CIS-TR-86-10

Detecting Structure by Symbolic Constructions on Tokens

Kent A. Stevens
Allen Brookes

Department of Computer and Information Science
University of Oregon

DETECTING STRUCTURE BY SYMBOLIC CONSTRUCTIONS ON TOKENS[1]

Kent A. Stevens and Allen Brookes

Department of Computer Science
University of Oregon
Eugene, OR 97403

Abstract. Geometric organization is readily detected in discrete textures such as dot patterns. A common proposal is that orientation-tuned receptive field mechanisms provide the local orientation information from which the global organizations emerge. Alternatively, the local orientation might be attributed to grouping constructions between adjacent tokens, each representing the position of a dot and its attributes such as color, size and contrast. Geometric organization would then emerge by grouping operations on selected tokens that are similar, adjacent and aligned. It is the ability to group on the basis of similarity that most strongly differentiates this from the energy-summating receptive field approach. Using dot patterns with rivalrous organization, we demonstrate grouping phenomena that are difficult to attribute to a broad class of energy summation detectors operating in the spatial frequency domain, which we therefore attribute to perceptual groupings on tokens. We discuss the computational differences between feature detection and structure detection, and suggest that orientation-tuned receptive field mechanisms, while appropriate for the former task, have little application to the latter.

1. Introduction

The notion of discrete grouping items or tokens was tacit, by and large, in the Gestalt demonstrations of similarity grouping [Wertheimer 1923; Koehler 1929; Koffka 1935]. Later, it was specifically proposed that groupings involved "place markers" [Attneave 1974] or "place tokens" Marr [1976, 1982], which individually carry information about position and attributes such as contrast, color, size and orientation (see also Ullman's [1979] grouping tokens for motion correspondence). Pairings between adjacent place tokens are represented by "virtual lines" [Attneave 1974; Marr 1976, 1982; Stevens 1978]. Similarly, Caelli and Julesz [1978] discuss "local dipoles" between neighboring dots in texture discrimination. Virtual lines are not illusory lines in the sense that subjective contours are illusory intensity edges; they do not exhibit contrast phenomena. Instead, perceived pairings make explicit the orientation, position, and separation of the similar and adjacent elements. Attneave [1955, 1974] has demonstrated position and orientation judgements that seem mediated by place tokens, and Beck and Halloran [1985] similarly suggest that virtual lines underlie some vernier acuity judgments. It is not clear, however, whether the position markers in attentive judgments of relative position and orientation under focal scrutiny are the same "place tokens" that have been proposed for early visual processing of texture. Alternative models have been proposed that do not involve explicit position markers, but instead, simple cell mechanisms that summate within even-symmetric simple-cell receptive fields.

Simple cells, it is argued, could detect the local pairings without distinguishing the constituent items, obviating the need to hypothesize symbolic grouping tokens. Hence while perceptual grouping in general seems to require some explicit representation of the constituent elements that are grouped, receptive field mechanisms seem well suited for detecting local pairings and collinear alignments without distinguishing the individual elements. Furthermore, it would follow that since such cortical cells presumably detect continuous line and edge intensity features, they might detect discrete versions of the same features, within certain limits.

---

[1] The research reported in this article was supported by Air Force Office of Scientific Research contract number F49620-83-C-0093.

This article sets out evidence for perceptual grouping in texture detected on the basis of local constructions on symbolic tokens. The evidence is somewhat indirect, by necessity. We have constructed several demonstrations in which energy[2] summation by receptive field mechanisms does not predict the perceptual outcome. Basically, in detecting line-like arrangements of dots either the individual dots are distinguished as individuals and grouped, or they are not, whereupon the individual constituents are spatially blurred or summated within a receptive field. Thus if energy summation mechanisms would fail to detect the local dot organization, the alternative explanation that remains, it would seem, is that the constituent elements are grouped explicitly.

It is of course simplistic to suggest that these two schemes are mutually exclusive. Perceptual grouping is neither exclusively symbolic nor detected in the luminance domain by simple cells. In fact, one purpose of this article is to tease apart various issues that are often confounded in studies of perceptual grouping. These include the undeniable role of low spatial frequency luminance changes (often obvious when dot patterns are even casually defocussed to the point where the dots are no longer resolvable but the overall organization remains), the difference between detection of local structure (such as pairings) and global organization (such as radial versus concentric arrangements of pairs), and issues of symbolic versus energy-summation types of processing.

While we will treat these points as we go along, at the outset we should comment a bit further on the notion of symbolic processing. Marr [1976, 1982] made quite clear the needs for symbolic processing, and particularly the distinction to be made between assertion and measurement. From his computational perspective, vision is inevitably symbolic, it is the development of assertions about the visual world, maintained within internal representations. Image events, both spatial and temporal, are noticed, measured, and described, yielding an internal description of external events. But from the neurophysiological perspective, cortical cells often have definite receptive fields, summate quite linearly their incident signals, and respond in a graded manner. Is this processing "analog" or "symbolic"? Is symbolic processing qualitatively distinct from that performed by receptive fields, only to be found later in the cortex? Or is vision symbolic soon after the transduction stage despite the fact that the signals appear continuous and graded? We suspect the latter — that it is expedient, even necessary to introduce symbolic processing at so early a stage. But the symbolic quality of vision that is obvious when one regards the intentionality and goals of visual information processing, will not necessarily be obvious in the functional behavior of the mechanisms.

If similarity of some visual quality or attribute, such as the color or contrast of adjacent dots, governs the apparent grouping of dots in a given pattern, it would seem clear that these attributes must have been explicitly available to, and distinguished by, the underlying grouping processes. The processing of visual information by attribute seems fundamentally symbolic. More clearly so is the generation of an "emergent feature" [Pomerantz, Sager & Stoever 1977], a "second-order" symbol that represents the orientation of the similar dots. The generation of symbolic groupings imposes affiliations between parts and wholes that reflects the corresponding physical arrangements. Now, if it is possible to achieve these computational tasks by summation mechanisms with receptive fields, so much the better. But that would not mean that the underlying processes are not symbolic.

The basic computational question regarding grouping seems to be: does local organization emerge by constructions imposed on representations of the constituent items? That would imply that the individual items are first marked within some representation, with various of their salient properties made accessible (perhaps by feature maps [Treisman 1982]). It seems doubtful that selectivity to particular attributes can be attributed to known receptive field mechanisms. It is also doubtful that detection operations within feature maps is the answer (see [Prazdny 1986] and

---

[2] By energy we mean some function approximating contrast integrated over the area of the dot or blob-like intensity feature.

below).

2. The Place-Token Hypothesis

At an early stage of processing, the visual system generates an array of local intensity descriptors, each of which describes the intensity change (e.g. bar or edge of given contrast, orientation, and so forth) occurring at the corresponding retinal position. Marr [1982] refers to this as the raw primal sketch (RPS). The RPS is local and unarticulated — the constituent assertions (e.g. of edge or bar) are not yet organized into larger ensembles, and their spatial arrangement is unknown (e.g. the local connectivity among adjacent edge or bar segments remains implicit). Using Marr's [1982] terminology, the full primal sketch (FPS) refers to the image description at which such structure is explicitly represented and the local intensity changes are organized in a manner that reflects their physical causes and arrangement. The FPS contents are organized by processes of both aggregation and segregation — aggregated so that elements having a common physical origin are associated, and segregated so that physically distinct elements are distinguished.

The task of imposing organization on the RPS is considerable: even simple spatial relations such as connectedness, closure, and inside-outside would be combinatorially prohibitive to compute across the RPS [Ullman 1984]). It is therefore not clear which basic organizing processes are performed routinely to generate the FPS from the RPS, but it is probable that certain, computationally tractable spatial relations, particularly collinearity and parallelism, are detected in the RPS over at least moderately global spatial extents. Collinearity and parallelism are singled out here as representative of simple geometric relations that are extracted early in vision, and which constitute part of the FPS. They are probably computed preattentively, since they are compellingly apparent in brief presentations under experimental conditions such that the given stimulus structure cannot be predicted.

For discussion, imagine a square grid of dots, which is readily organized into dotted lines, either rows or columns. The basic question we address is whether the tangents along the length of such a dotted line are "detected" by energy summating receptive fields or "constructed" by virtual lines spanning successive tokens that represent the individual dots. Most would agree that the visual appreciation for the spacing, regularity, and parallelism of the dotted rows or columns reflect perceptual processes of grouping or structure extraction. Moreover, a dotted line is seen as a perceptual whole, as a single entity. But there is some contention, however, as to whether the dots are treated as symbolic place tokens. While the Gestaltists used these patterns to illustrate grouping principles, the preference for seeing rows over columns when the dots in the rows are closer (or more similar in intensity by having alternate rows differ in intensity) lends support to an edge-detection explanation. But while it is straightforward to find so-called grouping phenomena that might be detected in the energy (or spatial frequency) domain, that exercise does not address the potential role of symbolic grouping in the perception. As we have tried to put across, the question seems to hinge on whether the very local organization depends on the attributes of the constituent elements, and not merely local energy measurements.

2.1 Constructs for Indexing and Representing Selection

Physically-related intensity changes in an image are likely to be similiar along various spatial dimensions (such as orientation, scale, sharpness, direction of motion, and contrast profile) and non-spatial attributes (such as color and intensity) [Marr 1982]. Thus by examining the local geometry among similar RPS elements the physically-correct association among the locally detected intensity changes can be recovered. Marr [1976] stresses the limitations of local intensity information in preserving connectivity (*e.g.* contours becoming discontinuous due to the contrast across an edge reversing relative to the background) and the need to reassemble the pieces. To impose any structure on a collection of elements on the basis of their geometry and attribute similiarity, requires two basic computational abilities:

i) a means to address or index elements by position and by attribute,

ii) a means to represent the selected subset of elements.

The spontaneous organization of a dot grid into an ensemble of parallel rows cannot be accounted for merely in terms of receptive field mechanisms, even assuming such mechanisms do detect the individual rows. The local orientation evidence for each individual row must be extracted and aggregated into a connected contour, the parallelism of the rows must be noticed, and the parallel rows grouped into an ensemble.

There is need, therefore, for a computational means for accessing elements by attribute and position and for addressing neighboring places where similar intensity changes occur. Marr suggests that place tokens provide this facility, and that virtual lines make explicit the spatial relationship between similar tokens.

2.2 The Scales of Intensity Change and of Structure are Independent

We suggest that ***the scale of geometric structure is substantially independent of the scale of the intensity changes that comprise it.*** Thus, for example, global organization is not necessarily carried by the more global features. Thin line segments that are detectable only at the finest scale of resolution might be arranged collinearly into contours or into parallel striations that are detectable only by examining these elements across a spatial scale, say, an order of magnitude larger than their component spatial frequencies. There are thus two distinct tasks: detecting intensity changes over a range of spatial scales, and detecting structure over a range of spatial scales. Of course, intensity changes may occur simultaneously at several scales within a given spatial area — minute edges and markings might be superimposed over large-scale intensity features. Likewise both fine-scale and larger-scale structures might superimpose within an image. In the texture of a herringbone fabric, for example, the smallest scale intensity changes correspond to the individual fibers. A slightly larger scale captures the parallel diagonal striations characteristic of the herringbone pattern, at a larger scale their vertical organization into columns is apparent. At a still larger scale one might observe folds and creases across the fabric.

The extraction of structure at any scale seems largely independent of the scale of the individual elements. We will use that observation as part of our case for the place token hypothesis, as it is difficult to reconcile such results with the simple cell model. But more generally, there are geometric relationships that require both acute sensitivity to position information and scale independence (which usually act in opposition). The resulting apparent organization is not captured by any single scale of image description. Rather it appears necessary to generate distinct structural assertions that make explicit or summarize the local geometry, *i.e.* local organization emerges by synthesis, not by detection.

3. Alternative Models

If the constituent elements in some geometric grouping are not explicitly marked, their local arrangement must be detected from their averaged spatial distribution, usually phrased in terms of blurring (low-pass filtering) or energy summation within elongated receptive fields. For example, a closely-spaced pair of dots, or a chain of collinear dots, has a power spectrum similar to that of an isolated line segment for spatial frequencies less than 1/s, where s is the dot spacing. From the point of view of an appropriately spatial frequency-tuned mechanism, the dotted line would be roughly equivalent stimuli to a continuous line of equal total energy.

Glass proposes that the local orientation is derived by correlating the activity of simple cells over small neighborhoods [Glass 1969; Glass & Perez 1973; Glass & Switkes 1976; Glass 1979]. Zucker [1983] similarly proposes a cooperative computation whereby the broad orientation tuning curves of individual receptive fields can be sharpened by combining the outputs of individual cells over local neighborhoods. Simple cells whose receptive fields are oriented with the dot pairs would presumably respond more vigorously, on average, than those cells at other orientations, so that local correlation (or similar computations) of their activity would reveal the orientation of the dot pairs in each vicinity. Similarly, Caelli and Julesz [1978] suggest that linear arrangements of dots in texture are detected by neural units with elongated receptive fields applied to the retinal image, either "a single neural feature extractor of the Hubel and Wiesel type", or a unit that "measures the quasicollinearity of adjacent dipoles by combining single neural units of a retinal neighborhood with slightly different orientation sensitivity" [Caelli & Julesz 1978, p. 172; see also Caelli *et al.* 1978; Julesz 1981]. Recently Prazdny [1984] also suggested that the dot pairings in Glass patterns (see below) are detected by "... measurements in the spatial and energy domain rather than logical operations on symbolic descriptions" (see also [Prazdny 1986]).

4. Dot Pattern Phenomenology and Pitfalls

A deceptively simple dot pattern that has become popular for examining grouping phenomena is the Glass pattern [Glass 1969]. Glass patterns (figure 1) are constructed by superimposing onto a random dot pattern a copy that has been transformed, *e.g.* by scaling or rotation. Each dot and its transformed counterpart in the superimposed copy defines a dot pair. If the copy is scaled, say, a globally radial pattern emerges, where the dot pairs are all radially aligned[3].

While the visual effect in the Glass pattern is usually attributed to the locally detected dot pairings or dipoles, there are clearly other, more global, factors contributing to the apparent organization, principally inhomogeneities in dot density that arise as artifacts of the process of generating a Glass pattern, as discussed below.

4.1 Large Scale Clusters and Density Inhomogeneities Dominate

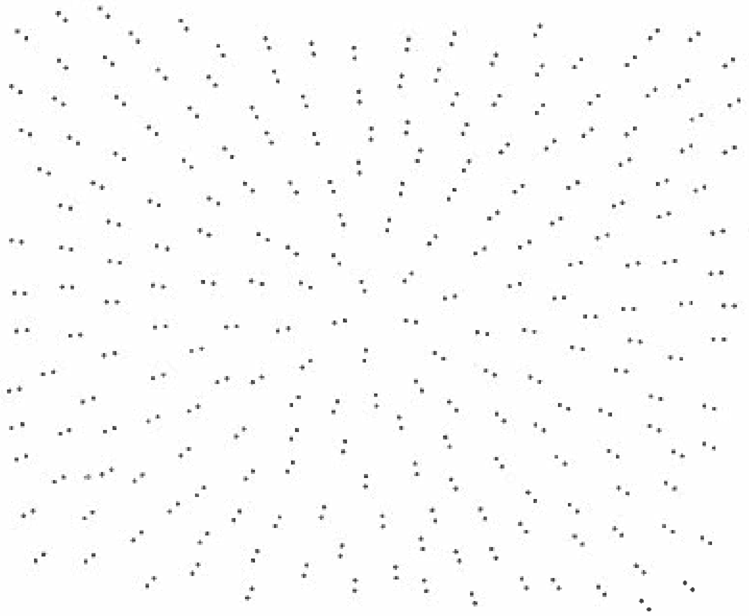

Figure 1. A radial Glass pattern based on a homogeneous density dot pattern, with homogeneous dot displacements.

---

[3] Glass's original patterns were produced by physical rotation or scaling of one copy of a dot pattern with respect to another, with the result that dot pairs were increasingly separated with distance from the center of rotation or expansion. The correlated dots were nearly touching at the center and too widely separated in the periphery for the pairings to be discerned. It is preferable to generate Glass patterns with homogeneous displacements between corresponding dots, resulting in distinct pairings across across the pattern [Stevens 1978].

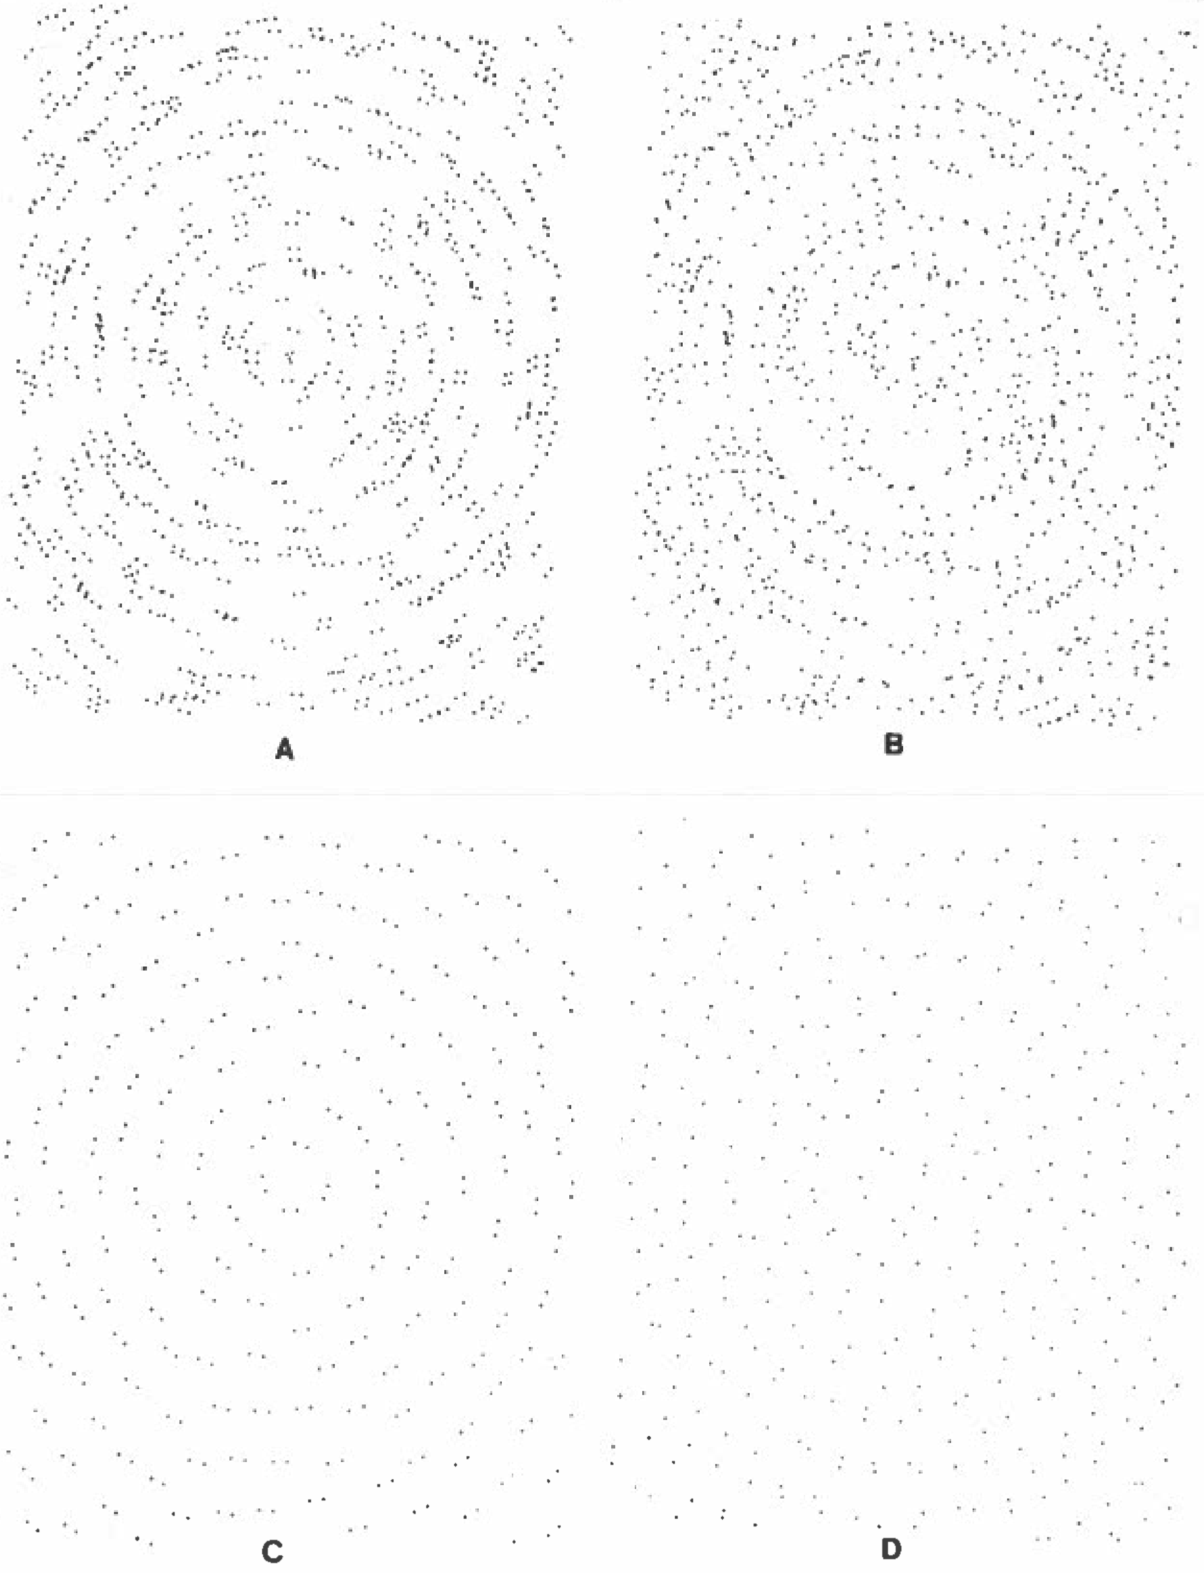

Figure 2. With a random basis pattern, low spatial frequency effects arise even when the dot pairings are no longer evident (compare organization in *a* and *b*). In contrast, *c* and *d* show how the organization in a homogeneous-density Glass pattern is carried by the pairings and not the low spatial frequencies.

Unfortunately, Glass patterns, by the way they are generated, tend to accentuate density inhomogeneities along the local transformation direction [Stevens 1978]. These low spatial frequency effects often carry a more effective orientation "signal" than the individual dot pairings. If n dots in the original dot pattern happen to align along the direction the copy will be translated, the resulting pattern will have $2n$ collinear dots. This commonly results in chains of four, six, or more dots as well as the simple pairings of dots, if not controlled for. Furthermore, inhomogeneities in dot density, which appear as clusters or voids, are selectively enhanced in the direction of the local transformation by the same process, resulting in an eccentuation of the boundaries of the clusters and voids along the same path as the individual dot pairs. Figure 2a, for example, shows a conventional Glass pattern of moderate density and dot displacements, based on a random dot pattern. While in figure 2a the concentric organization seems carried by the dot pairings, in figure 2b the displacement is so large that the pairings are no longer apparent but nonetheless the overall effect remains. The low spatial frequency components seem primarily responsible for the apparent organization.

When dealing with Glass patterns, therefore, it is necessary to make dot density as homogeneous as possible to avoid clusters and voids. We use dot patterns of constant nearest-neighbor distance. The corresponding Glass pattern (figure 2c) presents clear dot pairings, and the overall organization is virtually extinguished when the corresponding dots no longer appear paired (figure 2d).

4.2 Global Organizations may Dominate over Local Pairings

When examining the psychophysics of local dot pairings, it is important to factor out additional global effects that confound the local judgments. In particular, we have found that Glass patterns having foci (such as spiral, radial, and concentric patterns) have a more striking impression of organization than a pure translation pattern, which might well derive from factors other than the local dot pairings, as just discussed. Hence when using such patterns, the psychophysical judgment of global organization (radial versus concentric, for example) might not correspond to the perception of local pairings, if any. Since the current concern is how attribute similiarity influences local groupings, we use simple translation patterns for which the global organization is merely parallelism among the individual dot pairs. Later we return to the question of the effect of global organization (see below).

We have also used a pattern due to Marroquin [1976] which exhibits considerable global organization (figure 3). This pattern[4] has proven useful for examining global grouping tendencies, and in virtual line modelling (see section 6.2), has suggested the presence of long-range processes that detect collinearity among relatively isolated dot pairs. One may observe in this pattern various types of geometric organization, including circles, rectangles, and more complicated shapes. The pattern also exhibits clusters and voids.

5. Evidence for Place Tokens

We will discuss two types of evidence for place tokens. The first concerning the independence of the apparent groupings from the spatial frequency content of the patterns. Geometric organization can be seen in patterns that are devoid of low spatial frequencies, which is difficult to attribute to receptive field mechanisms. The second evidence is provided by demonstrating that organization on discrete items is dictated by the similarity of their properties, again contrary to

---

[4] The Marroquin pattern is generated from a square dot grid and two superimposed copies, one rotated 60° and the other 120° relative to the original. All rotations are about the center dot of the first pattern.

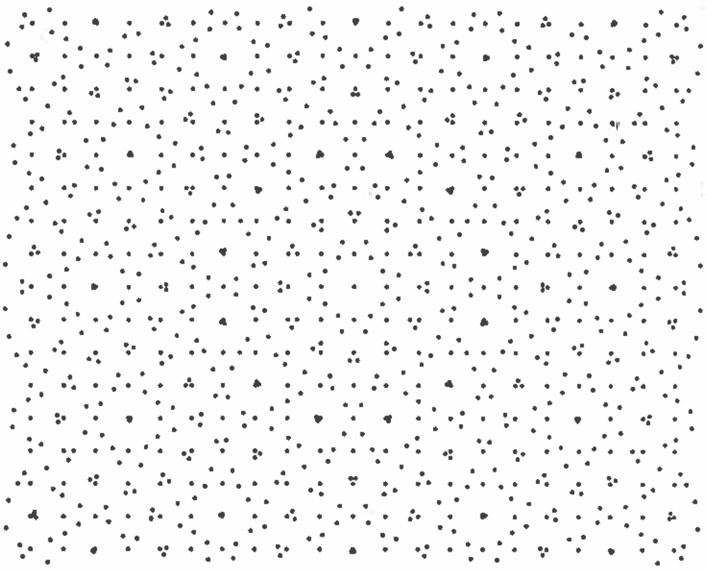

Figure 3. Marroquin pattern.

that predicted by linear energy summation models.

5.1 Patterns Devoid of Low Spatial Frequencies

Balanced Checkerboard Dots

It has been shown by Carlson *et al.* [1980] and Janez [1984] that dotted-line organization can be seen in high-pass spatial frequency filtered patterns. They conclude that some process method "more abstract" than low spatial frequency-tuned channels is involved. In a similar paradigm, we generated dot patterns with only high spatial frequency content. In the first study, the individual items were 3x3 pixels on a side — black-and-white checkerboards with the center pixel the same grey value as the background grey. The background grey matched the average luminance of the item, such that when viewed from sufficient distance the entire pattern appears a featureless grey. The individual energy-balanced "dots" were just visible when each black or white pixel subtended roughly .7′. Using balanced-energy dots in a Glass pattern, one was aware of local organization in stimuli presented for as little as 200 msec with masking. But since the features could not be resolved beyond a few degrees eccentricity, the pattern appeared a homogeneous grey except in the vicinity of the direction of gaze (the pattern scintillated parafoveally, particularly with saccadic eye movements). Wherever one fixated the correlated dots seemed paired, and the adjacent pairs were clearly parallel.

Prazdny [1986] could not replicate these observations in an experiment with similar visual stimuli[5]. When the background luminance matched the averaged Laplacian feature luminance the discrimination of global organization (spiral, circular, or ellipsoidal) approached chance level despite the fact that individual features were distinctly visible, at least in the vicinity of the point of gaze. The critical difference in task, which we believe reconciles the different results, is that

---

[5] His energy-balanced items were small Laplacian distributions, each a dark ring surrounding a bright 1.65′ diameter center on a grey background.

despite the fact that individual features were distinctly visible, at least in the vicinity of the point of gaze. The critical difference in task, which we believe reconciles the different results, is that discrimination of different global organizations required distinguishing global organizations parafoveally (over the 8.25° width of the display), where the energy-balanced features were barely discriminable. Although eye movements were allowed, it is reasonable to expect that the global organization would remain obscure when only a small portion of the pattern was clearly visible (see [Glass & Perez 1973]). The necessary control experiment would be to scale the size of the the energy-balanced features with eccentricity so that they could be simultaneously visible over a substantial field of view allowing the overall, global organization to be detected, as discussed next.

Difference-of-Gaussians Dots that Scale with Eccentricity

We used difference-of-Gaussian (DOG) dots [Carlson *et al.* [1980], but with their size carefully scaled to increase with eccentricity (analogously to the technique in [Wilson & Giese 1977]) — see figure 4 and [Stevens & Brookes 1987].

The pattern was viewed from a predetermined distance with a central fixation point. The scaling function was calibrated empirically such that the DOG dots were at threshold visibility at all eccentricities when holding fixation at the center of the display. The stimulus patterns were then viewed from a slightly closer distance, in order to provide adequate local "energy" to each DOG dot.

We found that a Glass pattern of scaled DOG dots presented global organization in much the same manner as a conventional pattern. While carefully maintaining focus on the fixation dot, when the Glass pattern was presented there was a discernable transition between viewing a

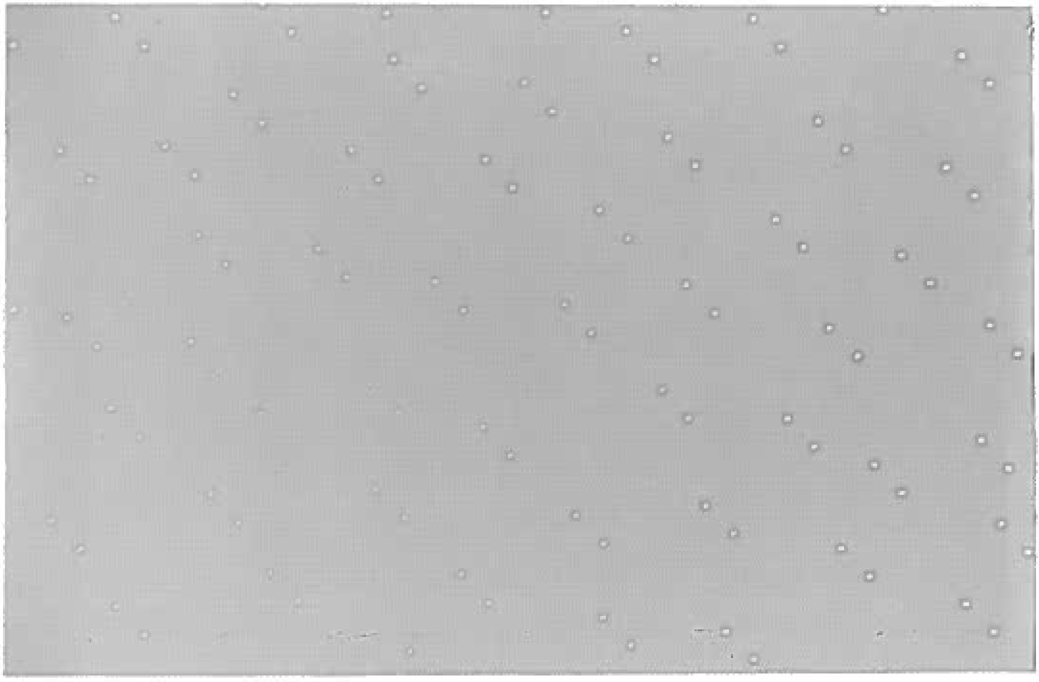

Figure 4. Portion of a Glass pattern composed of difference-of-Gaussian (DOG) "dots" that scale with eccentricity. The actual experimental stimuli were viewed with central fixation and a psychophysically matched scaling function.

disorganized field of small features and and an organized, global Glass pattern.

Spatial Summation of Energy-Balanced Patterns

Simple cells appear to be ill-suited for the detection of local pairings in energy-balanced dot patterns. The argument, already provided by Carlson *et al.* [1980] and Janez [1984] for similar "dotted-line" detection tasks, is that the stimuli have insignificant power in the range of spatial frequencies at which a correspondingly scaled simple cell would respond. We found that the impression of local pairings and parallelism is apparent when the correlated dots were separated by as much as 30′[6].

An isolated energy-balanced "dot" would maximally stimulate an (on-center) X-cell having central excitatory diameter ω equal to the size of bright spot, with the dark ring falling into the inhibitory surround. Larger X-cells would receive progressively weaker stimulation, and those with excitatory centers somewhat larger would produce insignificant response due to the cell's linearity within this range [Enroth-Cugell & Robson 1966]. At the viewing distances we used, the energy-balanced dots were "visible" only to ganglion receptive fields smaller than Wilson & Bergen's [1979] N channel ($\omega = 4.4'$), more on the order of midget ganglion cells (1.3′ at the central fovea) [Marr *et al.* 1980]. Given the linear summation exhibited by most simple cells [Movshon *et al.* 1978], an odd-symmetric simple cell (*e.g.* a S_1 cell [Schiller et al. 1976]), which might respond vigorously to a pair of regular dots located in its excitatory subfield, would receive negligible stimulation from a pair of energy-balanced dots. Such simple cells would not detect the local pairings in these patterns.

Allowing for some nonlinearity in the detection stages, it is still unlikely that receptive fields would be found that are both tuned to the very high spatial frequencies of the individual energy-balanced dots (ω approximately 1.3′) and accommodate their relatively large separations (30-40′) within their excitatory subfields. Specifically, does one find, in human vision, summation in a "receptive" or "perceptive" field that is over 20 times longer than its width? The following summarizes the negative evidence from adaptation psychophysics and the more equivocal evidence from single cell neurophysiology.

First, in adaptation studies, the summation area over which one finds threshold elevation is spatially limited to an area the size of which is reciprocally related to the spatial frequency, *i.e.* roughly 10 periods in length or width [Howell & Hess 1978; Wright 1982]. The area of functional summation is reciprocally related to spatial frequency over a large range (4-32 c/deg) of spatial frequencies. This would predict a maximum summation area of roughly 13′ for a channel driven by ganglion cells with ω=1.3′.

Similar results are found in studies of orientation sensitivity to small bars [Andrews 1967a, 1967b; Vassilev & Penchev 1976; Bacon & King-Smith 1977; Scobey 1982]. They conclude that the receptive fields in human vision that provide information about line orientation have a maximum length of about 9′ in the fovea. These results are directly relevant to the current question, since the stimuli are typically thin bars (2′ width), on the scale of channel that would be responding to the detail within the balanced-checkerboard dots. Burton & Ruddock [1978] have also shown that the threshold elevation effect is length selective when the (bright) bar length is less than roughly three times the bar width. In human vision the receptive fields of the relevant scale are apparently too short to span widely separated energy balanced dots and still be sensitive to the very high spatial frequency content that makes them visible against the background grey.

Turning to neurophysiology, the data are not as conclusive, as one might expect. Consider the dimensions of the receptive fields as mapped conventionally. The minimum width of the central excitatory region is probably the diameter ω of the constituent LGN X-cells [Hubel & Weisel 1962, 1968]. The overall dimensions of simple cell receptive fields, as conventionally mapped in

---

[6] Larger separations of as much as 1° can be tolerated, but the resulting patterns are so sparse that one can resolve only a very few pairs and the impression of preattentive groupings is less compelling, although the task is still performed with short presentation and masking.

monkey fovea, tend to be somewhat longer than they are wide, from .25° by .25° to .5° by .75° [Hubel & Weisel 1962, 1968; De Valois *et al.* 1982]. Within this overall extent, the central excitatory regions are commonly not more than 4ω long; for example, Poggio [1972] reports, for monkey simple cells in the foveal region, receptive field lengths between 6-24′ (divide by approximately 1.8 for human). On the other hand, Schiller *et al.* [1976, figure 17] show evidence for simple cells that have increasing response as the length of the bar stimulus increases up to the 6.4° examined (in monkey, at 2-5° eccentricity, therefore one should divide by roughly 4 to compare to human foveal). The evidence from direct receptive field mapping would suggest that foveal simple cell receptive fields of 30′ do not exist in human, but is not conclusive.

5.2 Examining Similarity Grouping

Rivalrous Glass Patterns

The basic paradigm for examining similiarity issues is to generate a rivalrous Glass pattern [Stevens 1978] in which two, differently-transformed, copies of a dot pattern are superimposed over the original (figure 5). By displaying the superimposed dots with differing color, intensity and displacements relative to the original dots one can pit, for example, intensity similarity against proximity. The evidence for similarity grouping demonstrated earlier in [Stevens 1978] is reiterated and extended here.

Because of the strong influence on the apparent local organization induced by certain global organizations we use simple translation patterns. We also use diagonal translations to avoid the known biases towards the vertical and horizontal and homogeneous dot density basis patterns for

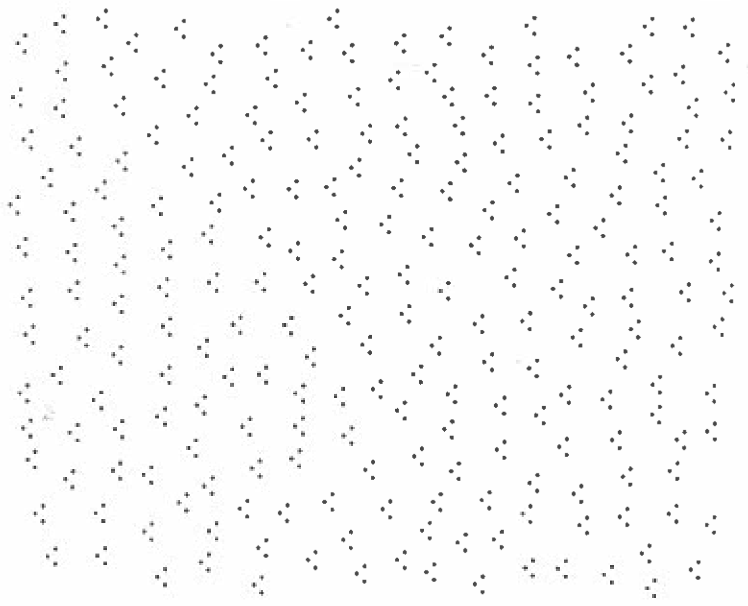

Figure 5. A homogeneous dot pattern with two differently-translated copies superimposed results in a rivalrous Glass pattern.

reasons discussed earlier. In the experiments we used masked tachistoscopic presentations.

Proximity and Intensity Similiarity

By manipulating intensity and proximity one can influence the subjective pairings. In figure 6a the impression is of diagonal lines leading upward to the right — the pairings are between dim dots rather than either combination of bright and dim dots. There is a preference for grouping dots of equal intensity even if there are adjacent dots of greater intensity [Stevens 1978]. This is, of course, contrary to energy-summation, which would predict pairings between the more energetic dim and bright dots, not between the dim dots.

Prazdny [1984, p. 474, also 1986, p. 158] was unable to replicate this result, for reasons which we attribute to his asking subjects to perform a different experimental task, namely the perception of global organizations rather than the perception of local groupings. Subjects saw apparent global organization (*e.g.* radial versus concentric) according to the "energy" (dim-bright) and not "similarity" (dim-dim) orientation. Prazdny concludes the global organization is extracted by local operations in the energy domain. But recall the low frequency effects demonstrated in figure 2, wherein the overall Gestalt may be apparent even when the local dot pairings are not — the overall effect is likely derived from energy measurements. Similarly, global organization in Prazdny's [1984, 1986] demonstration patterns is discernable even when they are defocussed. Furthermore, the correlated dots in his patterns are so widely separated that the "signal" is carried more by streaks or chains of dots than by local dot pairs *per se* (see also figure 2b).

To examine local groupings, we use homogeneous-density dot patterns to minimize the dominating low spatial frequency effects, and translation patterns to equate the local and global organizations. Under these conditions, which better reveal the local phenonema, ***we find that the pairings are based on similarity measures, not energy*** (figure 6a).

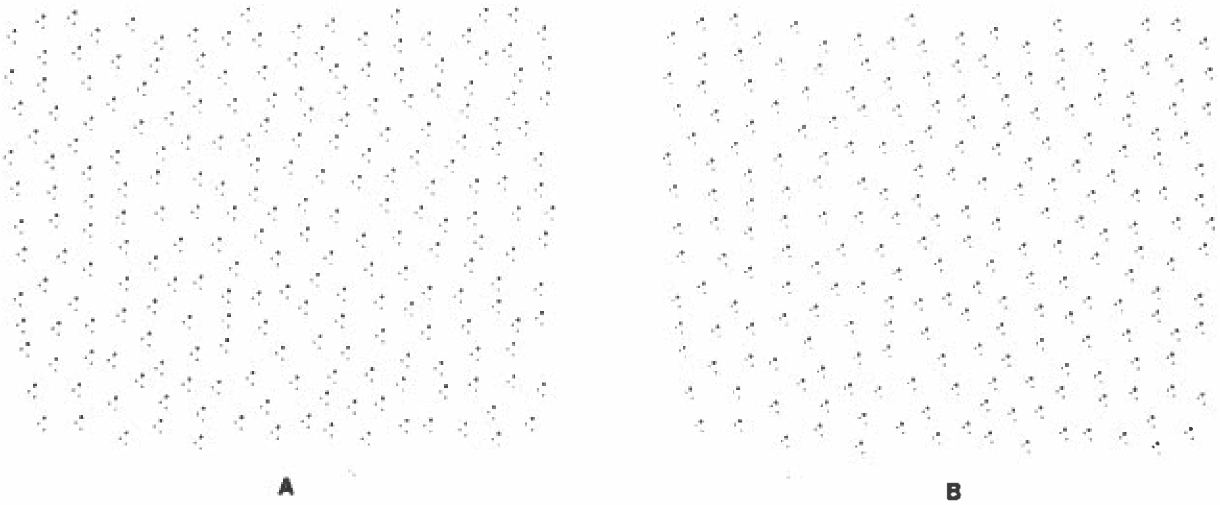

Figure 6. Similar intensity dots are paired, even when dimmer (*a*) and not nearest neighbors (*b*).

There is a well-known tendency to prefer proximity in grouping. The preference for intensity similarity, however, is stronger. In figure 6b the pitted against each other the pairing is between the dots of equal intensity (figure 6b). As the intensity of the bright dot is further increased the pairings between dim dots are disrupted, presumably by a tendency to attend to the more energetic signal — the global arrangement of bright dots. With even further increase in the intensity of the bright dots the dim dots are disregared altogether.

Color Similarity

Color (hue) similarity appears to be the strongest attribute for pairing place tokens. It is easily established that similarity of color strongly determines the apparent pairings in the rivalrous patterns. It is difficult, in fact, to override the pairings established by equal color. To group on the basis of intensity similarity over color similarity, the equal-intensity dots must have (much) higher intensity. We found that there is always a preference for color over proximity, which agrees with the earlier findings that proximity was a weak attribute. We also found that we could pit equal color against small differences in intensity proximity and size and still prefer the pairing of the equal-colored dots.

Prazdny [1986] observed correctly that the correlated features in a Glass pattern do not have to have the same hue for the organization to be apparent, provided they have the same contrast polarity against the background. He concludes that "the processes in the perception of Glass patterns are thus 'color blind' " [Prazdny 1986]. But tolerance to color dissimilarites does not demonstrate insensitivity to color similarity. Our rivalrous Glass patterns demonstrate an extraordinarily strong influence of color similarity on perceived groupings, with pairings seen between dots of equal color even when they are dissimilar in energy and not nearest neighbors.

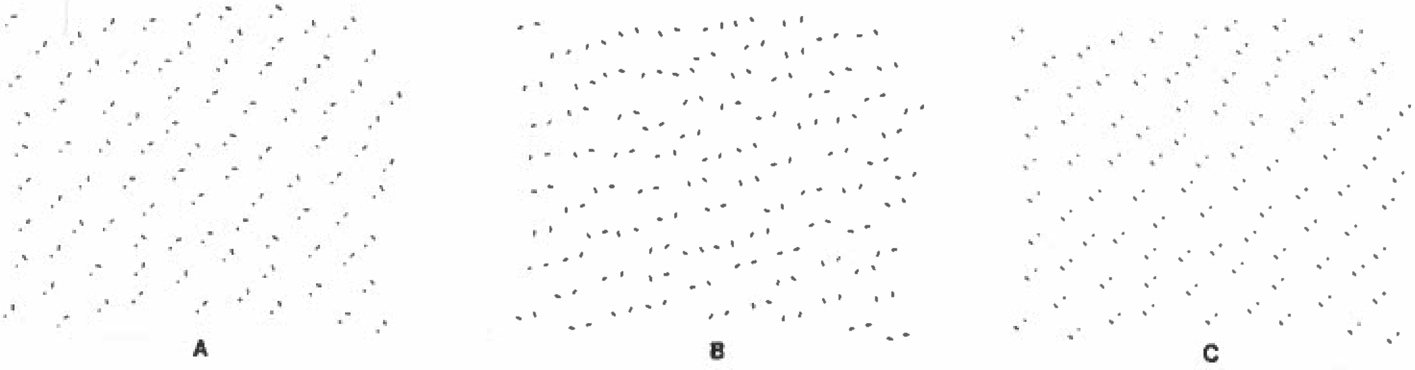

Figure 7. Dots and randomly-oriented bars (*a*) can be seen as paired in the absence of competing organizations, as can pairs of bars of random orientation (*b*). In *c* it is also possible to attend to the global parallelism of the bars, disregarding the adjacent dots.

Orientation Similarity

Up to now we have discussed pairings between dots. Figure 7a shows pairings between dots and randomly oriented bars, with a clear overall "/" diagonal structure defined by the orientation of the pairs. An adjacent dot and bar appears to group together regardless of the bar orientation. In the absence of competition, virtually any combination of adjacent items, regardless of shape, color, size, and orientation can be seen as grouped (assuming they have the same contrast sign relative to the background) — see [Marr 1976, figure 10] and figure 7b. But orientation does play a role in grouping, actually a complicated role. In figure 7c the adjacent dots and bars appear paired, each pair oriented "/", but one can also attend to the globally parallel orientation of the bars alone, disregarding the dot-bar pairings. (see also [Stevens 1978]). The psychophysics of bar patterns is thus complicated by the fact that one can attend to either the local groupings or, more globally, to the individual bars.

While isolated dots and bars might pair, their grouping affinity is not particularly strong for an arbitrarily oriented bar, and can be overriden easily. Figure 8a serves to demonstrate again the preference for proximity, with the resultant "/" orientation. In figure 8b the dominant pairings are seen between the adjacent dots, despite the closer proximity of the dot-bar pairs and the greater overall energy carried by those pairs. This suggests that similarity of type (*e.g.* dots pairing with dots) is also salient, along with similarity of intensity and color. In figure 8c an attempt is made to induce preference for pairings between dots and bars, by introducing collinearity between the bars and dots as well as proximity (and greater overall energy). Most observers find this pattern rivalrous, alternating between seeing dot-bar pairings (or perhaps merely attending to the globally parallel bars, disregarding the dots altogether), or dot-dot pairings (disregarding the adjacent bars).

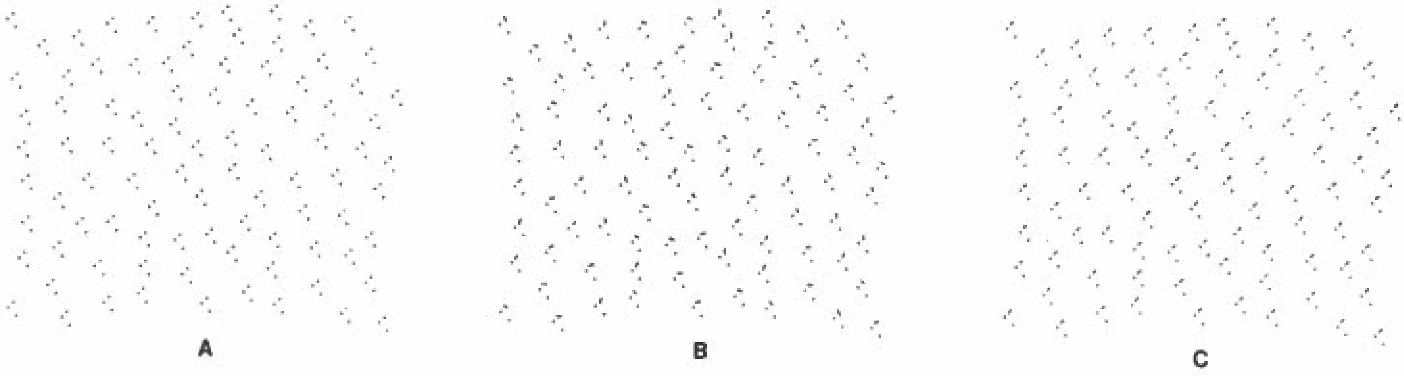

Figure 8. While there is a preference for proximity (*a*), one prefers in *b* to group together dots rather than to group adjacent dots and bars. In *c* the collinearity between the dots and bars increases the tendency to see dots and bars as paired.

In Figure 9a there is tendency to prefer the collinear dot-bar pairs (as in figure 8c) over the bars of dissimilar orientation. In Figure 9b the adjacent bars have identical orientation, and appear to pair together. Importantly, the "/" orientation of the pattern, as defined by the bar pairs, is perpendicular to the "\" orientation of the individual bars across the pattern (as well as the orientation of the bar-dot pairs). Thus despite several very local orientation "signals" indicating "/" orientation, there is a clear impression of a perpendicular orientation defined by the pairings between similarly-oriented bars.

As expected, the greatest pairing tendency occurs when the adjacent items consist of collinear bars, all other factors being constant. But surprisingly, introducing a modest difference in intensity or color (i.e. introducing a rivalrous pairing that is not collinear but similar in intensity or color) can defeat pairing on the basis of orientation or collinear similarity.

Effects of Global Organization

We mentioned earlier the strong effects introduced when the Glass pattern has a global organization other than mere parallelism (introduced by simple translation of one dot pattern with respect to another). In the rivalrous patterns composed of two translation patterns, the alternative global organizations were either diagonally oriented pairings all aligned either "/" or "\". Since there was no consistent bias for either alternative, the local pairings were not influenced by the global pattern. The following patterns, however, offer competing global organizations, which have the effect of influencing the apparent local structure. Figure 10a shows a rivalrous Glass pattern composed of the original homogeneous density dot pattern, one superimposed pattern translated and the other scaled. The displacement between correlated dots is the same in both the translation and scaling directions. The translated pattern forms diagonal pairings all oriented "/"; the scaled pattern produces a radial expansion effect. The dominant impression is of radial organization,

---

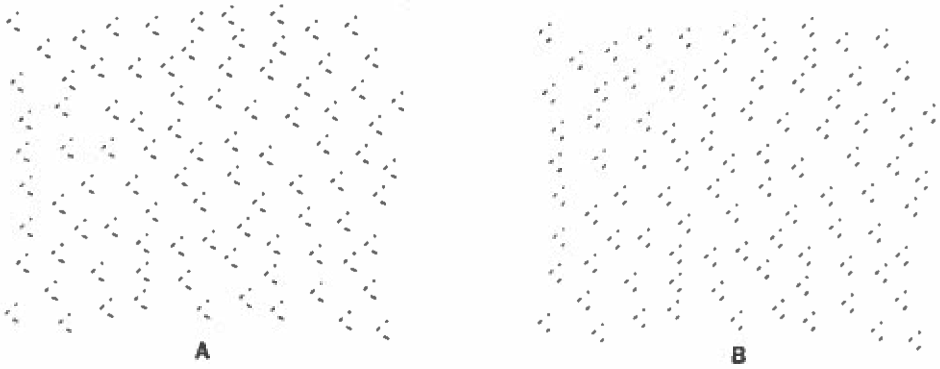

Figure 9. In *a* there is a tendency to prefer local pairings between the dots and bars on the basis of collinearity (as in figure 8c). In *b* the preferred pairings are between bars of equal orientation, resulting in a dominant "\"orientation.

even in Figure 10b where intensity similiarity would suggest the translation pattern (as in figure 6).

The presence of a focus of expansion provides a compelling focal point of attention as well. One is drawn to the center of the pattern, and, as Prazdny [1984] observed, that global organization is preferred (figure 10b).

A stronger global organization than the radial pattern is that of concentric curves. Figure 11 shows rivalrous patterns where one set of correlated dots are scaled, producing a radially pattern, and the other set is rotated about the same center. In figure 11a the correlated dots are displaced equally in the radial and concentric directions, and, while both local organizations are discernable, there is a clear preference for the globally concentric organization, even when the radial pairs are favored by proximity (figure 11b). There is even a ghost-like concentric pattern present when the pattern is designed to strongly emphasize the radial pairings (figure 11c), particularly in the parafovea.

The visual system's apparent propensity for detecting curved arcs is further demonstrated in figure 12a, where the rivalrous Glass pattern is compose of radial and spiral transformations. With scrutiny one can verify that in each triple two of the three dots are radially oriented, but it is extraordinarly difficult to attend to the global radial organization. The pattern appears globally as a complex of rosettes, of nested curves, each suggested by dot pairs aligned along spiral arcs. In figure 12b the dots corresponding to the radial pattern are brighter in order to demonstrate the radial pattern that is present but invisible in figure 12a.

6. Extracting Structure

6.1 Attribute selection

One of the principal computational problems of the primal sketch is to construct descriptions of global organization on the basis of local evidence. The task must be achieved, at least in part, by

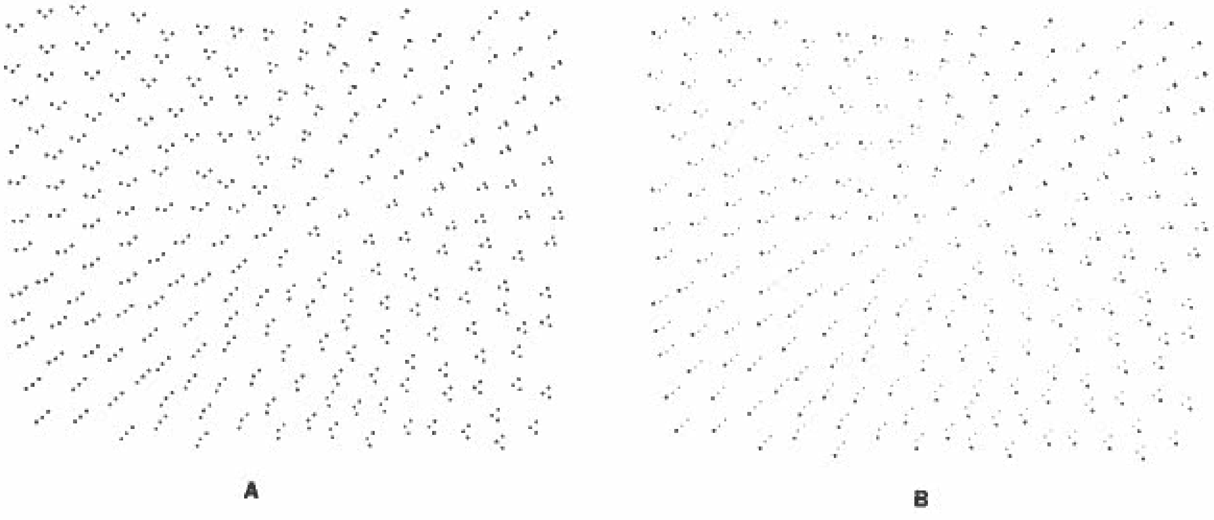

Figure 10. Examples of rivalry between pure translation and radial patterns.

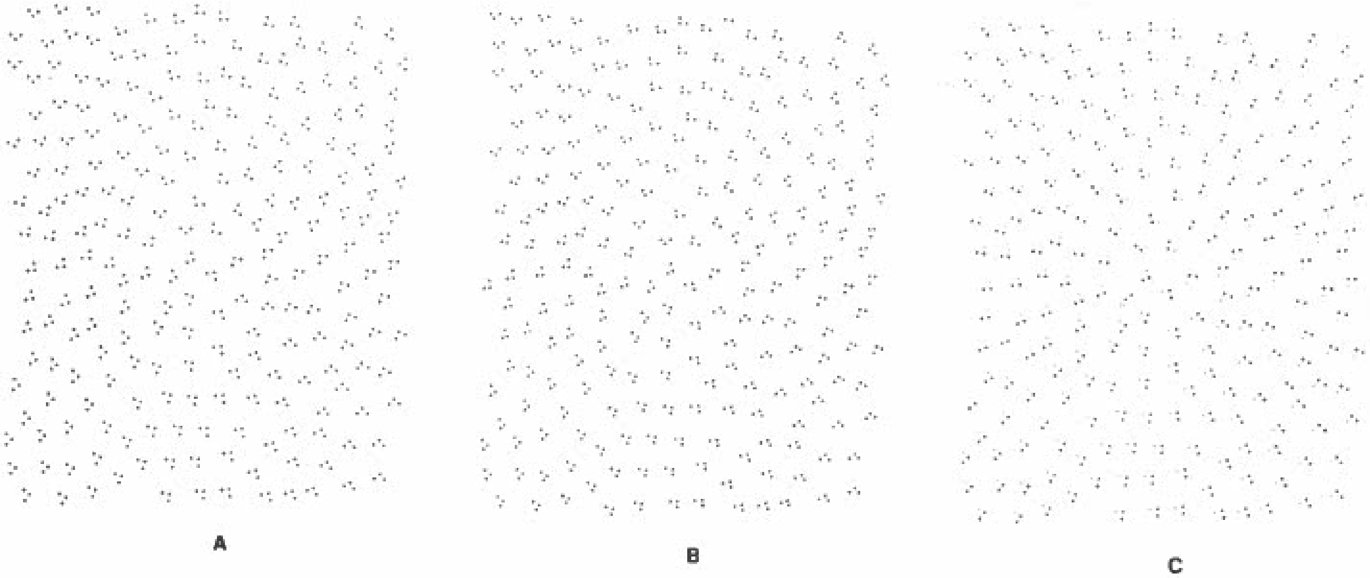

Figure 11. Examples of rivalry between radial and concentric patterns.

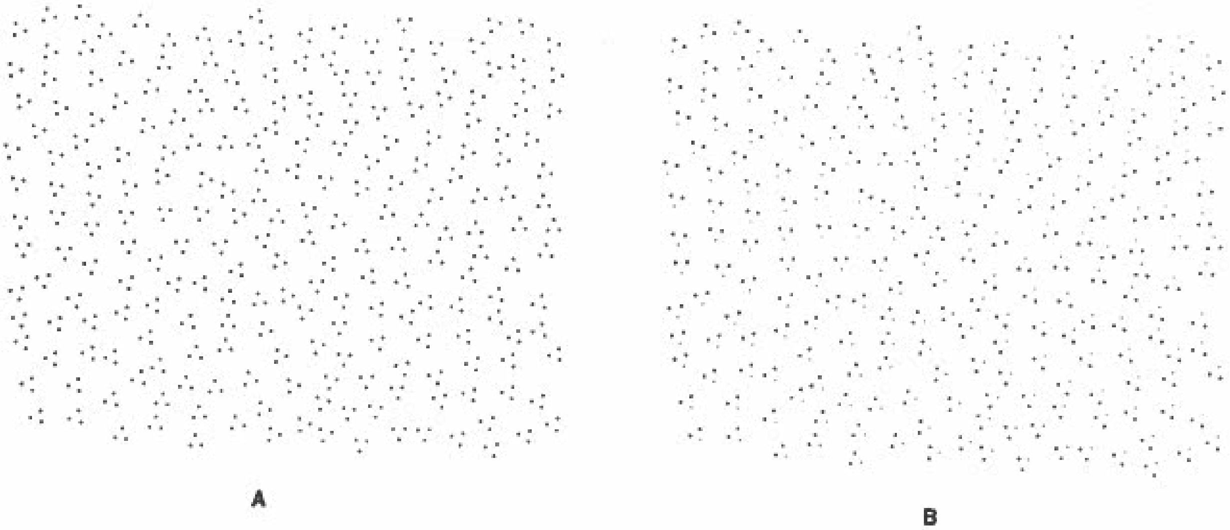

Figure 12. Examples of rivalry between radial and spiral patterns.

"bootstrapping". An issue is how to initiate this process. Marr [1982] observed that

The items generated on a given surface by a reflectance-generating processing acting at a given

> scale tend to be more similar to one another in their size, local contrast, color, and spatial organization than to other items on that surface [Marr 1982, p. 47].

As a strategy for detecting structure, we suggest that this observation can be inverted in the following way. Similarity and structure are generally correlated, i.e. *items at a given scale that are similar in size, local contrast and color tend to have a common spatial organization.* That is, (geometrically) structured intensity changes are probably similar along various (non-geometric) attribute dimensions, while uncorrelated intensity changes, such as those arising from different physical causes, are not expected to be similar except by coincidence. Hence: *similar intensity changes in any locality are likely to be structured as well. A good place to look for geometric organization is in those subpopulations defined by similarity.* Thus as an early step in the bootstrapping process of extracting structure we suggest the visual system seeks evidence of similarity in various non-geometric dimensions to partition the intensity changes into subpopulations which are subsequently analyzed for geometric organization. And the more similarity shared by a set of elements (edge segments, say) the more likely they are physically correlated.

Evidence for similarity is therefore sought by examining the spatial distribution of values of certain attributes (such as color, contrast, and so forth). A prominent band or peak within the distribution is likely to reflect related items. In terms of feature maps, wherein properties are mapped out in distinct representations, one might expect analysis of individual maps, with perhaps mutual support across maps. Recalling the earlier discussion of selection, by this proposal a distinguishable "signal" in any selectable attribute may initiate the bootstrapping.

6.2 Virtual lines represent pairings of selected items

Having selected a subpopulation of elements of similar attribute, individually represented as place tokens, their local arrangement must be represented. Virtual lines could make explicit a pair of locations, an angle, and potentially a length. Local operations on virtual lines would then begin to build geometric relations among tokens. Parallelism, as among dot pairs in the Glass pattern, could be detected by selecting those virtual lines that share the same orientation as the prominent virtual line orientation within a given spatial region [Stevens 1978]. Note the repetition of the general theme of computation, selection, computation, etc. Likewise, collinearity, such as among the dots in the chains also observed in Glass patterns, could be detected by pairwise collinearity of the corresponding virtual lines.

Virtual lines appear to be piecewise-linear constructions spanning the selected place tokens, not continuous curves (e.g. splines). This observation poses a particularly challenging issue for algorithms that construct virtual lines, because the connections or lines that are to be constructed need to connect those neighboring tokens that are collinear, but not necessarily isolated, and not necessarily at any given, fixed, separation. We have performed computational experiments to determine the utility of constructing virtual lines by spatial blurring of place tokens (*n.b.* we are not referring to low spatial frequencies in the image intensity domain, but to a blurring of the discrete token array). We have found that blurring schemes are successful only for isolated pairs or chains of dots. Collinear dots imbedded in a background of extraneous dots, as are the circles and squares observed in the Marroquin pattern, are not successfully extracted by such means. Spatial blurring (whether in the image intensity or place token domain) only serves to find density inhomogenieties: clusters and voids.

From empirical exploration of dot pairing algorithms, we conclude that the method for constructing pairings in human vision likely has the following properties: an isotropic (circular) support for detecting nearest neighbor to a given token, a proximity measure that scales (implicitly or explicitly) with the density of selected tokens. (Recall that brighter dots can be selected independently of the number of nearer dimmer neighbors, suggesting that the groupings among these dots occur after selection.) While dot pairings are substantially independent of scale, "relatively isolated" dots (compared to the mean dot density) are not paired with neighbors.

In [Stevens 1978] a simple virtual line algorithm was proposed in order to show that iterative, relaxation or cooperative processes are not needed to extract local parallelism, as exemplified by the Glass patterns. A local process first defines virtual lines between neighboring tokens, along the lines of [O'Callaghan 1974a, 1974b, 1975] where all neighbors within some factor k (typically betwen 1.3-1.5) times the nearest neighbor distance are connected. Then the virtual lines are histogrammed in each vicinity, resulting in a peak orientation if parallelism is present. It is then a simple matter to select those virtual lines that have approximately the same orientation as the local peak. The algorithm was not posed as a model for the detection of local parallelism in these patterns, but rather as a demonstration that a simple, noniterative computations might serve (see [Marr 1982] for discussion).

This virtual line algorithm was applied to the Marroquin pattern (figure 13a), and in a second step, those virtual lines that are pairwise collinear to within ±20-30° are selected (figure 13b). The algorithm is rather successful in making explicit the apparent connectivity among dots that one sees in the Marroquin pattern. Note that while this algorithm establishes only the local connectivity, subsequent selection of extended chains of virtual lines would reveal much of the various curvilinear and rectilinear figures seen in this pattern. The algorithm has also been applied to various random patterns including the experimental stimulus patterns in [Caelli *et al.* 1978; figure 4] (see figure 14). In figure 15 the algorithm constructs virtual lines and indicates with solid lines those that are pairwise-collinear.

This collinearity selection algorithm, like the earlier histogramming algorithm, is introduced only to demonstrate some computational principles. It shows that simple selection criteria applied in parallel but locally can bootstrap some global organization. This discussion treated collinearity selection merely as a geometric issue, and ignored issues of attribute similarity. Feature maps [Treisman 1982] offer a possible means for separating tokens of differing attribute, so that items that are similar are explicitly segregated from other items. It is not likely, however, that similarity preference is achieved merely by performing collinearity detection independently within separate feature maps (detecting the collinearity among red dots, ignoring the adjacent dots of different color, say). The observed tolerance for pairings between dissimilar elements (e.g. color dissimilarity [Prazdny 1986]) in the absence of competition would suggest that groupings can be

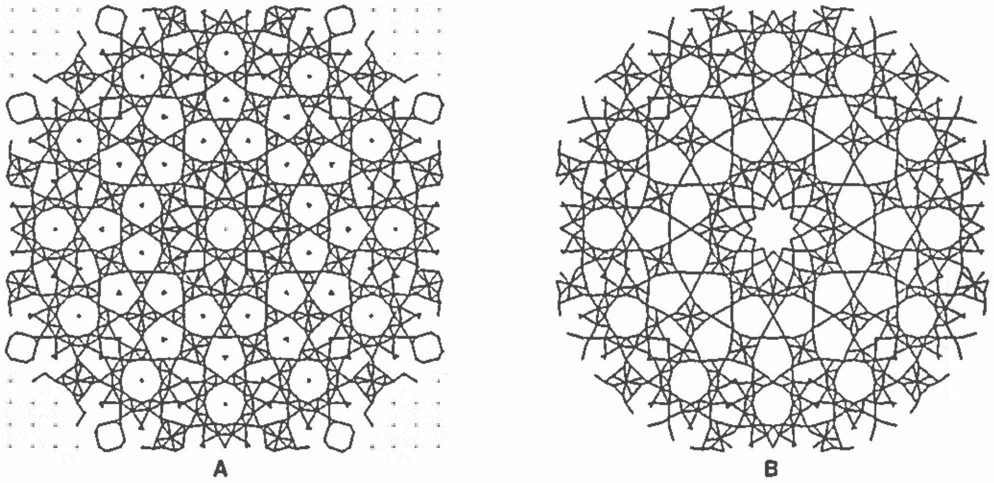

Figure 13. Of the virtual lines constructed on a Marroquin pattern in *a*, those that are pairwise collinear are selected in *b*.

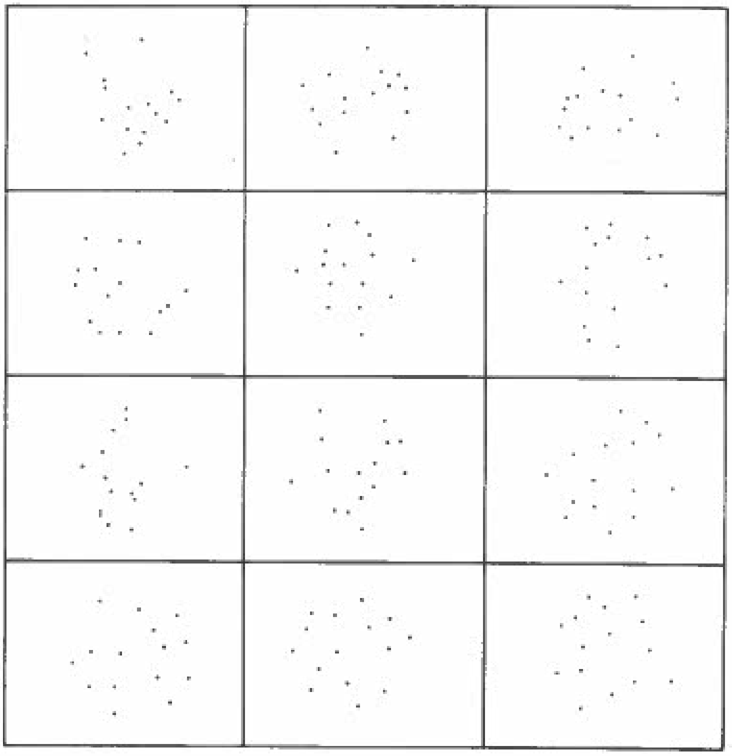

Figure 14. Dot grouping stimulus patterns from [Caelli *et al.* 1978, figure 4.].

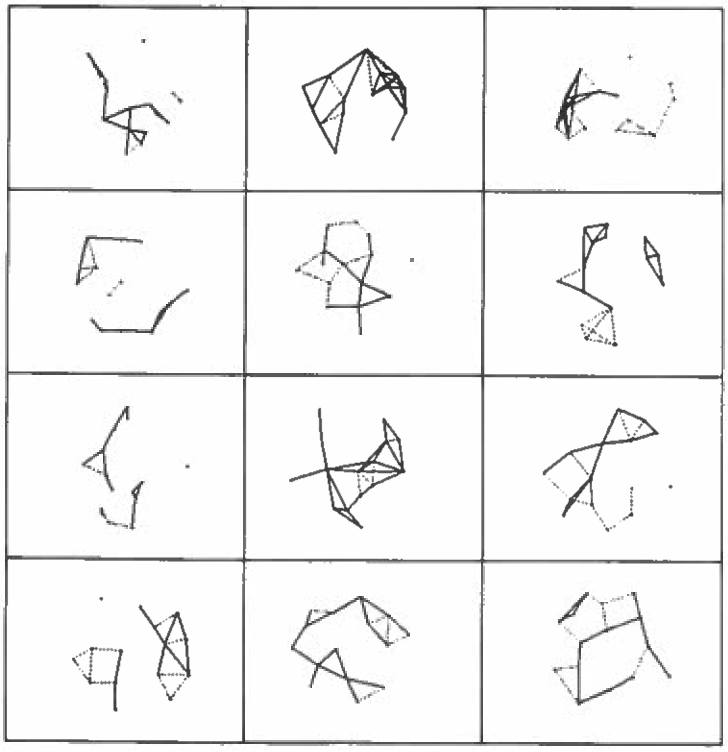

Figure 15. Pairwise-collinear virtual lines computed on the dot patterns from [Caelli *et al.* 1978,

figure 4.].

achieved across, as well as within, feature maps. The rather continuous trade-off of proximity with similiarity along various dimensions seems to rule out the strict model of first selecting similar tokens then performing pairing operations within this selected subpopulation, as we used to open this section. The grouping processes in human vision more likely determine local pairings on the basis of combined geometric and non-geometric factors.

7. Summary

1. To detect structure within a collection of elements on the basis of local geometry and attribute similiarity requires the addressing or indexing of the elements by position and by attribute, and the representing of those that are selected.

2. The scale of local geometric structure (*e.g.* collinearity, parallelism) is, we propose, substantially independent of the scale of the constituent intensity changes. There are thus two distinct computational problems: detecting intensity changes across spatial scales, and detecting structure across spatial scales.

3. Local structure can be detected in patterns containing spatial frequencies only at the scale of the individual elements, suggesting that local pairings can be constructed independently of the spatial frequency content of the elements. Global structure requires that the size of the elements (difference-of-Gaussians or Laplacians) scale appropriately with eccentricity.

4. Global organization is detected, in part, by "bootstrapping" from local organization (*e.g.* pairing) and in part by low spatial frequency features. In Glass patterns, visually compelling streaks and clusters due to dot density inhomogeneities are introduced as artifacts of the process of generating the Glass pattern. These low spatial frequency effects, unless controlled for, can suggest global organization independently of any discernable local grouping organization. Homogeneous (constant nearest-neighbor) dot density patterns with constant displacements between correlated dots, and global translation is recommended for the study of local grouping phenomena in Glass patterns.

5. Using rivalrous Glass patterns it can be shown that local pairings are constructed on the basis of similarity and not energy. Dim dots are seen as paired on the basis of intensity similarity despite greater proximity and "energy" in the alternative pairings. Color (hue) is seemingly the strongest similarity attribute for pairing. Orientation similarity also plays a role in grouping of adjacent bars, with greatest preference for bars that are collinear.

6. While global organization is, in part, "bootstrapped" from local groupings, the visual system seeks to organize the image into extended curvilinear (line-like) groupings, the evidence for which might be present only at scales much larger than the individual pairings. While parallelism might be global selection of local orientation features, collinearity seems to be a more constructive task of fitting a curve through the local orientation evidence.

7. A basic bootstrapping principle follows from the similarity and structure are generally correlated (items at a given scale that are similar in size, local contrast and color tend to have a common spatial organization). A good place to look for geometric organization is in those subpopulations defined by similarity of non-geometric attributes.

References

Andrews, D.P. 1967a Perception of contour orientation in the central fovea. Part I. Short lines. *Vision Research* **7**, 975-997.

Andrews, D.P. 1967b Perception of contour orientation in the central fovea. Part II. Spatial integration. *Vision Research* **7**, 999-1013.

Attneave, F. 1955 Perception of place in a circular field. *American Journal of Psychology* **68**, 69-82.

Attneave, F. 1974 Apparent movement and the what-where connection. *Psychologia* **17**, 108-120.

Bacon, J. & King-Smith, P.E. 1977 The detection of line segments. *Perception* **6**, 125-131.

Beck, J. & Halloran, T. 1985 Effects of spatial separation and retinal eccentricity on two-dot vernier acuity. *Vision Research* **25**, 1105-1111.

Burton, G.J. & Ruddock, K.H. 1978 Visual adaptation to patterns containing two-dimensional spatial structure. *Vision Research* **18**, 93-99.

Caelli, T.M. & Julesz, B. 1978 On perceptual analyzers underlying visual texture discrimination: Part I. *Biological Cybernetics* **28**, 167-175.

Caelli, T.M., Preston, G.A.N., & Howell, E.R. 1978 Implications of spatial summation models for processes of contour perception: a geometric perspective. *Vision Research* **18**, 723-734.

Carlson, C.R., Anderson, C.H. & Moeller, J.R. 1980 Visual illusions without low spatial frequencies. *Invest. Ophthal. visual Sci., Suppl.* **19**, 165.

De Valois, R.L., Albrecht, D.G., & Thorell, L.G. 1982 Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research* **22**, 545-559.

Enroth-Cugell, C. & Robson, J.G. 1966 The contrast sensitivity of retinal ganglion cells of the cat. *J. Physiol.* **187**, 517-552.

Glass, L. 1969 Moire effect from random dots. *Nature* **223**, 578-580.

Glass, L. 1979 Physiological mechanisms for the perception of random dot moire patterns. In *Pattern formation by dynamic systems and pattern recognition.* H. Haken, ed. Berlin: Springer-Verlag.

Glass, L. & Perez 1973 Perception of random dot interference patterns. *Nature* **246**, 360-362.

Glass, L. & Switkes, E. 1976 Pattern recognition in humans: correlations which cannot be perceived. *Perception* **5**, 67-72.

Howell, E.R. & Hess, R.F. 1978 The functional area for summation to threshold for sinusoidal gratings. *Vision Research* **18** 369-374.

Hubel, D.H. & Wiesel, T.N. 1962 Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106-154.

Hubel, D.H. & Wiesel, T.N. 1968 Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* **195**, 215-243.

Janez, L. 1984 Visual grouping without low spatial frequencies. *Vision Research* **24**, 271-274.

Julesz, B. 1981 Figure and ground perception in briefly presented isodipole textures. In *Perceptual organization*, Kubovy, M. & Pomerantz, J.R., eds. Hillsdale, New Yersey: Lawrence Erlbaum Associates, 27-54.

Koehler, W. 1929 *Gestalt Psychology.* New York: H. Liveright.

Koffka, K. 1935 *Principles of Gestalt Psychology.* New York: Harcourt, Brace.

Marr, D. 1976 Early processing of visual information. *Phil. Trans. R. Soc. Lond. B* **275**, 483-524.

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

Marr, D., Poggio, T. & Hildreth, E. 1980 Smallest channel in early human vision. *J. Opt. Soc. Am.* **70**, 868-870.

Marroquin, J.L. 1976 Human visual perception of structure. M.S. Thesis. Department of Electrical Engineering and Computer Science, MIT.

Movshon, J.A., Thompson, I.D. & Tolhurst, D.J. 1978 Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J. Physiology* **283**, 53-77.

O'Callaghan, J.F. 1974a Computing the perceptual boundaries of dot patterns. *Computer Graphics and Image Processing* **3**, 141-162.

O'Callaghan, J.F. 1974b Human perception of homogeneous dot patterns. *Perception* **3**, 33-45.

O'Callaghan, J.F. An alternative definition for "neighborhood of a point". *IEEE Trans. Computers*, Nov., 1121-1125.

Pomerantz, J.R., Sager, J.C., & Stoever, R.J. 1977 Perception of wholes and their component parts: some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance* **3**, 422-435.

Poggio, G.F. 1972 Spatial properties of neurones in striate cortex of unanesthetized macaque monkey. *Investve Ophth.* **11** 368-377.

Prazdny, K. 1984 On the perception of Glass patterns. *Perception* **13**, 469-478.

Prazdny, K. 1986 Some new phenomena in the perception of Glass patterns. *Biological Cybernetics* **53**, 153-158.

Schiller, P.H., Findlay, B.L. & Volman, S.F. 1976 Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *J. Neurophysiol.* **39**, 1288-1319.

Stevens, K.A. 1978 Computation of locally parallel structure. *Biological Cybernetics* **29**, 19-28.

Stevens, K.A. & Brookes, A. 1987 Perceptual grouping on symbolic tokens. *Biological Cybernetics* in preparation.

Treisman, A. 1982 Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance* **8**, 194-214.

Ullman, S. 1979 *The interpretation of visual motion*. Cambridge, MA: MIT Press.

Ullman, S. 1984 Visual Routines. *Cognition* **18**, 97-159.

Vassilev, A. & Penchev, A. 1976 Spatial and temporal summation in the perception of lines. *Vision Research* **16**, 1329-1335.

Wertheimer, M. 1923 Untersuchungen zur Lehre von der Gestalt. *Psych. Forch.* **4**, 301-350. Abridged translation: Principles of perceptual organization. In *Readings in Perception*, 1958, D.C. Beardslee & M. Wertheimer, Eds. New York: D. Van Nostrand.

Wilson, H.R. & Bergen, J.R. 1979 A four mechanism model for spatial vision. *Vision Research* **19**, 19-32.

Wilson, H.R. & Giese, S.C. 1977 Threshold visibility of frequency gradient patterns. *Vision Research* **17**, 1177-1190.

Wright, M.J. 1982 Contrast sensitivity and adaptation as a function of grating length. *Vision Research* **22**, 139-149.

Zucker, S.W. 1983 Computational and psychophysical experiments in grouping: early orientation selection. In *Human and machine vision*, Beck, J., Hope, B., and Rosenfeld, A., eds. 545-567.

Ullman, S. 1979 *The interpretation of visual motion*. Cambridge, MA: MIT Press.

Ullman, S. 1984 Visual Routines. *Cognition* **18**, 97-159.

Vassilev, A. & Penchev, A. 1976 Spatial and temporal summation in the perception of lines. *Vision Research* **16**, 1329-1335.

Wertheimer, M. 1923 Untersuchungen zur Lehre von der Gestalt. *Psych. Forch.* **4**, 301-350. Abridged translation: Principles of perceptual organization. In *Readings in Perception*, 1958, D.C. Beardslee & M. Wertheimer, Eds. New York: D. Van Nostrand.

Wilson, H.R. & Bergen, J.R. 1979 A four mechanism model for spatial vision. *Vision Research* **19**, 19-32.

Wilson, H.R. & Giese, S.C. 1977 Threshold visibility of frequency gradient patterns. *Vision Research* **17**, 1177-1190.

Wright, M.J. 1982 Contrast sensitivity and adaptation as a function of grating length. *Vision Research* **22**, 139-149.

Zucker, S.W. 1983 Computational and psychophysical experiments in grouping: early orientation selection. In *Human and machine vision*, Beck, J., Hope, B., and Rosenfeld, A., eds. 545-567.