the data from the unfamiliar field while ignoring a vast literature that is inconsistent with the author's theory.

In this article, I use visual architecture as a framework for combining concepts from diverse disciplines. Architecture provides a good level of analysis because it is a high-level construct, so a system can be described without having to specify too much detail. This may seem uncomfortable to psychologists, who are used to forming very specific microtheories. However, any attempt at interdisciplinary synthesis is going to require a less detailed level of analysis.

In this article, only feature-module architectures will be considered. The architecture of a system specifies the atomic components and their interconnections. This is perhaps the most basic part of any theory, because architecture provides the major constraints on a computational system. In the context of current search theories,

This research was supported by NSERC Grant OPB0046297 to M. Green. Correspondence should be addressed to Marc Green, Computer Studies Programme, Trent University, Peterborough, Ontario K9J 7B8, Canada.

architectural issues reduce to the following: (1) Do the feature modules communicate indirectly or directly? (2) Is there a separate featureless module for keeping track of locations (Treisman & Gelade, 1980)? And (3) is the integration of modules sequential or parallel? A complete theory would also address representational questions: Which features are represented in individual modules, and how are they represented? Lastly, a theory must also describe an algorithm, a set of step-by-step procedures, that runs on the architecture and uses the representation to perform a specific task.

By focusing on architecture, one can avoid detailed discussion of representation and algorithm. This is desirable, since as the discussion will show, current data are too difficult to interpret and too inconclusive to permit explanation of a detailed theory. Instead, architectural analysis focuses on a few large and fundamental issues that constrain more detailed theorizing. For example, it will be clear that current visual search theories depend on an unlikely visual architecture. I will offer as an alternative not a specific theory, but a class of theories.

The remainder of this article contains three sections. First, I review the physiological, computational, and psychological rationales for feature modules. In the subsequent section, I present alternative architectures and note their fundamental differences. Lastly, I examine evidence to evaluate the candidate architectures. The discussion will center on theories of visual search and especially feature-integration theory, because (1) it is an extreme example of a blackboard theory, and (2) it has generated a large body of research that can be used to evaluate architectures.

Feature-Module Rationales

Each discipline justifies feature modules in a different way. The following discussion is a brief outline of the rationales used in different research areas. The goal is merely to demonstrate the overlap of fundamental concepts and to provide a backdrop for the later discussion. The material will be presented uncritically, with editorial comments reserved for the subsequent sections.

Physiology. Physiological evidence suggests that the brain is organized to facilitate a divide-and-conquer strategy (DeYoe & van Essen, 1988; Hubel & Livingstone, 1985, 1987; Livingstone & Hubel, 1987, 1988; Phillips, Zeki, & Barlow, 1984). The early (Figure 1) part of the visual pathways is divided into two streams specialized for different image features. There is disagreement over the extent of specialization, but the extreme view says that the "parvocellular" pathway carries information on color and fine spatial detail, while the parallel "magnocellular" stream codes motion and contrast of coarse form. This is basically an extension of the older dichotomy (Breitmeyer & Ganz, 1976; Green, 1984) that postulated the functional division between a high spatial frequency, low temporal frequency "sustained channel" and a low spatial frequency, high temporal frequency "transient channel." It resembles the magno-parvo dichotomy, except that it makes no mention of color.

VISUAL STREAMS (Hubel & Livingstone)

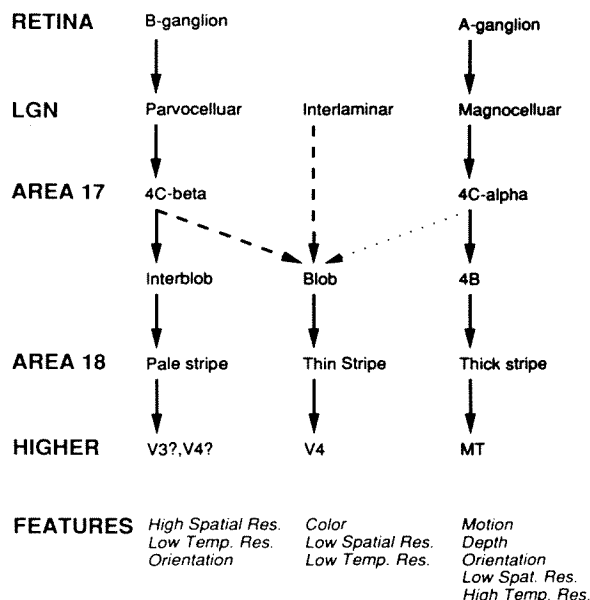


Figure 1. Visual streams proposed by Livingstone and Hubel (1987). Dashed and dotted lines represent weak and possible but unproved connections.

At the primary visual cortex (V1), the two streams become three as they project to extrastriate areas. The magnocellular stream synapses at layer 4B in V1 and then projects to both the middle temporal area (MT) and the "thick stripe" area of V2, which are connected to each other. Both MT and thick stripe then connect to V4. The parvocellular pathway splits and synapses at the "blob" and "interblob" regions of V1. From the interblob area, cells project to the "pale stripe" (also called the "interstripe" area) of V2 and then V4. The blob region, which may also receive some magnocellular input, projects to the "thin stripe" area of V2 and then to V4. There are other lesser connections, as well as reciprocal inputs between some areas. Although there is again dispute over both the physiological and the anatomical distinctness of the streams, the proposed specialization is shown in Figure 1. To oversimplify, there are distinct streams for (1) color, (2) fine spatial information, and (3) motion, binocular disparity, and coarse spatial information.

This organization reflects the limited connectivity (Barlow, 1981, 1986; Cowey, 1979) possessed by cortical neurons; one cell can easily communicate with another only if the two cells are in physical proximity. Image segmentation, however, often requires the linking of pieces from different areas of the retina and hence distant parts of the retinotopic map. Feature modules solve this problem by collecting information on each feature in a single cortical site.

Computation. Feature modules also reduce the complexity of computation. A common technique (Ballard, 1984, 1985, 1986; Ballard, Hinton, & Sejnowski, 1983; Barlow, 1986) is to divide retinotopic information into feature modules by means of a Hough transform, a method of mapping from one coordinate system to another. In vision, images are initially represented in a spatial coordinate system. Figure 2, for example, shows a collection of line segments in a coordinate system defined by x and y directions in space. A Hough transform maps this topographic representation into a new system where the axes are image parameters, rather than spatial coordinates. Since lines can be specified by the equation $y = mx + b$, they can be represented into an m/b space. Note that the new coordinate system has lost all spatial information.

The Hough transform brings together information on similar features located throughout the visual field. This is useful for trying to segment images when an occluding surface has fragmented the scene. The two segments, with different x/y locations but identical m/b locations, might be interpreted as a single line that is broken by an occluding surface. The new representation is therefore resistant to both noise and occlusion.

There is no theoretical limit on the number of possible dimensions. For example, the coordinate space could have additional axes for color, binocular disparity, velocity, and so forth, as well as for x and y spatial locations. As the number of features and dimensionality of the space increases, however, the number of coding units in the feature module grows exponentially. If all features of an image were coded in a single parameter space—say, in the retinotopic map—the resulting computation would be intractable. This problem can be solved by dividing the single large feature module into a collection of smaller subspaces, where each subspace represents a limited number of image properties. For example, a space that coded 10 values on 10 feature dimensions would require 10,000,000,000 ($= 10^{10}$) coding units. If the 10-dimensional space were divided into two 5-dimensional spaces, then the coding would require only 200,000 ($= 10^5 + 10^5$) units. Durbin and Mitchison (1990) tie physiological and computational rationales together by suggesting




Figure 2. Example of a Hough transform. Images represented in a spatial coordinate system are mapped to a new coordinate system where the axes are feature values.

that the visual pathways from the retina to the cortical surfaces are organized to reduce the parameter space.¹

Psychology. Lastly, feature-module architectures are commonly used in psychological theories of early vision, especially search and image segmentation. The dominant theory is called “feature-integration theory.” Although this theory is usually credited to Treisman and Gelade (1980), Attneave (1974) first stated the basic premise:

If, as we have reason to believe, color and form are processed in separate parts of the nervous system, why does one not simply perceive circle, triangle, blue, green without knowing which form has which color? The simple answer, I think, is that blue and circle are tagged to the same spatial location. (p. 109)

Attneave, however, did not specify how the features are combined or where the spatial information resides. Treisman and Gelade (1980) extended Attneave’s idea by suggesting that focal attention combines features in a “master map of locations.” Their theory suggests that early vision first analyzes images into separate features, each represented in its own “map.” It is unclear, however, exactly what is mapped. Each map seems to code not a single feature but a single feature value (Treisman & Sato, 1990). Moreover, the axes of the map are unspecified. It is tempting to assume that they are spatial dimensions, but the theory is vague about this point. Sometimes there are suggestions that feature maps are topographic (Treisman & Gormican, 1988), but the theory also says that features are free-floating with respect to location. Topographic information resides in a separate master map of locations, which has links into the feature maps. The two maps are independent until focal attention scans the master map. If there is topographic information in feature maps, it has no role in either feature detection (which uses pooled activity) or feature integration (which requires focal attention).

The theory also offers an algorithm for search. Feature maps register only pooled activity, which is the sum of activity produced throughout the visual field by the relevant feature. When searching for a unique feature, the observer can apparently detect the target from pooled activity within a map. Search for specific combinations of features—conjunction search—cannot be performed by examining a single map. The feature maps must be integrated by focusing attention in the master map. When the attentional spotlight rests on a place in the master map, links to the feature maps are activated, and features become tied to specific locations. This somehow “glues” the features into an object, which can then be perceived. In brief, attention is a beam (Eriksen & Hoffman, 1972) that serially scans a scene and constructs objects from raw features.

The theory rests primarily on three pieces of evidence (Treisman & Gelade, 1980). First, the search for features is parallel, while search for conjunctions is serial. Conjunction searches require the observer to move attention from location to location in order to conjoin features and examine objects. Second, identification can be better than

localization. Features are represented independently of their location, so observers can identify a feature (from pooled activity?) without knowing where it is. Lastly, features can miscombine into illusory conjunctions. Features are represented independently of their locations and can be combined only by attention. It is possible, if attention is "overloaded," to mistakenly construct an object from features at different locations. Moreover, spatial separation of the features does not affect the probability of miscombination (Treisman & Schmidt, 1982), which further shows that there is absolutely no topographic information in feature maps.

There are several variants to the original integration theory. One is "guided search" (Hoffman, 1978; Pashler, 1987; Treisman & Sato, 1990; Wolfe, Cave, & Franzel, 1989; see also Krose & Julesz, 1989; Zohary & Hochstein, 1989), which preserves most of the original theory but adds a new twist: the output of preattentive parallel search can direct attentive serial search. Treisman and Sato (1990), for example, suggest that observers can use attention to inhibit entire feature maps. Acting "upward" through the links to the master location map, inhibition reduces the signal arising from particular places/objects in the visual field. Serial search is more efficient, because the low amplitude signals can be ignored. The Wolfe et al. (1989) version is similar, except that preattentive search heightens the activity of likely targets rather than reduces the activity of unlikely ones. It is unclear whether map suppression and spotlight attentional mechanisms are the same or different mechanisms.

The major alternative to feature-integration theory has no specific name, so I will call it "interrupt theory." It arises from studies (Atkinson & Braddick, 1989; Johnston & Pashler, 1991; Nordruft, 1985; Sagi & Julesz, 1985) showing that observers can localize targets without identifying them. Sagi and Julesz suggest that in preattentive search, observers detect targets by a "difference signal" arising from a discontinuity in a feature gradient—the distribution of features across space. The featureless signal is an interrupt that automatically attracts focal attention to its location. The source of the signal can be recognized only after the observer moves attention to the interrupt's location.

This theory is also based on three pieces of evidence. First, there is the superiority of localization over identification. Second, observers can localize targets in preattentive search (Sagi & Julesz, 1985). Third, search performance may improve with increased numbers of distractors (Green, in press; Sagi, 1990; Sagi & Julesz, 1987). This could be interpreted to mean that observers make comparisons between neighboring items. The larger the display, the closer the items and the faster the comparison. All three show that location is accessible in parallel search.

The Sagi and Julesz (1985) theory (see also Nothdurft, 1985) differs from feature integration in several ways. Observers detect texture/feature discontinuities in global gradients, the spatial distribution of feature/texture ele-

ments, rather than pooled activity from feature maps. In addition, Sagi and Julesz (1985) believe that feature modules are topographic. The discontinuities arise from comparison of neighboring display elements. In contrast, feature-integration theory postulates nontopographic feature maps.

Johnston and Pashler (1990) suggest a compromise theory. Like Sagi and Julesz (1985), they have concluded that observers could sometimes localize but not identify targets. However, Johnston and Pashler (1990) used nonuniform backgrounds, so that detection could not arise from a simple gradient discontinuity. Instead, observers presumably detected an interrupt signal created by the pop-out feature. However, the interrupt does not identify the source feature. As in the Sagi and Julesz (1985) theory, observers must direct attention toward the interrupt location to identify the source.

The two theories resemble feature-integration theory in maintaining the dichotomy between attentive and preattentive processing. Sagi and Julesz (1985) are vague about the precise role played by attention. However, Johnston and Pashler (1990) take an additional step away from feature-integration theory by saying that attention does not affect feature computations. Instead, attention merely transfers feature identity to "central processes."

To summarize, physiological, computational, and psychological research all lead to the notion that early vision represents images in feature modules. The rationales may be different, but the message is clear: the visual system operates with a divide-and-conquer strategy. These questions remain: "What is the system architecture?" And "what puts Humpty Dumpty together again?"

To answer these questions, in the following section I will examine three candidate architectures for feature modules. After proposing the possible architectures, I will use the psychological evidence, along with physiological and computational data, to evaluate the alternatives.

Feature-Module Architectures

Feature modules are powerful because the divide-and-conquer strategy can handle the high dimensionality of visual information. While divide-and-conquer strategies reduce computational complexity, however, they introduce a new problem; the computations performed by different feature modules must be combined into a consistent global solution. The system architecture constrains the methods for efficiently integrating local computations.

Possible architectures can be divided into three classes. In the first (Figure 3), each module operates independently and in parallel. However, the integration of the modules occurs at a central locus and is controlled by a sequential executive processor. This will be termed a *blackboard*, after an analogous architecture used in speech understanding (Erman, Hayes-Roth, Lesser, & Reddy, 1980).

The basic organization is a collection of independent "knowledge sources" that may interact only by "writing" on a common blackboard and whose interaction is

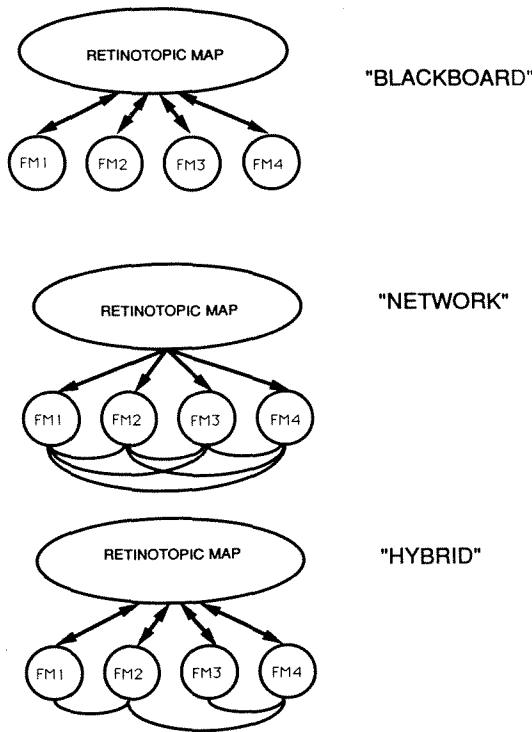


Figure 3. Feature-module (FM) architectures. Each architecture contains feature modules projecting from the retinotopic map. In blackboards, feature modules communicate only through a central map. In networks, feature modules may communicate directly. Hybrid blackboard/network architectures are also possible.

controlled by a single “scheduler.” Blackboards have many more complexities, so the analogy should not be taken too literally. For the present purposes, it merely is important to note that blackboards have a processing bottleneck, due both to the necessity to perform all operations on the central blackboard and to the control of all processing by a single control operator.

The major search theories assume a blackboard architecture, and feature-integration theory is probably the most clear-cut example. Features reside in independent modules and can be integrated only when focal attention, acting as the executive, accesses the master location map, the blackboard. Although their theories are less specific, Sagi and Julesz (1985) and Atkinson and Braddick (1989) also apparently subscribe to a blackboard architecture, whereas Johnston and Pashler (1990) and Wolfe et al. (1989) are more difficult to classify. (See also Navon, 1990a, for more blackboard possibilities.)

There are many possible variations on this architecture. For example, Figure 3 shows the retinotopic map acting as the blackboard. Some theories (Treisman & Gormican, 1988) suggest that the retinotopic map projects to an additional featureless module that records only topographic information (see Mishkin, Ungerleider, & Mako, 1983; Ungerleider & Mishkin, 1982). Feature modules would then project to this specialized location map rather

than back to the retinotopic map. However, these architectures would be computationally equivalent.

The second architecture class is the network (Figure 3), in which feature modules are directly connected together and control is distributed. Features are detected and integrated throughout the visual field in a single, parallel operation. Network architectures are common in the computational literature and are sometimes used in physiological theories. However, no major search theories are explicitly based on a network.

There are numerous variations on this architecture too. For example, Poggio, Gamble, and Little (1988) offer a more complicated architecture, in which each module projects to a separate “discontinuity map.” To combine features, the system uses constraints from the discontinuity maps rather than directly from the feature modules.

The following qualitative account will show that a network produces a very different view of feature detection and integration. To start, consider segmentation within a single feature module (Figure 4A), consisting of a topographic array of units. Although not shown, it is assumed that the units have overlapping receptive fields, creating some positional uncertainty. Segmentation proceeds by parallel, cooperative processing among neighboring units. One possible computation would be relaxation labeling (Barnard & Thompson, 1980; Zucker, 1976), a class of

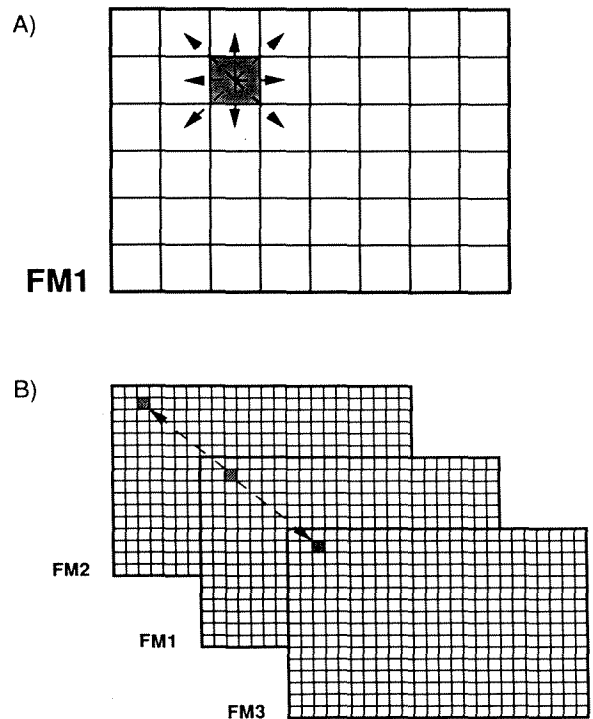


Figure 4. Schematic representation of relaxation labelling within and across feature modules. Each module contains feature detectors in two-dimensional array. Detectors are presumed to have overlapping receptive fields (not shown). Therefore, position can be coded only by combined output from several detectors.

cooperative computations that uses the smoothness constraint to find contours. The basic assumption is that neighboring units should be sensing similar features.

In relaxation, each unit "queries" its neighbors to see whether they signal a consistent segmentation. For example, if a unit were looking for luminance orientation—say, vertical contours—the neighbors directly above and below it should also be signaling vertical, not horizontal, diagonal, or no contour. Finding that the top and bottom neighbors sense vertical would increase the strength of the vertical signal coming from the unit in the middle. Conversely, conflicting information from the neighbors would decrease the strength of the vertical signal.

The definition of *neighbor* is variable. It could simply mean the immediate neighbor (as shown in Figure 4A), or it could extend out to cover a large area of the module. Whatever the size, consistent neighbors increase the unit's confidence in the vertical contour, whereas inconsistent neighbors reduce it. The neighbors themselves check the consistency of their neighbors. This propagates the influence of each unit beyond its own neighborhood. Each unit is influenced directly by units in its neighborhood and also indirectly by distant units to which it has no direct connection.

The unit continuously assesses consistency of the vertical hypothesis (as well as other possible orientation hypotheses) with its neighbors' activities and adjusts the probability that it is indeed sensing a vertical discontinuity. After many iterations, the units would presumably converge on a unique solution for both contour orientation and precise location.

Relaxation computations proceed concurrently within and across modules. Units query neighbors at the same locations in other modules (Figure 4B) as well as neighbors within the same module. That is, while performing its own computations, the luminance module may be querying the binocular disparity module to see if it also is sensing a vertical discontinuity at the same location. If so, the module segmentations are consistent and mutually reinforce each other. If not, there may be multistability or some other ambiguity. Another solution is to have some modules exert stronger constraints than others and to resolve inconsistencies in favor of the stronger module. For example, there is evidence that motion and binocular disparity exert strong constraints on color (see below).

The within- and across-module computations occur in parallel. Features are detected and integrated in a single procedure. In fact, consistency across modules is used to ascertain that a particular feature is present in a given module. Suppose there is uncertainty in the luminance module as to whether a contour exists at a particular location. If there is strong evidence in the binocular disparity module for a contour at that location, the binocular disparity would reinforce the vertical contour in the luminance module. This does not happen in blackboards, because feature detection and integration are two distinct and sequential operations.

However, the connections within a module are shorter than those across modules (Cowey, 1979), so it seems likely that within-module computations would proceed faster. Interruption early in processing, such as by a mask, may cause the local segmentations to become partially completed but not firmly coupled together (see the discussion of illusory conjunctions below).

This is not meant to be a detailed model. There are many variations on the relaxation theme (see, e.g., Ballard, 1984; Barnard & Thompson, 1980; Barrow & Tenenbaum, 1978). It merely describes how a distributed network might handle feature detection/integration and underscores the differences between the operations in networks and blackboards. The architectures do not differ in the assumption of parallel feature processing. Both assume that retinal information is divided into modules. They do not necessarily differ on the features that are represented in different modules. Although it may be important for detailed theories, discussion of architecture does not require specification of the primitives represented in FM1, and so forth, in Figure 3.

The major difference between architectures consists in how features may be efficiently combined—are feature modules independent or interconnected? If feature modules directly connect, there need be no common blackboard or central executive process, such as focal attention, to glue features together. In blackboard architectures, and in the theories they engender, feature detection and feature combination are serial stages. The first is preattentive, whereas the second is governed by attentive control. In network architectures, the visual system must still detect and combine features, but the two tasks are performed in parallel, without central control. In fact, the two operations can be mutually supportive. There is no need to postulate a dichotomy between focal and preattentive processing in feature integration.

A network architecture does not necessarily imply that attentional mechanisms are fiction. However, as an architecture for visual search, networks suggest that focal attention plays no major role in feature integration. I will return to possible roles for attention in my final remarks.

Finally, the third possible architecture is a network/blackboard hybrid (Figure 3). This is a blackboard with direct connections between subsets of modules. After all, there is no reason to assume that all feature interactions are handled identically. To highlight key issues, however, in the remaining discussion I will continue to contrast networks and blackboards.

Evaluation of Blackboard and Network Architectures

Because current theories generally assume a blackboard architecture, there is relatively little behavioral evidence about networks. The initial discussion will therefore center on arguments against blackboards. In particular, I will present arguments against feature-integration theory, simply because it is the most extreme blackboard theory and because it has led to a large body of relevant experimen-

tal research. In the first three sections, I will discuss the three key pieces of evidence for feature-integration theory; a subsequent section will contrast predictions of blackboards and networks. In the final sections, I will examine other psychophysical and physiological evidence in order to evaluate architectures.

1. *Detection, localization, and identification occur in parallel.* Feature-integration theory says that features are stored separately from their location codes. This aspect of the theory is supported by data showing that observers can detect targets without localizing them (Treisman & Gelade, 1980). However, this finding has not been replicated. Moreover, Johnston and Pashler (1991) point out possible artifacts in Treisman and Gelade (1980) and in similar studies. For example, many studies suffer from the "negative information problem." If there are two possible targets, and if they are not equally detectable, observers can perform above chance with a simple strategy; when no target is visible, they can choose the less detectable alternative.

In my own study (Green, in press), I compared detection, localization, and identification and found no difference in performance. Furthermore, a review of the literature showed no clear evidence that localization and identification are performed at different levels of accuracy. Roughly speaking, if the observer can detect the target, he can localize and identify it. This finding contradicts both feature-integration and interrupt theories, because it argues that the three tasks do not represent distinct processing stages. However, it is consistent with a network in which feature detection and object synthesis occur in a single parallel operation. Green (in press) outlines the argument in more detail.

2. *Conjunctions can be detected in parallel.* Blackboard-based theories require a central blackboard and a central control process to integrate feature modules. In feature-integration theory, the master map is the blackboard, and attention serves as the control process for gluing features together. The primary supporting evidence is that response time in conjunction search grows linearly as the number of display items increases. Presumably, the observer must perform an item-by-item serial scan. To draw this inference, feature-integration theory relies on three axioms: (1) search curves are diagnostic of serial versus parallel processing; (2) feature searches are always parallel and conjunction searches are always serial; and (3) attention is needed to glue features together. The evidence does not clearly support any of these axioms.

To start, there is doubt that the slopes of search curves have much ability to reveal details about the underlying process. Townsend (1972) noted that the slopes of curves do not necessarily distinguish between serial and parallel processes. There are various types of limited-resource parallel systems that behave much like serial systems. For example, relaxation processing is parallel, but depending on the circumstances, convergence time could grow with display size.

Another problem is that linear curves could be predicted by many different theories. Sagi (1988) offers a different

explanation for the positive slopes found in conjunction search. Serial search occurs because each feature map is confused by irrelevant feature gradients (see item 4 below). After removing the irrelevant gradients, Sagi demonstrated that conjunctions of spatial frequency and orientation, which presumably required serial search (Walters, Biederman, & Weisstein, 1983), would instead pop out.

Navon (1990a) argues that conjunction search would produce linear curves even if features were preattentively conjoined. In fact, he argues that because of the computational considerations, feature-integration theory really predicts positively accelerating, not linear, curves. The resulting debate (Navon, 1990b; Treisman, 1990) over semantics underscores the difficulty in making strong inferences from the slopes of search functions.

Empirical evidence further blurs the distinction between serial and parallel search. Whereas initial evidence (e.g., Treisman & Gelade, 1980) suggested that observers must scan each item serially to detect conjunctions, more recent studies often report exceptions. Several studies (e.g., Egeth, Virzi, & Garbart, 1984; Mordkoff, Yantis, & Egeth, 1990; Wolfe et al., 1989; Zohary & Hochstein, 1989) suggest that observers need not examine every item in serial search. One explanation is that search can be restricted to a subset of items by ignoring potential targets containing a particular feature. This has given rise to "guided search" theories (e.g., Treisman & Sato, 1990), which suggest that observers use the information gained from an initial parallel search to make the subsequent serial search more efficient. Such theories, however, are still blackboard theories, because attention is the control mechanism that glues features together.

Other studies show that observers normally detect many conjunctions, such as color or shape with binocular disparity (Nakayama & Silverman, 1986; Steinman, 1987) and motion (McLeod, Driver, & Crisp, 1988), independently of display size. In fact, parallel conjunction searches may be the rule rather than the exception. Several studies (Moraglia, 1989a; Steinman, 1987; Wolfe et al., 1989) have reported that large target-distractor value differences and/or practiced observers produce shallow or even flat slopes for most conjunctions.

Interpretation is further clouded by the similarity of conjunction search to the search for feature targets that are similar to the background. In fact, Moraglia (1989a; see also Dehaene, 1989) demonstrated that the distinction between serial and parallel search correlates poorly with the distinction between feature and conjunction search. He measured search time for a conjunction of spatial frequency and orientation and for each feature individually. When target-distractor differences were large, both features and conjunctions were detected in parallel. If target-distractor differences were small, search was serial for both features and conjunctions. Serial search does not correlate with the need to glue features together.

Treisman and Sato (1990) try to explain why there is such a poor correlation between serial search and conjunction targets. They note that the pop-out conjunctions

typically contain features—binocular disparity and motion—that cue depth. Perhaps observers divide items into different depth planes, so the conjunction search effectively becomes a feature search at two spatial locations. Although this strategy is reasonable, there is no evidence that observers actually employ it. Moreover, this notion does not explain Moraglia's (1989a), Sagi's (1988) or many of Wolfe et al.'s (1989) results.

As noted above, Treisman and Sato (1990) also postulate a new attentional mechanism—map suppression. Conjunctions can be detected easily (in parallel or with a shallow slope) if attention can suppress one feature map. The notion that observers can selectively attend to visual features is not new and has been used in many theories (e.g., Garner, 1974). However, other studies (Posner, 1978; see also item 4 below) have shown that observers often cannot selectively attend one feature to the exclusion of another.

This modification, even if true, substantially weakens feature-integration theory. First, it adds new baggage and greatly reduces the theory's economy, which is supposedly a major virtue (Treisman & Sato, 1990). More importantly, much of feature-integration theory is based on the distinction between flat and sloped search functions because they correspond to parallel/feature versus serial/conjunction search. The revised theory implies that search curves do not reveal much about processing. There is no sharp distinction between flat and sloped curves, and the slopes may be due to different combinations of attentional strategies.

3. *Attention is not necessary to combine features.* To summarize the argument so far, blackboard-based theories, such as feature-integration theory, rely largely on three pieces of empirical data: superiority of identification over localization, the serial functions of conjunction search, and the existence of illusory conjunctions. The superiority of identification has not been replicated and may have been due to artifacts. Observers perform many conjunction searches in parallel, and in any event, the feature-integration explanation is in doubt. This leaves only illusory conjunctions to support feature-integration theory.

Assuming that illusory conjunctions exist, two important questions arise: (1) Do illusory conjunctions occur more often outside the focus of attention? And (2) if so, does this prove that attention is the perceptual glue that connects feature modules together?

With regard to the first question, several authors (Briand & Klein, 1987; Cohen & Ivry, 1989; Prinzmetal, Presti, & Posner, 1986; Treisman & Schmidt, 1982) have concluded that attention does integrate features. However, the interpretation of these studies is unclear. Illusory conjunctions are measured indirectly, so the experimenter must make many assumptions: the locus and area of focal attention must be known, feature errors must act as an adequate baseline, and so forth.

Cohen and Ivry's (1989) study exemplifies the difficulty of interpreting illusory conjunction studies. Cohen and

Ivry claim that outside of attention, illusory conjunctions occur only between neighboring features with less than 1° separation. Inside the attentional focus, illusory conjunctions can occur between any two features. Cohen and Ivry used the ratio of feature errors to conjunction errors to measure the illusory conjunction rate. Observers detected letters ("X" or "F") in one of four colors, with an "O" of a different color as a distractor. When the observers correctly reported the letter and not the color, there were three possible color errors. Feature errors occurred when the guessed color was not in the display, whereas conjunction errors occurred when observers guessed the color of the distractor. With three possible colors, pure guessing would predict a 1:2 ratio of color feature errors to conjunction errors. A lower ratio would show that observers guessed the distractor color at more than chance, and it would suggest that illusory conjunctions must have been occurring.

Unexpectedly, some conditions produced "exclusionary conjunction" errors, where the ratio of feature to conjunction errors was above chance. There is no satisfactory explanation for this result. Moreover, the authors never explain the inverse relation between overall accuracy and the feature-conjunction error ratio. Lastly, this study may contain the "negative information problem" (Johnston & Pashler, 1990).

There is another reason for doubting Cohen and Ivry's (1989) interpretation of their experiment. As in all other illusory conjunction experiments, Cohen and Ivry had to make the implicit assumption that features can be identified but not localized. As already noted above, only Treisman and Gelade (1980) obtained this result. In other studies (reviewed in Green, in press), localization has been as good or better than identification.

Tsal (1989a, 1989b) has pointed out further flaws in the design and interpretation of the experiments (e.g., Briand & Klein, 1987; Treisman & Schmidt, 1982) that purportedly confirm Treisman's account of illusory conjunctions. In fact, he notes that one (Prinzmetal, et al., 1986) can easily be interpreted as providing counterevidence. Tsal (1989a, 1989b) examined feature (wrong feature) and conjunction (right features but wrong combination) errors both inside and outside of attention. Feature-integration theory predicts that illusory conjunctions should be more frequent outside of focal attention. Although there were somewhat more conjunction errors outside of attention (35% vs. 25%), the number of conjunction errors both inside and outside of attention was much greater than the number of feature errors. What caused all those conjunction errors inside of focal attention and why were they so much more frequent than feature errors? If attention is required to conjoin features and make them perceptible, how could one even measure conjunction errors outside of attention?

Authors tangled in these issues (see the lively debate by Briand & Klein, 1989; Tsal 1989a, 1989b) have all been plagued by the same problem: How can one really know if a stimulus is being attended? The whole notion

of inside versus outside of attention depends on the spotlight metaphor and the assumption that there is a sharp dichotomy between attending and not attending. Suppose there were a single, graded process, so that it might be possible to "attend" stimuli with varying degrees of resources, or suppose there were no beam. In fact, data (Gathercole & Broadbent, 1987) do show that the effects of distractors depend not on whether or not they are inside a beam, but on their distance from the target. It simply takes more time for the influence of distant distractors to reach the target. This could explain many phenomena previously attributed to an attentional spotlight.

The second major question is whether attentional effects, even if they could be demonstrated, would prove that attention is the perceptual glue. To merely say that illusory conjunctions occur because features are mislocalized is to say very little. Poor registration of location is not necessarily antagonistic to network architectures.

In fact, illusory conjunctions would be expected to occur simply because there is uncertainty inherent in localization. Many studies (e.g., Hiltz, Rentschler, & Brettel, 1981; Morrone, Burr, & Spinelli, 1989) have demonstrated that, due in part to coarser receptor sampling, cortical magnification, and misfocus, humans are poor at localizing single features in the peripheral field. Since color is used as a feature in most experiments, it is particularly noteworthy that color localization is poor throughout the visual field (Troscianko & Harris, 1986).

Mislocalization is further encouraged by the combination of eye movements with the different processing rates for different features. For example, low spatial frequencies are processed faster than high spatial frequencies (Breitmeyer & Ganz, 1976; Green, 1984), and they form faster than color (e.g., Schwartz & Loop, 1982). If the display is presented during an eye movement, the different processing times may cause features to be coded at different retinotopic locations. Brief exposure may reduce this artifact but cannot eliminate it. True, it takes about 200–250 msec (Hallet, 1986) to initiate an eye movement, but the eyes could already be in motion before flash onset. Most studies have been done with unpracticed observers who had little experience in holding their eyes still during fixation.

In short, any mechanism responsible for combining features represented in separate modules is going to make errors, especially in the periphery and especially with color. Since observers typically foveate the location being attended to, it is not surprising that there are more illusory conjunctions outside of attention (i.e., in the periphery). As Prinzmetal and Keysar (1989) correctly note, it would be surprising if illusory conjunctions did not occur.

Tsal (1989a) has also argued that even if attention were related to illusory conjunctions, this does not necessarily mean that attention glues features together. Attention facilitates many tasks, such as simple luminance detection. Suppose that features are integrated by some independent process. Might not attention facilitate this

process as well? If so, attention can no more be the mechanism of feature-integration than it can be the mechanism of luminance detection.

Attention could affect feature combination in many ways. Tsal (1989a) suggests that memory loss may cause illusory conjunctions. Perhaps attention can prevent the memory decay. This is supported by the observation (Green, in press) that when subjects must perform two concurrent tasks, they perform slightly worse on the task measured by the second response. Wolfe et al. (1989) seem to argue that attention detects activity levels at different spatial locations (see below). Prinzmetal and Keysar (1989) assume that initial feature registration is coarsely localized. To prevent illusory conjunctions, the feature must be more accurately pinned to position. Any factor—and lack of attention is only one of many—that prevents finer localization will therefore encourage illusory conjunctions.

There are other arguments against the attention-as-glue theory. Illusory conjunctions, even under the most favorable conditions, are relatively uncommon. Features generally manage to become integrated, with or without attention. If I try as hard as possible, I still cannot create illusory conjunctions by narrowing my attentional focus. Tsal (1989a) further notes that if attention were required for one to integrate features, the attentional spotlight would have to be narrowed to fit each object exactly and perfectly. This, Tsal argues, seems untenable.

Regardless of the illusory conjunction data, and their explanation, it is clear that focal attention is unnecessary to combine features. Houck and Hoffman (1986) created orientation-specific color aftereffects in one part of the display while the observer attended another. Since the aftereffect requires conjunctive processing of oriented luminance edges and color (Savoy, 1987), these features must have been joined outside of focal attention.

Treisman and Gormican (1988) attempt to dismiss Houck and Hoffman (1986) by saying that the aftereffects were created at an earlier stage of vision where features are still conjoined. Such explanations are hard to evaluate, because they are not based on any real evidence. (In fact, the source of form-specific color aftereffects is poorly understood.) However, this argument can be partially refuted by noting that color-coded cells, even in early vision, typically have little orientation tuning (see, e.g., Hubel & Livingstone, 1987). Similarly, color aftereffects can be made contingent on direction of motion (Stromeyer & Mansfield, 1970), although color and motion are already presumably segregated prior to the lateral geniculate nucleus (LGN).

The binocular rivalry paradigm provides more evidence that features combine without attention. In many studies (e.g., Blake & Fox, 1974), it has been possible to create aftereffects specific to particular feature combinations even when the stimuli were presented to an eye suppressed by binocular rivalry. The observer does not consciously perceive the adapting stimuli, so it is hard to believe that the observer would be applying focal attention. Rivalry exper-

iments are especially convincing, because they require no assumptions about spatial distribution of attention.

Can networks explain illusory conjunctions? Recall that cooperative processing within a feature module should occur faster than across modules. This would allow the feature segmentations to begin forming before localization was precise and before different modules were firmly coupled. Illusory conjunctions could occur between spatially adjacent features, although not between distant ones (Cohen & Ivry, 1989).

4. *Variation in the "irrelevant" dimension degrades search and segregation.* There are some clear differences in predictions made by networks and blackboards. Network architectures suggest that feature modules are directly connected to provide mutual constraint. If the segmentations from different modules are inconsistent, it should be more difficult for each module to compute a solution. If the modules are independent, as in blackboards, the computation in one module should have no effect on the others during preattentive vision.

Data show that features interfere with one another. In most studies, the target and distractors are identical except for a single feature. Several experimenters have asked whether orthogonal variation in an "irrelevant" (for the detection task) dimension degrades performance. If feature maps functioned independently, the orthogonal variations should have no effect. Results (Callaghan, 1984, 1989; Callaghan, Lasaga, & Garner, 1986; Pashler, 1988) often show, however, significant loss of performance when an irrelevant feature is randomly varied.

5. *Features sum to produce better detection.* Just as inconsistent segmentations from different feature modules degrade performance, consistent segmentations would support each other and enhance performance. Several experiments have also demonstrated this effect.

Studies show that observers can more easily detect features that are conjoined rather than separate. Caelli and Moraglia (1985) found that dimensions could sum in preattentive texture segregation. Observers viewed textures composed of Gabor signals that varied in spatial frequency and/or orientation from the distractors. Speed of segregation depended on the difference in feature value between target and background on each dimension; small differences produced poor segregation, whereas larger differences produced faster responses. Most importantly, however, simultaneous small variations of orientation and spatial frequency produced the same effect as did a large difference in either dimension alone. Two features, neither of which popped out by itself, could sum to produce rapid detection.

Farell (1984) reported a similar result using a different paradigm. He tested for detection of two features in two separate display items or the same two features conjoined in the same item. Observers detected conjoined features (two features, same location) more readily than they did the same features in unspecified combinations (two features, different locations). If features are independent, de-

tection of conjoined features and separate features should be equivalent. On the contrary, direct connection of feature modules predicts better detection of features at the same spatial location.

Wolfe et al. (1989) compared detection of "triple conjunctions" (three features) and "simple conjunctions" (two features). The triple conjunctions were easier to detect and often produced flat functions of response time versus display size (see also Quinlan & Humphreys, 1987). Wolfe et al. (1989) suggest that each preattentive map creates its own segmentation, which marks possible target locations. These maps are apparently topographic and superimposed, so that the activity from each map sums location-by-location with the others. If activity at the target position is sufficiently high, and if noise in the system is sufficiently low, attention is immediately directed to the correct location, and search time is independent of display size. However, low signal and high noise may cause a few misdirections of attention, resulting in a function with a small positive slope.

If the preattentive/attentive dichotomy is removed, this model agrees well with a network architecture that uses simple summation rather than relaxation. Each module segments the image and combines its evidence with other maps. The more modules that constrain the segmentation, the easier the search. Focal attention is excess baggage in this model, since it plays no significant role. It does not explicitly glue maps together; feature activity apparently sums simply because the maps are topographic and are superimposed on one another. Focal attention appears to act merely as a mechanism to detect the summed activity. It is not clear why attention must be moved to the target to accomplish this. The only real purpose served by attention is that misdirections provide an explanation for the occasional positive slope. However, Wolfe, et al. (1989) offer no proof that this is the correct explanation. It would be simpler to dispense with focal attention altogether and just assume, as many have done (e.g., Green, Terman, & Terman, 1979), that response time lengthens with smaller sign-noise differences.

Lastly, features sum to produce improved localization as well as detection. Rivest and Cavanagh (1991) found that observers were better at locating edges created by combinations of color and luminance than by either feature alone. Again, this would be predicted from mutually constraining, topographic feature modules.

6. *Feature modules are topographic.* Feature-integration theory says that feature maps contain no usable topographic information. Evidence showing that feature modules are topographic serves as indirect evidence favoring network architectures. Recall that the original motivation for blackboard theories was apparent independence of codes for features and their spatial position. Although there is no principled reason why blackboard architectures could not have topographic feature modules, feature-integration theories (Briand & Klein, 1987; Treisman & Gelade, 1980; Treisman & Gormican, 1988) assume that

features float freely in their map and that locations are stored in a separate map of locations. The two can only be connected by focal attention.

On the other hand, a network must have topographic information in each feature module because there is no separate location map. The goal of a network is to integrate local computations into a consistent global solution. A likely definition of a consistent solution is that it parses visual space into a unique set of segments. To know whether the local solutions are consistent or inconsistent, each module would need to know the location of its own features as well as the location of features in other modules. This is not necessary in a blackboard, because the features are presumably tied to a master location map.

Many experiments suggest that feature modules are topographic. In several studies (Green, *in press*; Sagi, 1990; Sagi & Julesz, 1985, 1987), search has become more accurate with increased numbers of distractors. Sagi and Julesz (1987) have interpreted this to mean that observers perform preattentive search by making local comparisons in a feature gradient. As the display size increases, the items become closer together and the comparison easier.

This conclusion is reinforced by several other studies. Moraglia (1989b) showed that search is much easier if the distractors are arranged to create a structured background (see also Gorea & Julesz, 1990). Gathercole and Broadbent (1987) found that the effects of distractors depended on their distance from the target. Lastly, Fahle (1990) examined preattentive search for lines with vernier offsets amid straight line distractors. The vernier targets popped out, even though the only cue was a very small relative shift in position of two line segments.

Lastly, Cohen and Ivry (1989), taken at face value, reported that illusory conjunctions occur only between nearby features. This result refutes the assertion (Treisman & Schmidt, 1982) that illusory conjunctions occur independently of feature proximity and therefore the notion of pooled activity. It also shows that coarse location information must be available in early vision.

As Sagi and Julesz (1985) note, search results suggest that local differences are important and that in search, local constraints are more important than global ones (see final remarks). In the kind of network model that is presented above, processing arises strictly from local computations within a neighborhood of detectors. Close proximity would therefore help speed the computation.

7. Other psychophysical experiments show that features are linked. A general theory of early vision should be applicable, not only to search, but to other behavioral paradigms. I have already mentioned the Houck and Hoffman (1986) experiment, which showed that features combine in orientation-contingent color aftereffects. Many other studies, although not directly concerned with the distinction between attentive and preattentive processing, have resulted in close coupling among features. The literature is too large to cover in depth, so the following is merely a sampler.

Several studies have yielded interactions between isoluminance (defined by color) and luminance (defined by brightness) contours. Switkes, Bradley, and DeValois (1988; see also DeValois & Switkes, 1983) found that isoluminant gratings can mask luminance gratings but that luminance gratings can facilitate isoluminant ones. In other studies, luminance edges have interacted with isoluminant edges to aid in color identification (e.g., Eskew, 1989).

Luminance and color also interact in motion perception. For example, isoluminant gratings appear to move much slower than luminance gratings (Cavanagh, Tyler, & Favreau, 1984; Troscianko & Fahle, 1986). Most moving objects have both color and luminance contrast, so the luminance edges should leave color lagging behind. Since this does not happen, there must be a constraint keeping form and color joined.

Others (Kleffner & Ramachandran, 1988; Ramachandran, 1987) have found that a moving luminance pattern "captures" an isoluminant pattern and makes it appear to move. Ramachandran (1987) suggests that the constraint is unidirectional, from motion to color. Green and Odom (1991) also found that luminance gratings capture isoluminant ones even when separated by large distances. However, there is also evidence of a mutual constraint. Apparent velocity of colored gratings lies between the apparent velocity of the isoluminant and luminance components alone (Cavanagh et al., 1984), and Green (1991) has found that moving isoluminant patterns can capture black dots. Kleffner and Ramachandran (1988) have similarly reported a capture of color by binocular disparity.

Interaction is also found in motion aftereffects and stroboscopic motion. Adaptation to moving luminance gratings produces motion aftereffects in isoluminant gratings (Cavanagh & Favreau, 1985). Green (1989) found interaction in correspondence matching during stroboscopic motion. Observers viewed a series of frames containing red and green dots on an isoluminant yellow background. Dots in successive frames moved to neighbors of the same color, producing coherent motion. Addition of luminance contrast with ambiguous matches destroyed the coherent motion of the colored dots. The color failed to uncouple from luminance and move in its own path.

Feature interactions are not restricted to color. Evidence (see, e.g., Breitmeyer & Ganz, 1976) suggests that motion (the transient/magnocellular stream) can inhibit detailed form (the sustained/parvocellular stream) and possibly vice versa. Moreover, many adaptation, masking, subthreshold summation, detection-discrimination, and figural aftereffects data (reviewed by numerous authors; see, e.g., Braddick, Campbell, & Atkinson, 1978) show detectors selective to combinations of spatial frequency, orientation, direction of motion, binocular disparity, and so forth.

All of these studies have shown feature linking in early vision. It might be argued that the effects occurred either before features were segregated or after they had been glued together by focal attention. Treisman and Sato

(1990) made the first of these arguments with respect to Houck and Hoffman (1986), and I have already mentioned some counterarguments.

There is no obvious way to determine conclusively whether search operates on the same representations as adaptation, motion capture, and so forth. It is possible, however, to compare known properties of the representations used in different tasks. It is interesting that the bandwidths of orientation channels defined by many different tasks are similar. Using the standard half-height/half-width criterion (see Braddick et al., 1978), many orientation-specific adaptation studies (see Braddick for a review) have yielded orientation bandwidth, with values in the 12°–15° range. Similar bandwidths have been found for correspondence matching (Green, 1986) and motion capture (Green & Odom, 1991). Bergen and Julesz (1983) measured search performance with a range of orientation differences between target and mask. It is clear from their data that orientation bandwidth for search falls between 10° and 20°. Of course, the similarity could be explained away post hoc; but with no convincing argument, why not accept the simpler explanation—that all three tasks have the same representational substrate?

In addition, it is clear that the conjoining of features does not correlate with focal attention. The same pair of features can appear joined without focal attention and independent with focal attention. For example, color seems loosely tied to stereo depth (Lu & Fender, 1972), even though the observers inspect the display with focal attention. On the other hand, observers can perform parallel search for conjunctions of color and depth. This is exactly the reverse of the feature-integration theory predictions.

My conjecture is that color and binocular disparity are always directly linked but that the constraint flow is stronger in one direction; binocular disparity and luminance constrain color more than color constrains them. However, the asymmetrical constraint is apparently sufficient to support parallel search.

8. *Physiology (combined with behavioral data) refutes blackboards.* Many psychologists and computationalists use physiological data to reinforce their theories. This is not strictly necessary in the current discussion, because architecture is an abstract computational concept that could be implemented in software as well as hardware. Further, history shows that attempts to link physiological data to perceptual phenomena are at best tenuous and at worst misguided. Nevertheless, physiological data provide a major rationale for feature-module theories.

Unfortunately, the physiological evidence on feature modules is inconclusive. Most researchers believe that the primary visual pathway is organized in two or more streams from the retina to the LGN and then in three major streams from the primary visual cortex to extrastriate cortex and beyond. However, there is disagreement over the functional and anatomical integrity of the streams. One view (e.g., that of Hubel & Livingstone, 1987; Livingstone & Hubel, 1987, 1988) sees the streams as highly

distinct, with little overlap in cell properties or anatomical locus. Others (e.g., DeYoe & van Essen, 1988) believe that, although somewhat specialized, the streams are fuzzy, showing some common feature sensitivities.

If we assume for the moment that the streams are sufficiently distinct to be the physiological substrate of feature modules, there are three critical questions: (1) Are the feature streams topographic? (2) Is there evidence for a master map of locations? And (3) do the streams interconnect?

Feature-integration theory asserts that features are not tied to location. However, there is little evidence of non-topographic mapping in the visual system. The neurons of extrastriate cortex have larger receptive fields than do neurons of lower visual areas, but they are still localized. Bigger receptive fields may suggest coarser localization, but it is possible for the combined output of many broadly tuned mechanisms to give very precise information. Color vision is a perfect example.

Feature-integration theory puts topographic information in a master map of locations. Some (e.g., Treisman & Gormican, 1988) use physiological data (e.g., Mishkin et al., 1983; Ungerleider & Mishkin, 1982) to place the master map of locations in the parietal lobe. Investigations (e.g., Levine, Warach, & Farah, 1985; Newcombe & Russell, 1969) of patients with parietal lobe damage have yielded deficits in localization, but no loss in object recognition. This confirms a role for the parietal lobe in localization, but it provides no evidence that the parietal lobe integrates features into perceptible objects. An alternative scheme (Treisman & Gormican, 1988) places the master map in the primary visual cortex (see Figure 3). This seems an odd suggestion, because the neurons in striate cortex are tuned to multiple features—that is, they have some features already glued together. Why would attention be necessary for feature integration?

Perhaps the key question is whether streams interact. The physiological data are inconclusive, but it seems clear that the streams cannot be viewed as three pipelines running from the eye to higher visual centers. There are many cross- and reciprocal connections, with convergence and divergence of input/output. The extrastriate cortex receives direct LGN input as well as striate cortex input (Bullier & Girard, 1988). Although the significance of the LGN input is unclear, it may act as a “gating” mechanism for different areas to influence one another (Girard & Bullier, 1989). There are also reciprocal connections from higher levels back to lower ones. Zeki (1989) suggests that these may connect the different processing streams and allow integration. The streams also have lateral cross-connections like those suggested in Figure 3 (DeYoe & van Essen, 1988). If these streams are to be the substrate of feature modules, early vision more closely resembles a network than a blackboard.

It would perhaps be best to leave “neuromythology” at this point. However, many authors support perceptual theories with the casual invocation of physiology. There-

fore I cannot avoid the temptation to ask whether close examination actually reveals any parallels between physiological and behavioral data.

As noted above, there is no real evidence for a master map of locations. However, the most obvious issue is whether physiology predicts which conjunction searches will be easy and which will be difficult. If streams are feature modules, feature-integration theory makes the strong prediction that conjunctions will require serial search when the features come from different streams. There is also a weaker prediction that effortless search will occur if the conjunctive features come from the same stream. The second prediction is weaker, because it is possible to say that distinct feature modules exist within each stream.

The fuzzy stream model makes prediction difficult, so most authors use Livingstone and Hubel's (1987) distinct stream model. However, predictions based on the distinct stream model fail. For example, color-binocular disparity search is easy (Nakayama & Silverman, 1986), although color is processed by one stream (parvocellular-blob-thin) and binocular disparity is processed through a different stream (magnocellular-4B-thick/MT). Similarly, observers can perform parallel search on conjunctions of form and motion (McLeod et al., 1988), although these features are processed by separate streams.

On the other hand, it may be difficult to detect conjunctions of features processed within the same stream. Spatial frequency and orientation, although their conjunction is difficult to detect in standard displays (Walters et al., 1983), would both be processed in parvocellular-interblob-pale stream if the frequency were high or in the magnocellular stream if the frequency were low.

In sum, the physiological data are inconclusive but show some evidence of stream interaction. However, the distinct stream model, if accepted at face value, does not predict the difficulty of conjunction search.

Final Remarks

In this paper, I have used a wide range of evidence to evaluate architectures for visual search. Although no one piece of evidence is incontrovertible, the balance heavily favors network theories over blackboard theories such as feature integration. A network architecture places new constraints on theories of visual search: the feature modules must be directly connected, spatially local constraints must dominate, and focal attention must play no major role in gluing features together.

Experimental evidence challenges the three cornerstone observations of both feature-integration and guided search theories: (1) There is no evidence that identification can be performed without localization. (2) Conjunction searches are often parallel; there is no sharp distinction between flat and sloped curves; and search functions are difficult to interpret. And (3) illusory conjunctions may have many explanations; there is no strong evidence to support the feature-integration account. Moreover, other evidence shows direct linking between features in parallel

tasks. Lastly, the physiological evidence used to support feature-integration theory is, if anything, contradictory. There is no convincing evidence that features are free-floating or that attention glues them together.

Interrupt theory fares somewhat better. It acknowledges that feature modules are topographic, but, like feature-integration theory, interrupt theory still suggests that attention is needed for feature identification. Sagi and Julesz (1985) do not explain how attention permits identification, so the theory might possibly be made compatible with a network architecture. The Johnston and Pashler (1990) theory could easily be viewed as a network theory. It assumes that feature computations occur in parallel throughout the visual field and that attention does not direct processing. However, Johnston and Pashler still base their theory on the assumption that targets may be localized but not identified in parallel search. Other evidence (Green, in press) does not support this belief. Nevertheless, their theory seems to be in the right ballpark.

The feature modules described here are similar to what others (Ballard, 1984; Barrow & Tenenbaum, 1978) call intrinsic images—that is, topographic feature maps that can compute spatially local constraints through relaxation (Zucker, 1976) or some similar cooperative process. The computations within each feature module and the integration of feature modules should not be viewed as distinct stages, nor should a control process be assumed.

I have identified only the minimal architecture needed to account for visual search. A complete visual architecture would be much more complex than any of those suggested here. One of the major conclusions is that visual search is based on local processing. A network of intrinsic images is therefore sufficient to account for current search data.

The following question remains: *What is the role for attention?* If attention is not glue that combines raw features into objects, then what is it? Current data do not present a clear picture. Further, most psychological theories are very vague about attention's role in perception. For example, Johnston and Pashler (1990) suggest that attention "transfers identity to central processes." What does this mean? For that matter, exactly *how* does attention act as "perceptual glue"?

One recent idea (Moran & Desimone, 1985) suggests that attention is better conceived as a mask than as a beam. Attention does not promote processing in one region but rather inhibits it in another. This is consistent with our introspections that attention is used not so much to enhance processing of the attended input, but to block out extraneous input. In a cooperative processing model, this could be accomplished by limiting the size of the cooperative neighborhoods (Kienker, Sejnowski, Hinton, & Schumacher, 1986).

Green and Odom (1991) provide some support for the mask theory. They have examined how the visual system resolves ambiguity inherent in the aperture problem. A grating moving through a circular aperture has a family of possible motions. The ambiguity can be resolved by

surrounding the aperture with a second grating of unambiguous direction. The aperture grating links to the surround to form a single segment moving in the direction of the surround. However, observers can reduce or even eliminate the effect by tightly focusing attention on the center of the display. The surround grating does not disappear, but simply loosens its grip on the aperture grating. Attention seems to control segmentation, not by combining features into objects, but by controlling which objects link to form segments.

It is unlikely, however, that attention has a single explanation. Recent evidence shows that attention is not one, but several distinct mechanisms. For example, studies (e.g., Luck, Hillyard, Mangun, & Gazzaniga, 1989; Nakayama & Macken, 1990) suggest that there are separate transient and sustained attentional mechanisms. The transient mechanism is bottom-up and automatic, and it directs the gaze to an "interesting" location in space. The sustained mechanism is top-down and under voluntary control. There are several other proposed attentional mechanisms, including Treisman and Sato's (1990) feature-map inhibitor. The study of attention would also seem to benefit from the divide-and-conquer strategy.

REFERENCES

- ATKINSON, J., & BRADDICK, O. (1989). "Where" and "what" in visual search. *Perception*, **18**, 181-189.
- ATTNEAVE, F. (1974). Apparent motion and the what-where connection. *Psychologia*, **17**, 108-120.
- BALLARD, D. (1984). Parameter nets. *Artificial Intelligence*, **22**, 235-267.
- BALLARD, D. (1985). *Cortical connections and parallel processing: Structure and function* (Tech. Rep. No. TR133). New York: University of Rochester.
- BALLARD, D. (1986). Cortical connections and parallel processing: Structure and function. *Behavioral & Brain Sciences*, **9**, 67-120.
- BALLARD, D., HINTON, G., & SEJNOWSKI, T. (1983). Parallel visual computation. *Nature*, **306**, 21-26.
- BARLOW, H. (1981). The Ferrier lecture: Critical limiting factors in the design of the eye and visual cortex. *Proceedings of the Royal Society of London: Series B*, **212**, 1-34.
- BARLOW, H. (1986). Why have multiple cortical areas? *Vision Research*, **26**, 81-90.
- BARNARD, S., & THOMPSON, W. (1980). Binocular disparity analysis of images. *IDEE Transactions on Pattern Analysis & Machine Analysis*. **PAMI-2**, 333-340.
- BARROW, H., & TENENBAUM, J. (1978). Recovering intrinsic scene characteristics from images. In E. Riseman (Ed.), *Computer vision systems* (pp. 3-26). New York: Academic Press.
- BERGEN, J., & JULESZ, B. (1983). Rapid discrimination of visual patterns. *IEEE Transactions on Systems, Man & Cybernetics*, **SMC-13**, 857-863.
- BLAKE, R., & FOX, R. (1974). Adaptation to invisible gratings and the site of binocular suppression. *Nature*, **249**, 488-490.
- BRADDICK, O., CAMPBELL, F., & ATKINSON, J. (1978). Channels in vision: Basic aspects. In R. Held, H. Leibowitz, & H.-L. Teuber (Eds.), *Handbook of sensory physiology* (pp. 3-37). Berlin: Springer-Verlag.
- BREITMEYER, B., & GANZ, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression and information processing. *Psychological Review*, **83**, 1-36.
- BRIAND, K. A., & KLEIN, R. M. (1987). Is Posner's "beam" the same as Treisman's "glue"? On the relation between visual orienting and feature integration theory. *Journal of Experimental Psychology: Human Perception & Performance*, **13**, 228-241.
- BRIAND, K. A., & KLEIN, R. M. (1989). Has feature integration theory come unglued? A reply to Tsal. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 401-406.
- BULLIER, J., & GIRARD, P. (1988). Visual responses of neurons in area V2 and in the superior temporal sulcus of the macaque monkey during reversible inactivation of area V1. *Society of Neuroscience Abstracts*, **14**, 602.
- CAELLI, T., & MORAGLIA, G. (1985). On the detection of Gabor signals and discrimination of Gabor textures. *Vision Research*, **25**, 671-684.
- 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., & FAVREAU, O. (1985). Color and luminance share a common motion pathway. *Vision Research*, **25**, 1595-1601.
- CAVANAGH, P., TYLER, C., & FAVREAU, O. (1984). Perceived velocity of moving chromatic gratings. *Journal of the Optical Society of America A*, **1**, 1428-1438.
- COHEN, A., & IVRY, R. (1989). Illusory conjunctions inside and outside the focus of attention. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 650-663.
- COWEY, A. (1979). Cortical maps and visual perception: The Grindley memorial lecture. *Quarterly Journal of Experimental Psychology*, **31**, 1-17.
- DEHAENE, S. (1989). Discriminability and dimensionality effects in visual search for featural conjunctions: A functional pop-out. *Perception & Psychophysics*, **46**, 72-80.
- DEVALOIS, K., & SWITKES, E. (1983). Simultaneous masking interactions between chromatic and luminance gratings. *Journal of the Optical Society of America*, **73**, 11-18.
- DEYOE, E. A., & VAN ESSEN, D. C. (1988). Concurrent processing streams in monkey visual cortex. *Trends in Neuroscience*, **5**, 219-226.
- DURBIN, R., & MITCHISON, G. (1990). A dimension reduction framework for understanding cortical mappings. *Nature*, **343**, 645-647.
- EGETH, H., VIRZI, R., & GARBART, H. (1984). Searching for conjunctively defined targets. *Journal of Experimental Psychology: Human Perception & Performance*, **10**, 32-39.
- ERIKSEN, C. W., & HOFFMAN, J. E. (1972). Temporal and spatial characteristics of selective encoding from visual displays. *Perception & Psychophysics*, **12**, 201-204.
- ERMAN, L., HAYES-ROTH, F., LESSER, V., & REDDY, R. (1980). The Hearsay II speech understanding system: Integrating knowledge to resolve uncertainty. *Computing Survey*, **14**, 213-253.
- ESKEW, R. (1989). The gap effect revisited: Slow changes in chromatic sensitivity as affected by luminance and chromatic borders. *Vision Research*, **29**, 717-729.
- FARELL, B. (1984). Attention in the processing of complex visual displays: Detecting features and their combinations. *Journal of Experimental Psychology: Human Perception & Performance*, **10**, 40-64.
- FAHLE, M. (1990). *A new elementary feature of vision*. Unpublished manuscript.
- GARNER, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.
- GATHERCOLE, S., & BROADBENT, D. (1987). Spatial factors in visual attention: Some compensatory effects of location and time of arrival of targets. *Perception*, **16**, 433-443.
- GIRARD, P., & BULLIER, J. (1989). Visual activity in area V2 during reversible inactivation of area 17 in the macaque monkey. *Journal of Neurophysiology*, **62**, 1287-1302.
- GOREA, A., & JULESZ, J. (1990). Context superiority in a detection task with line element stimuli: A low-level effect. *Perception*, **19**, 5-16.
- GREEN, M. (1984). Masking by light and the sustained-transient dichotomy. *Perception & Psychophysics*, **35**, 519-535.
- GREEN, M. (1986). What determines correspondence strength in apparent motion? *Vision Research*, **26**, 599-608.

- GREEN, M. (1989). Color correspondence in apparent motion. *Perception & Psychophysics*, **45**, 15-20.
- GREEN, M. (1991). *Form captures color, color captures form*. Manuscript submitted for publication.
- GREEN, M. (in press). Visual search: Detection, identification and localization. *Perception*.
- GREEN, M., & ODOM, J. V. (1991). *Occlusion, segmentation and the aperture problem*. Manuscript submitted for publication.
- GREEN, M., TERMAN, M., & TERMAN, J. (1979). Comparison of yes-no and latency measures of auditory intensity discrimination. *Journal of the Experimental Analysis of Behavior*, **33**, 363-372.
- HALLET, P. E. (1986). Eye movements. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance*. New York: Wiley.
- HILZ, R., RENTSCHLER, I., & BRETTEL, H. (1981). Insensitivity of peripheral vision to spatial phase. *Experimental Brain Research*, **43**, 111-114.
- HOFFMAN, J. E. (1978). Search through a sequentially presented visual display. *Perception & Psychophysics*, **23**, 1-11.
- HOUCK, M., & HOFFMAN, J. (1986). Conjunction of color and form without attention: Evidence from an orientation-contingent color aftereffect. *Journal of Experimental Psychology: Human Perception & Performance*, **12**, 186-199.
- HUBEL, D., & LIVINGSTONE, M. (1985). Complex-unoriented cells in a subregion of primate area 18. *Nature*, **315**, 325-327.
- HUBEL, D., & LIVINGSTONE, M. (1987). Segregation of form, color and stereopsis in primate area 18. *Journal of Neuroscience*, **7**, 3378-3417.
- JOHNSTON, J. C., & PASHLER, H. (1990). Close binding of identity and location in visual feature perception. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 843-856.
- KIENKER, P., SEJNOWSKI, T., HINTON, G., & SCHUMACHER, L. (1986). Separating figure-from ground with a parallel network. *Perception*, **15**, 197-216.
- KLEFFNER, D., & RAMACHANDRAN, V. (1988). Interaction of motion, color and depth. *Investigative Ophthalmology & Visual Science* (Suppl.), **30**, 299.
- KROSE, B., & JULESZ, J. (1989). The control and speed of shifts of attention. *Vision Research*, **29**, 1607-1619.
- LEVINE, D., WARACH, J., & FARAH, M. (1985). Two visual systems in mental imagery: Dissociation of "what" and "where" in imagery disorders due to bilateral posterior cerebral lesions. *Neurology*, **35**, 1010-1018.
- LIVINGSTONE, M., & HUBEL, D. (1987). Psychophysical evidence for separate channels for the perception of form, color, motion and depth. *Journal of Neuroscience*, **7**, 3416-3468.
- LIVINGSTONE, M., & HUBEL, D. (1988). Segregation of form, color, movement and depth: Anatomy, physiology and perception. *Science*, **240**, 740-749.
- LU, C., & FENDER, D. (1972). The interaction of color and luminance in stereoscopic vision. *Investigative Ophthalmology*, **11**, 482.
- LUCK, S., HILLYARD, S., MANGUN, G., & GAZZANIGA, M. (1990). Independent hemispheric attentional systems mediate visual search in split brain patients. *Nature*, **342**, 543-545.
- MCLEOD, P., DRIVER, J., & CRISP, J. (1988). Visual search for a conjunction of movement and form is parallel. *Nature*, **332**, 154-155.
- MISHKIN, M., UNGERLEIDER, L., & MACKO, K. (1983). Object vision and spatial vision. *Trends in Neuroscience*, **6**, 414-417.
- MORAGLIA, G. (1989a). Display organization and the detection of horizontal line segments. *Perception & Psychophysics*, **45**, 265-272.
- MORAGLIA, G. (1989b). Visual search: Spatial frequency and orientation. *Perceptual & Motor Skills*, **69**, 675-689.
- MORAN, J., & DESIMONE, R. (1985). Selective attention gates visual processing in extrastriate cortex. *Science*, **229**, 782-784.
- MORDKOFF, J. T., YANTIS, S., & EGETH, H. E. (1990). Detecting conjunctions of color and form in parallel. *Perception & Psychophysics*, **48**, 157-168.
- MORRONE, M., BURR, D., & SPINELLI, D. (1989). Discrimination of spatial phase in central and peripheral vision. *Vision Research*, **29**, 433-445.
- NAKAYAMA, K., & MACKEN, J. (1990). Sustained and transient components of focal visual attention. *Vision Research*, **29**, 1631-1647.
- NAKAYAMA, K., & SILVERMAN, G. (1986). Serial and parallel processing of visual feature conjunctions. *Nature*, **320**, 264-265.
- NAVON, D. (1990a). Does attention serve to integrate features? *Psychological Review*, **97**, 453-459.
- NAVON, D. (1990b). Treisman's search model does not require feature integration: Rejoinder to Treisman (1990). *Psychological Review*, **97**, 464-465.
- NEWCOMBE, F., & RUSSELL, W. (1969). Dissociated visual perceptual and spatial deficits in local lesions of the right hemisphere. *Journal of Neurology, Neurosurgery & Psychiatry*, **32**, 73-81.
- NOTHDURFT, H. (1985). Sensitivity for structure gradient in texture discrimination tasks. *Vision Research*, **25**, 1957-1968.
- PASHLER, H. (1987). Target-distractor discriminability in visual search. *Perception & Psychophysics*, **41**, 285-292.
- PASHLER, H. (1988). Cross-dimensional interaction and texture segregation. *Perception & Psychophysics*, **43**, 307-318.
- PHILLIPS, C., ZEKI, S., & BARLOW, H. (1984). Localization of function in cerebral cortex: Past, present and future. *Brain*, **107**, 327-361.
- POGGIO, T., GAMBLE, E., & LITTLE, J. (1988). Parallel integration of visual modules. *Science*, **242**, 436-439.
- POSNER, M. (1978). *Chronometric explorations of the mind*. Hillsdale, NJ: Erlbaum.
- PRINZMETAL, W., & KEYSAR, B. (1989). Functional theory of illusory conjunctions and neon colors. *Journal of Experimental Psychology: General*, **118**, 165-190.
- PRINZMETAL, W., PRESTI, D., & POSNER, M. (1986). Does attention affect visual feature integration? *Journal of Experimental Psychology: Human Perception & Performance*, **12**, 361-369.
- QUINLAN, P. T., & HUMPHREY, G. W. (1987). Visual search for targets defined by combinations of color, shape, and size: An examination of the task constraints on feature and conjunction searches. *Perception & Psychophysics*, **41**, 455-472.
- RAMACHANDRAN, V. (1987). Interaction between colour and motion in human vision. *Nature*, **328**, 645-648.
- RIVEST, J., & CAVANAGH, P. (1991). Cross media cooperation in localization of contours. *Investigative Ophthalmology & Visual Science* (Suppl.), **32**, 1024.
- SAGI, D. (1988). The combination of spatial frequency and orientation is effortlessly perceived. *Perception & Psychophysics*, **43**, 601-603.
- SAGI, D. (1990). Detection of an orientation singularity in Gabor textures: Effects of signal density and spatial frequency. *Vision Research*, **30**, 1377-1388.
- SAGI, D., & JULESZ, J. (1985). "Where" and "what" in vision. *Science*, **228**, 1217-1219.
- SAGI, D., & JULESZ, J. (1987). Short-range limitations on detection of feature differences. *Spatial Vision*, **2**, 39-49.
- SAVOY, R. (1987). Contingent aftereffects and isoluminance: Psychophysical evidence for separation of color, orientation and motion. *Computer Vision, Graphics & Image Processing*, **37**, 3-19.
- SCHWARTZ, S., & LOOP, M. (1982). Evidence for transient luminance and quasi-sustained color mechanisms. *Vision Research*, **22**, 445-447.
- SIMON, H. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- STEINMAN, S. (1987). Serial and parallel search in pattern recognition. *Perception*, **16**, 389-398.
- STROMEYER, C. F., III, & MANSFIELD, R. J. W. (1970). Colored aftereffects produced with moving edges. *Perception & Psychophysics*, **7**, 108-114.
- SWITKES, E., BRADLEY, A., & DEVALOIS, K. (1988). Contrast dependence and mechanisms of masking interactions among chromatic and luminance gratings. *Journal of the Optical Society of America A*, **7**, 1149-1162.
- TOWNSEND, J. (1972). Some results concerning the identifiability of parallel and serial processes. *British Journal of Mathematical & Statistical Psychology*, **25**, 168-199.
- TREISMAN, A. (1990). Variations on the theme of feature integration: Reply to Navon (1990). *Psychological Review*, **95**, 460-463.
- TREISMAN, A., & GELADE, G. (1980). A feature-integration theory of perception. *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., & SATO, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 459-478.
- TREISMAN, A., & SCHMIDT, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, **14**, 12-31.
- TROSCIANKO, T., & FAHLE, M. (1986). Why do isoluminant stimuli slow down? *Perception*, **15**, A29.
- TROSCIANKO, T., & HARRIS, J. (1986). Phase discrimination in chromatic gratings. *Perception*, **15**, A18.
- TSAL, Y. (1989a). Do illusory conjunctions support feature integration theory? A critical view of theory and findings. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 394-400.
- TSAL, Y. (1989b). Further comments on feature integration. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 407-410.
- UNGERLEIDER, L., & MISHKIN, M. (1982). Two cortical systems. In D. Ingle, M. A. Goodale, & R. J. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549-586). Cambridge, MA: MIT Press.
- WALTERS, D., BIEDERMAN, I., & WEISSTEIN, N. (1983). The combination of spatial frequency and orientation are not effortlessly perceived. *Investigative Ophthalmology & Visual Science*, **24**(Suppl.), 238.
- WOLFE, J., CAVE, K., & FRANZEL, S. (1989). A modified feature-integration model for visual search. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 419-433.
- ZEKI, S. (1989). Functional specialization and multi-stage integration in the visual cortex. *Perception*, **18**, A263.
- ZOHARY, E., & HOCHSTEIN, S. (1989). How serial is serial processing in vision? *Perception*, **18**, 191-200.
- ZUCKER, S. (1976). *Relaxation labelling and the reduction of local ambiguities* (Tech. Rep. No. TR451). College Park, MD: University of Maryland.

NOTE

1. One difference between computational and biological vision resides in their appeals to ultimate causation. In the biological sciences, people usually attribute the state of the world to evolution, whereas computer scientists use a computational teleology. Why does an organism have a particular physiology, behavior, and so forth? A biologist or psychologist will answer "evolutionary advantage." A computer scientist, however, will answer the same question with "reduction of computational complexity."

(Manuscript received January 18, 1991;
revision accepted for publication June 3, 1991.)

Notices and Announcements

21st Annual Meeting of the Society for Computers in Psychology San Francisco, California November 21, 1991

The 21st Annual Meeting of the Society for Computers in Psychology will be held at the Hyatt Regency Hotel at Embarcadero Center in San Francisco on November 21, 1991, the day before the annual meeting of the Psychonomic Society. The meeting will include presentations, discussions, tutorials, and times for software and hardware demonstrations. All areas of psychology are featured, including research, education, clinical practice, and industrial applications.

For further information regarding the conference, contact William L. Palya, Department of Psychology, Jacksonville State University, Jacksonville, AL 36265 (BITNET address FWLP@JSUMUS, phone (205) 782-5641, FAX (205) 782-5680).

32nd Annual Meeting of the Psychonomic Society San Francisco, California November 22-24, 1991

The 32nd Annual Meeting of the Psychonomic Society will be held in San Francisco, November 22-24, 1991. The meetings will begin Friday morning and continue until Sunday at noon. The headquarters hotel will be the Hyatt Regency San Francisco at Embarcadero Center.

The program and hotel reservation cards have been mailed to members and associates. A copy of the program will be published in the November issue of the *Bulletin of the Psychonomic Society*. Additional programs will be available at the registration desk for \$6.00.

For further information, please contact the secretary-treasurer of the Society: Cynthia H. Null, P.O. Box 7104, San Jose, California 95150-7104 (telephone: 415-604-1260).