Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research 45 (2005) 2511–2522
Vision Research
www.elsevier.com/locate/visres

Grouping local orientation and direction signals to extract spatial contours: Empirical tests of "association field" models of contour integration

Timothy Ledgeway [a,*], Robert F. Hess [b], Wilson S. Geisler [c]

[a] School of Psychology, University of Nottingham, University Park, Nottingham, Nottinghamshire NG7 2RD, UK
[b] McGill Vision Research, Department of Ophthalmology, McGill University, Montreal, Que., Canada H3A 1A1
[c] Center for Perceptual Systems and Department of Psychology, University of Texas at Austin, Austin TX 78712, USA

Received 15 September 2004; received in revised form 24 February 2005

Abstract

Over the last decade or so a great deal of psychophysical research has attempted to delineate the principles by which local orientations and motions are combined across space to facilitate the detection of simple spatial contours. This has led to the development of "association field" models of contour detection which suggest that the strength of linking between neighbouring elements in an image, is determined by the degree to which they aligned along smooth (first-order) curves. To test this assumption we used a path detection paradigm to compare the ability of observers to identify the presence of contours defined by either spatial orientation, motion direction or by specific combinations of both types of visual attribute. The relative alignment of the local orientations and/or directions with respect to the axis of the depicted contour was systematically varied. For orientation-defined contours detection was best when the elements were aligned along (parallel with) the contour axis, approached chance levels for obliquely oriented elements and then improved for elements that were orthogonal to the contour axis (i.e., performance was a U-shaped function of degree of orientation misalignment). This pattern of results was found for both straight and curved contours and is not readily explicable in terms of current association field theories. For motion-defined contours, however, performance simply deteriorated as the relative directions of the constituent path elements were progressively misaligned with respect to the contour. Thus the rules by which local orientations are linked to define spatial contours are qualitatively different from those used for linking local directions and each may be mediated by distinct visual mechanisms. When both orientation and motion cues were simultaneously available, contour detection performance was generally enhanced, in a manner that is consistent with probability summation. We suggest that association field models of orientation linking may need to be extended in light of the present findings.

Keywords: Orientation; Direction; Spatial-linking; Contours

1. Introduction

It is well established that the perceived overall (global) structure of a scene depends heavily on the configuration and properties of its constituent local visual features. The Gestalt laws of perceptual organisation attempted to formalise this important aspect of vision by specifying the rules by which local proximity, similarity, continuity and common motion determine object perception (Wertheimer, 1923). However these laws are largely descriptive in nature (based on subjective experience), have little predictive power and lack an explicit model or mechanism of the visual processes involved.

* Corresponding author. Tel.: +44 (0) 115 8467343; fax: +44 (0) 115 9515324.
E-mail addresses: timothy.ledgeway@nottingham.ac.uk, txl@psychology.nottingham.ac.uk (T. Ledgeway).

0042-6989/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2005.04.002

Recent research has attempted to address these deficiencies and has provided important insights into the underlying visual circuitry that could mediate a diverse range of perceptual grouping phenomena. These phenomena include context-dependent modulation of contrast sensitivity, pre-attentive "pop-out", and the extraction of smooth contours and motion trajectories from noisy backgrounds (e.g., see Bex, Simmers, & Dakin, 2003; Hess & Field, 1999; Kovacs, 1996; Morgan & Solomon, 2000; Polat & Sagi, 1993; Verghese, McKee, & Grzywacz, 2000).

Over the last decade or so psychophysical studies in particular have investigated how local orientation signals are combined across space to define simple spatial contours. For example Field, Hayes, and Hess (1993) measured the ability of observers to detect elongated spatial contours (paths) that were constructed from oriented, bandpass elements (Gabor patches) embedded in an array of similar but randomly oriented background elements. They reported that detection performance was best for straight paths and fell gradually as the degree of curvature of the depicted contours increased. They interpreted their results in terms of an "association field" which responds to contours in the image by integrating (linking) information across neighbouring local filters or receptive fields tuned to similar orientations (Fig. 1A). The strength of linking between adjacent filters is greatest for those that have the same orientation preference and are co-linear in visual space. This model is economical in that an interactive network composed of a relatively small ensemble of local detectors could, in principle, encode a wide range of contours. Physiological studies indicate that in striate cortex intrinsic, long range, horizontal connections link cells with similar preferences but non-overlapping receptive fields (e.g., Das & Gilbert, 1995; Gilbert & Wiesel, 1979, 1983, 1989; Malach, Amir, Harel, & Grinvald, 1993; Rockland & Lund, 1982, 1983; Schmidt, Goebel, Lowel, & Singer, 1997). This intra-cortical circuitry could enable V1 to integrate information over relatively large regions of visual space and could potentially be an anatomical substrate for an association field mechanism.

A number of recent studies have also examined the structure of the visual environment (in particular "natural" images) to investigate the presence of consistent statistical properties that could be utilised by contour-integration mechanisms (Elder & Goldberg, 2002; Kruger, 1998; Sigman, Cecchi, Gilbert, & Magnasco, 2001). These studies have measured the co-occurrence (joint) statistics of oriented filter responses in natural images and have established that there are indeed statistical correlations (involving both geometric and luminance-based regularities) between pairs of oriented edge elements within scenes, suggestive of the Gestalt cues of proximity, continuity and similarity. Furthermore research by Geisler, Perry, Super, and Gallogly (2001) provides ecological validity for the concept of an association field for contour extraction in human vision. They measured edge co-occurrence statistics in photographs representative of naturalistic scenes. That is how likely pairs of local edge elements belong to the same contour as a function of separation distance, orientation difference and polar direction. The most striking feature of this analysis was that edge elements that are consistent with a smooth, continuous curve are more likely to belong to the same physical contour, especially for elements that that are nearby and co-linear. Moreover human performance for detecting naturalistic contours embedded in noisy backgrounds was closely coupled to the likelihood of edge co-occurrence in the natural world. This provides evidence that human vision may exploit these statistical regularities of the visual environment and is consistent with the general linking principles embodied in the association field model of contour integration.

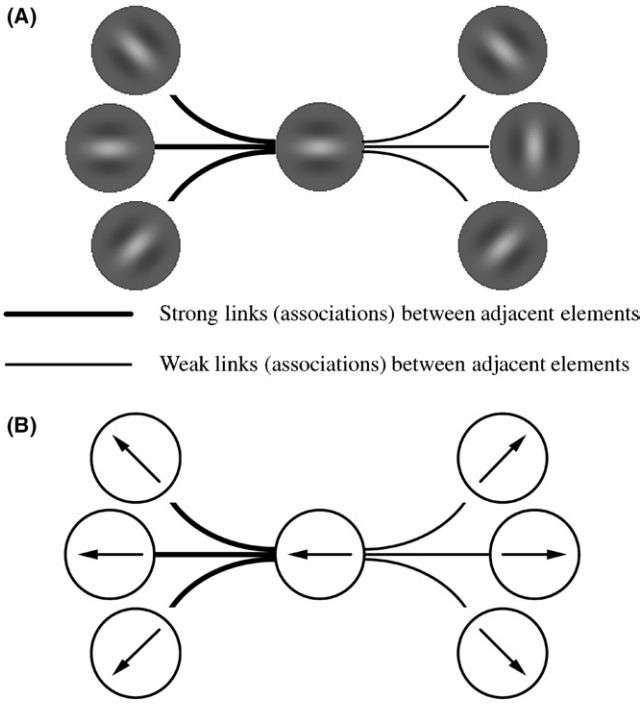

Fig. 1. (A) Schematic illustration of the association field model of Field et al. (1993) which detects elongated spatial contours in the image by grouping or linking information across adjacent local filters (cf. receptive fields) tuned to similar orientations. The diagram shows a central reference element from which the association field radiates in all directions. Strength of linking (association) is depicted by the thickness of the lines emanating from the reference element. The thick lines on the left of the diagram represent strong links between edge elements that conform to smooth, first-order curves. Those on the right of the diagram represent weak links. (B) Illustration of the specific rules of association for local direction signals (represented by the arrows) that define simple spatial contours (based on the results of Ledgeway & Hess, 2002). Strength of linking (association) is also depicted by the thickness of the lines emanating from the central reference element.

However some results are difficult to interpret with the framework of the association field model. For example the model predicts that neighbouring elements in a path will only be linked or grouped together if their relative orientations and positions are both consistent with a smooth (first-order) curve. Consequently path detection performance should deteriorate when the orientations of the elements are misaligned with respect to the path axis and no longer conform to a smooth contour. In agreement with this prediction Field et al. (1993) found that when the path elements comprising a contour were rotated so that they were each randomly oriented ±30° relative to the path axis, and hence were no longer aligned, performance deteriorated and was close to chance levels. However when the degree of orientation misalignment was increased further so that the elements were oriented 90° (orthogonal) with respect to the path, detection was surprisingly good (80–90% correct for straight paths). Subsequent studies (e.g., Bex, Simmers, & Dakin, 2001; Hess, Ledgeway, & Dakin, 2000) have also noted that elements orthogonal to a contour can give rise to above chance performance on contour-integration tasks. This interesting aspect of performance is not readily explicable in terms of the original association field model and warrants further systematic study. It is possible that the relatively good task performance found with orthogonal (with respect to the path) elements simply represents a special case of contour integration or some qualitatively different segmentation process distinct from that encapsulated by the association field.

A number of recent studies have extended the path detection paradigm to investigate the rules by which other local visual cues such as motion direction can be integrated across space to disambiguate the presence of elongated contours (e.g., Hess & Ledgeway, 2003; Ledgeway & Hess, 2002). For example Bex et al. (2001) measured the ability of observers to detect moderately curved (20° curvature) paths composed of Gabor elements in which the local orientation information was either aligned with or orthogonal to the path's axis. In addition although the elements themselves were stationary their internal sinusoidal luminance structure could be static, moving or flickering. Paths composed of drifting or flickering elements were significantly more detectable than those containing static spatial form and performance was best when the element orientations were aligned along the path (so that motion direction was perpendicular to the depicted contour) rather than orthogonal to it. Although the results clearly demonstrate that local motion (and flicker) cues can augment contour extraction, the elements always contained explicit local orientation cues so it was not possible to isolate the role that motion plays in contour detection.

To address directly the issue of spatial linking of motion cues Ledgeway and Hess (2002) required observers to detect the presence of simple motion-defined paths embedded in a field of otherwise random local motions. There were no explicit local orientation cues and contours were defined solely on the basis of the local direction signals present along the path's length. This was achieved by using elements composed of spatially two-dimensional (2-d) random noise that could be made to drift in any desired direction (within the confines of a stationary Gaussian aperture). When the motion directions of the elements were aligned along the contour in a consistent manner detection performance was good for straight and moderately curved paths (~90% correct) but fell to chance for curvatures greater than about 60°. Thus there are rules for combining local motion directions across space that, to a first approximation, appear to be analogous to those that determine how local orientation signals are combined to encode elongated contours (Field et al., 1993). In other words there is an association of motion direction and spatial location that is consistent with smooth, first-order curves (Fig. 1B).

However the precise rules by which local direction signals are grouped across space to encode extended spatial contours have not been well studied and the degree to which they share similarities with those proposed for orientation linking has yet to be firmly established. To address this important issue more fully and establish the principles by which local visual cues such as orientation and motion direction are grouped across space, a fruitful approach is to investigate how contour detection depends on the relative alignment of each of these cues with respect to the contour's axis. In the present study we therefore compared the ability of observers to identify the presence of contours defined by either spatial orientation, motion direction or by specific combinations of both types of visual attribute. This will enable us to: (1) Investigate further the relationship between detection performance and the relative orientation of edge elements with respect to the contour axis on which they lie. (2) Refine the association field model of orientation linking so that it accurately reflects the perceptual abilities of human observers. (3) Establish the extent to which the mechanisms that encode motion-defined contours utilise similar principles to those used for orientation-defined contours. (4) To examine further how orientation and direction signals interact to govern contour detection when both cues are simultaneously present.

2. Methods

2.1. Observers

Two of the authors (TL and RFH) served as observers in the experiments and both had corrected-to-normal acuity with no history of ocular ill health.

2.2. Apparatus and stimuli

Stimuli were computer generated and displayed on a *Sony Trinitron Multiscan E500* monitor (frame refresh rate of 75 Hz) that was carefully gamma-corrected with the aid of internal look up tables. To obtain precise control of luminance contrast the number of intensity levels available was increased from 8 to 12 bits using a video attenuator device (Pelli & Zhang, 1991). All stimuli were viewed binocularly at a distance of 0.74 m and presented within a square window at the center of the display that subtended 16.9° both horizontally and vertically. At this viewing distance each screen pixel subtended $1.6 \times 1.6'$. The remainder of the display area was homogenous and had a luminance of $\sim$56 cd m^{-2}.

Each stimulus was generated anew immediately prior to its presentation (on any one trial) and consisted of a dense spatial array of 158, non-overlapping micropattern elements, similar to those used in previous studies of contour (path) detection (e.g., Field et al., 1993; Hess & Ledgeway, 2003; Ledgeway & Hess, 2002). Each element was composed of a patch of random noise texture (Michelson contrast 0.5) presented within the confines of a stationary, 2-d, Gaussian spatial window (standard deviation 0.13°, truncated at ±0.4°). The mean spatial separation between the centers of adjacent elements was 1.6° (± a uniform random deviate of 0.53°).

Three experiments were conducted and the spatiotemporal properties of the noise within each micropattern were tailored individually for each experiment (illustrated schematically in Figs. 2 and 3): in Experiment 1, to assess contour integration based purely on the spatial linking of local orientations, each micropattern contained an independent sample of static, one-dimensional (1-d) oriented noise. The width of the individual noise "bars" was 3.2′. Each micropattern could be assigned any arbitrary spatial orientation (spanning the 180° range) independently of the orientations assigned to other elements within the display. In Experiment 2, to measure contour integration based solely on the linking of motion signals, each micropattern contained spatially 2-d noise that could be made to drift smoothly at a speed of 4° s^{-1}, in any desired direction spanning the 360° range. The individual noise checks subtended $3.2 \times 3.2'$. The use of 2-d, rather than 1-d, noise in this experiment ensured that there were no explicit local orientation cues present. In Experiment 3, to investigate how orientation and direction signals govern contour detection when both cues are simultaneously present, micropatterns contained spatially 1-d noise that could be made to move in a direction perpendicular to its orientation.

2.3. Procedure

A contour detection task directly analogous to that used by Field et al. (1993) was employed. Using a two-alternative-forced-choice (2AFC) procedure observers were required to choose which of two temporal intervals (separated by 1 s) contained an elongated spatial contour (path). One interval, chosen at random on each trial (duration 507 ms), contained 158 micropatterns of random position (background micropatterns). In the other interval (path plus background) some (8) of the micropatterns were constrained to lie along the invisible backbone of an elongated contour that was constrained to pass through a central circular region of the display area of radius 0.8°. There were no density differences between the two stimulus intervals and both contained exactly the same number and type of micropattern elements.

Performance was measured for contours of varying curvature, expressed conventionally in terms of the path angle (either 0°, 20° or 40° plus a small amount of random jitter of ±5°). A path angle of 0° indicates a straight path and a path angle of 40°, for example, indicates a highly curved path. In Experiment 1 (orientation-defined contours) a uniform orientation difference or misalignment could be imposed between the constituent edge elements (1-d noise) and the local curvature of the depicted path. Seven relative orientation differences were used ranging from 0° (orientations perfectly aligned with and parallel to the contour axis) to 90° (orientations orthogonal to the contour) in equal steps of 15°. The magnitude and sign (chosen randomly on each presentation to be either clockwise or counter-clockwise relative to the contour axis) of the orientation difference were identical for all micropatterns along the contour's length. In a similar manner in Experiment 2 (motion-defined contours) a uniform difference could be imposed between the directions of local motion of the 2-d noise and the curvature of the depicted path. A set of seven relative directions, ranging from 0° (motion in a consistent direction always along the contour's axis) to 90° (directions orthogonal to the contour) were tested. In Experiment 3 (contours defined by orientation and motion cues) both the relative orientation and direction of motion were co-varied in seven equal steps of 15°. In this case it is important to note that when the local orientations of the drifting 1-d noise were aligned with the depicted contour (0° relative orientation), the motion directions were necessarily constrained to be orthogonal to its axis (90° relative direction) and vice versa.

Observers completed at least 2 runs of 100 trials for each condition of the three experiments and testing of different runs was carried out in a pseudo-random order. The overall percentage correct detection and the associated standard error (SE) were calculated separately for each observer and condition (see Appendix A.1 for the calculation of the SE of a percentage score).

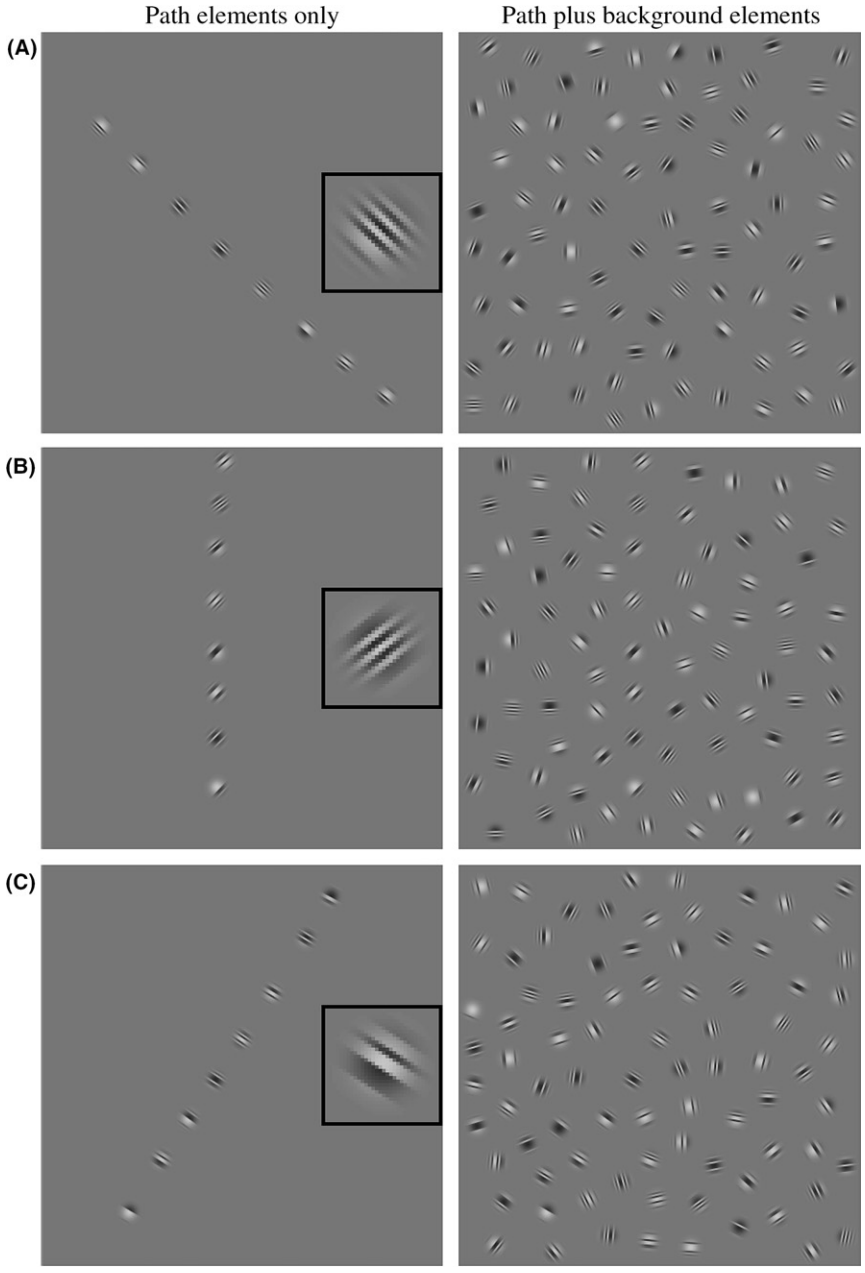

Fig. 2. Illustration of a number of orientation-defined straight contours (paths) similar to those used in Experiment 1 (and 3). The three images on the left show only the micropatterns belonging to the contour and the insets shows a magnified view of one of the elements. The three images on the right show the same contours embedded in background elements having random orientations. (A) Path elements are oriented so that they are aligned and parallel with the axis of the depicted contour (0° relative orientation). (B) Path elements are oriented obliquely with respect to the contour (45° relative orientation). (C) Path element orientations are orthogonal to the contour axis (90° relative orientation). In Experiment 1 the 1-d, oriented noise was static and in Experiment 3 the noise could be made to drift in a direction perpendicular to its orientation.

3. Results

3.1. Experiment 1: The detection of orientation-defined contours

In Fig. 4 the ability of the two observers to detect orientation-defined contours is plotted as a function of the relative orientation of the constituent micropattern elements, for path angles (curvatures) of 0°, 20° and 40° (depicted by the different symbols). The observers exhibit similar patterns of performance and two main features of the results are readily apparent. First, detection performance is best for straight paths (0° path angle) and progressively deteriorates as the degree of contour curvature increases. Second, and more importantly, it is clear that for a given path angle performance is a non-monotonic function of the relative orientation of the micropattern elements comprising the contour.

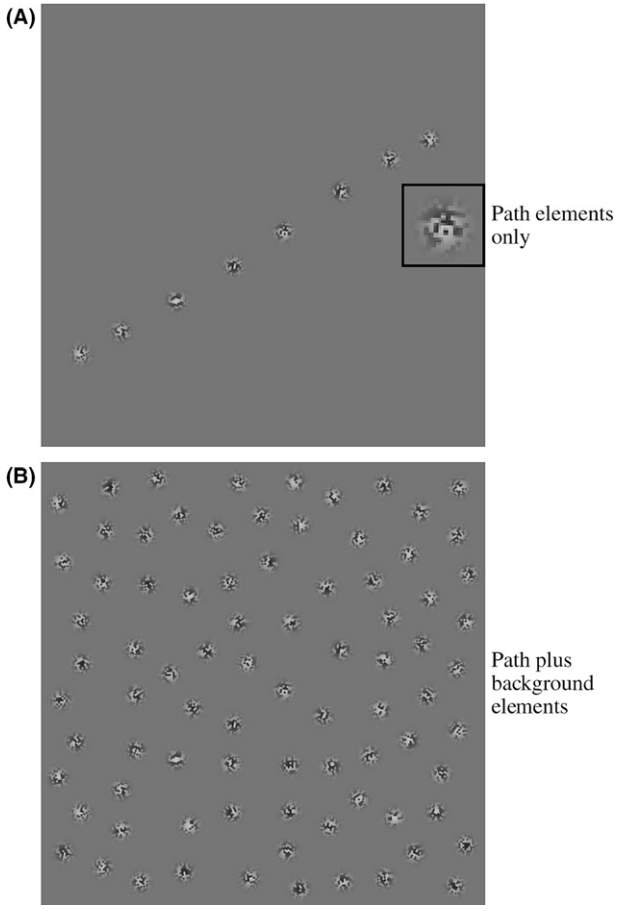

Fig. 3. Illustration of the basic spatial layout of a motion-defined straight contour (path) similar to those used in Experiment 2. (A) The path micropattern elements in isolation together with an inset showing a magnified view of one of the elements. (B) Path embedded in background elements. In the actual experiment the 2-d noise within each micropattern could be made to drift in any direction spanning the 360° range. See text for further details.

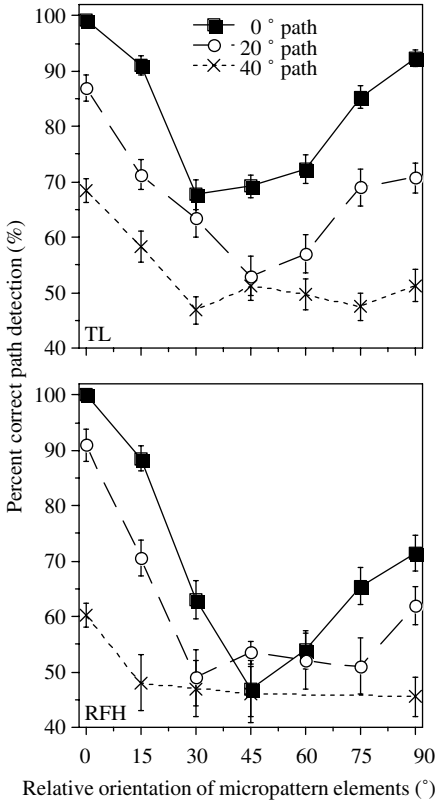

Fig. 4. Percent correct performance for detecting static orientation-defined contours is plotted against the relative orientation of its constituent elements for three different contour curvatures (specified as path angle and depicted by the different symbols and line styles). Results are plotted separately for observers TL and RFH. Each data point is based on at least 2 runs of 100 trials and the vertical bars above and below each value represent ±1 SE.

That is the probability of correct detection is a U-shaped function of the relative orientation of the path elements. For example for straight paths performance is close to 100% correct for both observers when the relative orientation is 0° (orientations parallel to the contour axis), gradually deteriorates (especially for RFH where performance is at chance) as the relative orientation approaches 45° (orientations oblique with respect to the contour axis) and then gradually improves again until the relative orientation reaches 90° (orientations orthogonal to the contour). However performance for the 90° relative orientation condition is never as good (70–90% correct) as that for the 0° relative orientation condition. This general pattern of results is also apparent for the curved paths but overall performance is typically worse, especially for the 40° curved paths where detection only exceeds chance levels (50% correct) when the relative orientation is less than about 15°. Thus the relative orientation of the elements defining a simple spatial contour has a profound influence on the ability to extract that contour in manner that is inconsistent with the association field model of Field et al. (1993).

3.2. Experiment 2: The detection of motion-defined contours

Fig. 5 shows path detection performance for the two observers plotted as a function of the relative direction of the micropattern elements, for path angles (curvatures) of 0°, 20° and 40° (depicted by the different symbols). The results for the two observers are in good agreement and demonstrate a number of important characteristics. Varying the relative direction of the micropattern elements has a marked impact on the efficacy of detecting motion-defined contours but the effects are qualitatively different from those found using orientated-defined contours (in Experiment 1) when relative orientation was varied. For motion-defined contours the probability of correct detection tends to simply decrease (in a monotonic manner) as the relative direction of the micropattern elements representing that contour increases. For example for straight paths

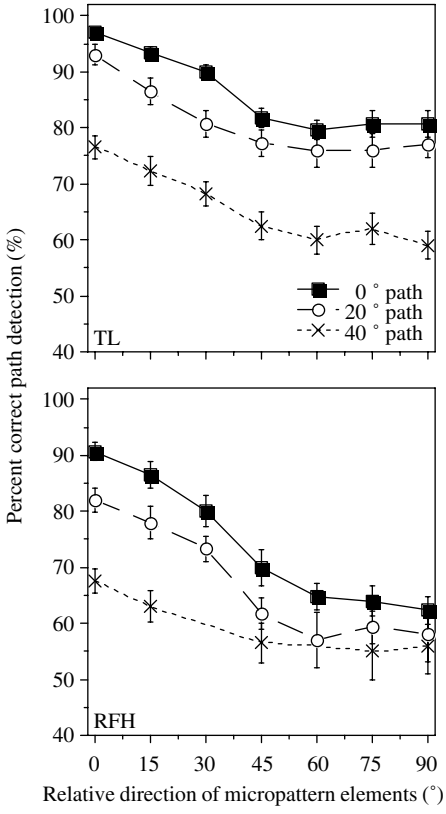

Fig. 5. Percent correct performance for detecting motion-defined contours is plotted against the relative direction of its constituent elements for three different contour curvatures (specified as path angle and depicted by the different symbols and line styles). Results are plotted separately for observers TL and RFH. Each data point is based on at least 2 runs of 100 trials and the vertical bars above and below each value represent ±1 SE.

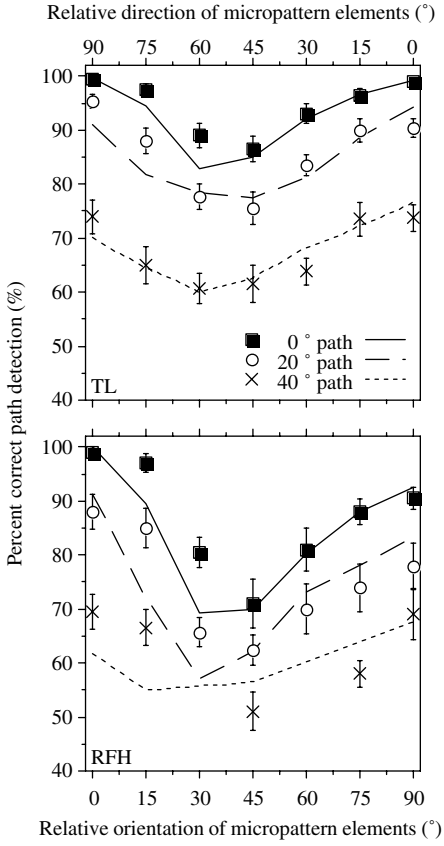

Fig. 6. Percent correct performance for detecting contours defined by both orientation and motion is plotted against the relative orientation/direction of its constituent elements for three different contour curvatures (specified as path angle and depicted by the different symbols). Results are plotted separately for observers TL and RFH. Each data point is based on at least 2 runs of 100 trials and the vertical bars above and below each value represent ±1 SE. In addition the different line styles indicate the predicted performance for detecting contours containing both orientation and motion cues, based on d' summation analysis of the outputs of two independent mechanisms: one specialised for encoding orientation-defined contours and the other motion-defined contours. The detection performance measured for each condition in Experiments 1 and 2 was used to predict the performance found in Experiment 3 (shown by the symbols) for observers TL and RFH. See Appendix A.2 for further details of the probability summation model and the equations used to derive the predictions.

performance is maximal for both observers when the relative direction is 0° (directions parallel to the contour axis), declines markedly as the relative direction approaches 45° (directions oblique with respect to the contour axis) and then tends to level off, or at least fall at a much shallower rate, with further increases in relative direction. Importantly there does not appear to be any advantage for detecting paths in which the element directions are 90° (orthogonal to the contour). This general pattern of results is also shown for the curved paths (20° and 40°) but overall performance levels are uniformly worse than for the straight paths. Thus the relative alignment of direction signals along the axis of simple spatial contours also affects their detectability, but in a distinctly different manner to that found for contours defined by local orientation cues.

3.3. Experiment 3: The detection of contours defined by both orientation and motion cues

Fig. 6 shows path detection performance for the two observers plotted as a function of both the relative orientation (bottom abscissae) and relative direction (top abscissae) of the micropattern elements, for path angles of 0°, 20° and 40° (depicted by the different symbols). The patterns of results shown by the two observers are comparable and once again demonstrate that straight paths (0° path angle) are more much easily detected than curved ones. Interestingly for all path angles performance levels are considerably better overall (by ~10% on average) than those found in the previous experiments (where the contours were defined solely by the spatial linking of either orientation cues or motion cues in isolation). It is also apparent that in the present experiment although both orientation and motion cues were

simultaneously present, the pattern of results is similar to those found in Experiment 1 using orientation-defined contours. This is evidenced by the fact that detection is a U-shaped function of relative orientation/direction and exhibits a clear performance minimum (for all path angles) when the relative orientation and direction are both 45° (oblique with respect to the contour axis).

4. Discussion

In the present study we were able to directly compare the principles by which local visual cues such as orientation and motion direction are grouped across space to define simple spatial contours. By systematically investigating how contour detection depends on the relative alignment of each of these cues with respect to the contour's axis, we were able to test the adequacy of association field models that have been proposed to account for contour extraction in human vision.

4.1. The detection of orientation-defined contours and the association field model of orientation linking

For contours defined solely by static orientation cues (Experiment 1), detection performance was highly sensitive to the curvature of the depicted contours. In line with previous studies that have employed oriented Gabors (e.g., Field et al., 1993) or line segments (e.g., Geisler et al., 2001) to construct contours (rather than 1-d noise) performance deteriorated as the contour curvature increased. This aspect of the results is readily explicable in terms of conventional association field models of contour integration but other features of the results are not. Crucially the largest orientation misalignments (between the micropattern elements and the depicted contours) did not result in the worst performance. Indeed the relationship between contour detection and the relative orientation of the elements was a U-shaped function. Contours were easiest to detect when the element orientations were aligned (parallel) with its axis, marginally worse when they were orthogonal and hardest to detect when they were oriented obliquely (e.g., 45°) with respect to the contour. It is important to emphasise that this phenomenon is not simply due to an anisotropy in contour saliency, in that horizontal and vertical contours are more easily detected than oblique ones (Li & Gilbert, 2002). In the present study the contour configuration (absolute orientation and position) was randomised from trial to trial and the orientations of its constituent elements were defined with respect to the local contour axis (not in terms of visual space or screen co-ordinates).

The association field model of orientation linking (Field et al., 1993) cannot account for the pattern of results found in Experiment 1. In the original model adjacent elements were only strongly associated (linked) if they satisfied the joint constraints of position and orientation along smooth, first-order curves. To account for the present results within the same theoretical framework this assumption would need to be relaxed: Adjacent elements that are positioned along smoothly varying curves but oriented orthogonal to those curves are also strongly associated (albeit to a lesser degree than edge elements that are perfectly aligned along contours). This is illustrated in Fig. 7A which represents diagrammatically the modified rules of association between oriented elements (for clarity only elements that fall upon straight curves are shown but the same principles apply for other degrees of curvature).

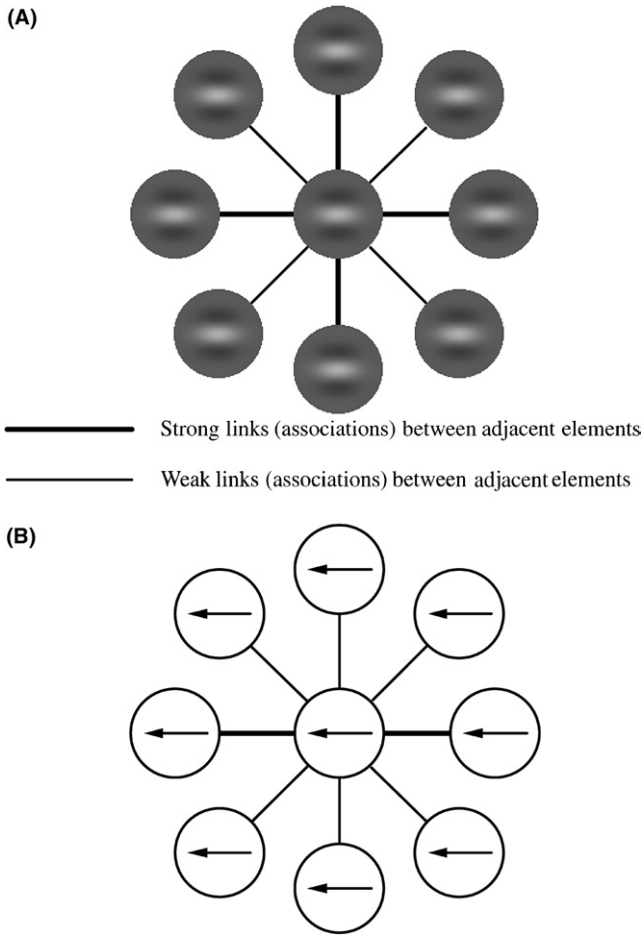

Fig. 7. (A) Schematic illustration of the revised rules of association between oriented elements based on the results of Experiment 1 (for clarity only elements that fall upon straight curves are shown but the same principles apply for other degrees of curvature). Strength of linking (association) is depicted by the thickness of the lines emanating from the central reference element. Thick lines represent strong links between edge elements thin lines represent weak links. (B) Illustration of the rules of association for local direction signals (represented by the arrows) based on the results of Experiment 2 (again for clarity only elements that fall upon straight curves are shown). Associative strength is depicted by the thickness of the lines emanating from the central reference element.

It is important, however, to distinguish between the diagram of association strengths depicted in Fig. 7A and the underlying mechanism(s) by which this spatial linking or grouping is achieved. It is possible, for example, that the relatively good performance found when elements are orthogonal (relative orientation 90°) to the contour they represent reflects the operation of a different mechanism to the case when the element orientations are aligned (relative orientation 0°) with the contour axis. The former could be mediated by (say) a general image segmentation process that serves to spatially group together proximal elements with similar absolute orientations, without regard to either their spatial configuration or orientation alignment with respect to a contour. Such a segmentation mechanism would necessarily be distinct from the specialised contour extraction mechanism originally encapsulated by the association field model. However there are a number of potential problems with this suggestion: First, a segmentation mechanism based on linking similar local orientations should be able to encode a contiguous ensemble of elements that have oblique (e.g., 45°) orientations (relative to a contour) as readily as those that have orthogonal, or even aligned, orientations which is clearly not the case. This suggests that it is not simply the degree of orientation similarity between adjacent path elements that governs performance but rather the alignment of those orientations along spatially elongated contours. Second, although overall performance levels are typically uniformly lower for elements that are orthogonal to a contour's axis than those that are aligned with that axis, in each case performance is maximal for straight contours and falls off with increasing contour curvature (path angle) in a qualitatively similar manner.

4.2. Alternatives to the association field model of orientation linking

It is informative to consider alternative computational models of contour extraction that are more explicit in terms of their architecture than the original association field model proposed by Field et al. (1993). In particular the cortical-based model described by Yen and Finkel explicitly includes excitatory, long range, horizontal connections both between adjacent oriented receptive fields that sample regions of space along smoothly varying curves (''co-axial'') and also those that have the same preferred orientation but are orthogonal to those curves (''trans-axial''). Excitatory connections in this model are confined to these two configurations and are absent between receptive fields that have (say) oblique (e.g., 45°) preferred orientations relative to a contour. Furthermore the trans-axial connections are more spatially focused, with the degree of excitation falling off much more rapidly with distance than for the co-axial connections. Thus the model of Yen and Finkel (1998) has much in common with the modified rules of association depicted in Fig. 7A, is broadly consistent with the pattern of psychophysical results found in the current study and can account for a range of perceptual grouping phenomena (e.g., Morgan & Solomon, 2000; Polat & Sagi, 1993). However the relatively good performance found in this and previous studies (Bex et al., 2001; Field et al., 1993; Hess et al., 2000) when elements are orthogonal (relative orientation 90°) to the contour they represent may pose difficulties for computational schemes that suggest that inhibitory, rather than excitatory, trans-axial connections exist between neighbouring receptive fields (e.g., Li, 1998; Ursino & La Cara, 2004).

4.3. Are the patterns of association strength related to statistical properties of natural images?

Presumably the pattern of association strengths in Figs. 1A and 7A, reflect the operation of mechanisms that are useful for encoding scenes in the normal environment. Thus it may be useful to consider whether there are statistical properties of natural images that are consistent with these patterns of association strength. Geisler et al. (2001) found that simple edge co-occurrence statistics in natural images were dominated by two trends, one presumably due to the smoothness of natural contours and the other to the parallel relationship between contours in natural images. The first trend is reflected in the fact that for any given distance and orientation difference between two edge elements the most likely direction of one of the elements from the other is the direction approximately consistent with co-circularity (Geisler et al., 2001; Sigman et al., 2001). Further, analysis of edge elements that have been grouped (by hand) into physically meaningful contours shows that elements that are co-circular (which include co-linearity as a special case) are more likely to belong to the same physical contour than elements that are not (Elder & Goldberg, 2002; Geisler et al., 2001). This result demonstrates the potential usefulness (ecological validity) of association strengths like that in Fig. 1A and the co-linear association strengths in Fig. 7A.

The second trend is reflected in the fact that for any given distance and direction between two elements the most likely orientation difference between the elements was zero (i.e., they were most likely to be parallel). An obvious question is whether this trend is consistent with the strong association strength in the orthogonal directions shown in Fig. 7A. The simplest analysis suggests that it is not. The co-occurrence probability for edge elements that are parallel in orientation is maximal in the co-linear direction and minimal in the orthogonal direction, rather than U-shaped. However, a more definitive analysis (which we have not yet attempted) would be to

measure the distribution of edge orientation within elongated natural objects (branches, leaves, animals, etc.) that have been hand segmented from natural images. As such ecological correlates of the psychophysical findings of Experiment 1 may not be found in the domain of contours per se, but instead may be related to texture flows (e.g., hair, fur, wood grain, surface markings, etc.) and therefore involve higher-order statistical measurements of image structure. It is entirely possible that texture-based edge orientations that are parallel and perpendicular to the long axis of natural objects are more likely to occur than those at other orientations. However such an analysis would raise the potential problem of deciding a priori what exactly constitutes a valid, "elongated natural object" in a visual scene. One interesting possibility is to select image structures perceived as salient by observers (we thank an anonymous reviewer for this suggestion).

4.4. The detection of motion-defined contours and the association field model of direction linking

Interestingly when contours were defined solely on the basis of motion cues (Experiment 2), the results were markedly different from those found using orientation-defined contours. The ability to detect motion-defined contours tended to deteriorate in a more or less uniform manner as the relative direction of the elements was progressively misaligned with the elongated axis of the contour. The most difficult contours to detect were those in which the direction of motion was oblique or orthogonal to the contour. This pattern of results is generally consistent with the properties expected of a simple association field for direction linking (e.g., Ledgeway & Hess, 2002) and clearly suggests that there are some potentially important differences between the processes that underlie the grouping of static orientation cues and direction cues. This is highlighted in Fig. 7B which, for comparison purposes, illustrates how local direction signals are preferentially linked for straight contours. In this case grouping only occurs when neighbouring motions satisfy the joint constraints of direction and spatial position that define simple smooth (motion-defined) contours.

The principal finding that motion-defined contours of moderate curvature are more easily detected when the constituent element directions are aligned along the axis of the contour, than when they are uniformly misaligned with the contour axis is in good agreement with previous research using directly comparable stimuli (Ledgeway & Hess, 2002). The issue still remains, however, why this is the case. From an ecological perspective there is at present no evidence to suggest that in the natural environment motion along the elongated axis of a contour or an object is a consistent and prevalent property of the visual world. Indeed, intuitively at least, one might expect objects in general motion to generate all directions with roughly equal probability but this has yet to be substantiated. Nonetheless the pattern of empirical results found is consistent with a number of other findings in human vision that suggest that there are distinct directional anisotropies in the ability to utilise motion cues that are dependent on the spatial configurations of stimuli. For example Chang and Julesz (1983) using random-dot-kinematograms (RDKs) demonstrated that the ability to detect the coherent motion direction of a displaced rectangular region of dots (target) was best ($D_{\max}$ values were higher) when the major axis of the target was orientated parallel, rather than orthogonal, to the direction of the movement. In a similar manner Verghese et al. (2000) reported that triplets of dots moving in a consistent direction are more detectable if that movement is along a common axis (co-linear motion) than when it is perpendicular to it (but see Bex et al., 2003). These studies among many others (e.g., Fredericksen, Verstraten, & van de Grind, 1994; Verghese, Watamaniuk, McKee, & Grzywacz, 1999; Watamaniuk, McKee, & Grzywacz, 1995) suggest a role for interactions between local motion detectors that may be specialised for extracting motion trajectories and/or predicting the future motion paths of objects.

4.5. The nature of the interactions between local orientations and local directions

To determine how orientation and direction signals interact to govern contour detection when both cues are simultaneously present we employed (in Experiment 3) micropatterns composed of drifting 1-d noise. The results are commensurate with previous studies (e.g., Bex et al., 2001) which have shown that contours containing both orientation and motion information are much more detectable than those defined only by static oriented elements.

The nature of the interactions between local orientations and local directions in contour detection is an important issue. One possibility that is not inconsistent with results of Experiments 1 and 2 is that separate visual processes mediate the extraction of contours defined by orientation signals and those defined by motion signals (i.e., each is subserved by a spatial linking mechanism with different underlying properties). Consequently when both mechanisms are simultaneously active performance could be determined by the most sensitive mechanism in each case (i.e., a winner-take-all type strategy). However the finding that overall performance levels were better when both orientation and motion were present, than when each was present in isolation, clearly argues against any simple winner-take-all approach. An alternative possibility is that probability summation determines detection performance when multiple visual cues define a spatial contour. This is

depicted in Fig. 6 in which the plotted lines show the predictions of a probability summation model based on d' (e.g., Graham, 1989; Green & Swets, 1966) when contours are defined by both orientation and motion. The percent correct scores obtained in each of the conditions of Experiments 1 and 2 were used to predict the results, shown by the symbols in Fig. 6, found in Experiment 3 (see Appendix A.2 for details of the probability summation model and the equations used to derive the predictions). As can be seen the predictions are remarkably good (given that there are no free parameters) and the value of the Pearson correlation coefficient for the predicted performance versus the observed performance is 0.98 for observer TL and 0.90 for RFH.

In conclusion the present experiments have shed further light on the principles by which local orientations and motions are grouped across space to define simple contours. This is an important first step in object recognition and scene perception. The results suggest that different rules of association apply to orientation cues and motion cues and place important constraints on models of the underlying mechanisms involved. In particular we have demonstrated that when edge elements are oriented obliquely with respect to the contour on which they lie, this can render that contour undetectable. In contrast when the same elements are simply rotated so that they are either parallel or orthogonal to the contour, its detection can be facilitated. Interestingly Bex et al. (2001) suggested that the most effective camouflage for an object may be achieved with orthogonal texture. The present results, however, suggest that surface markings that are oriented obliquely with regard to an object or animal's body may serve as a better form of camouflage.

Acknowledgement

Supported in part by NSERC grant #0GP0046528 to RFH and by NIH grant EY11747 to WSG.

Appendix A

A.1. Calculation of the standard error (SE) of percentage correct detection scores

The SE of the percentage correct detection score obtained for each condition was calculated using the following standard equation:

$$\mathrm{SE} = \sqrt{\frac{P \times (100 - P)}{(n - 1)}} \tag{1}$$

where P is percentage correct detection and n is the total number of trials completed for that condition.

A.2. Deriving the predictions of the probability summation model based on d′

Assuming an equal-variance Gaussian model of signal detection performance, then for a 2AFC testing procedure the d' values obtained when the contours were defined solely by either orientation (d'_{O}) or motion (d'_{M}) cues are given by

$$d'_{\mathrm{O}} = 2 \times \Phi^{-1}\left(\frac{P_{\mathrm{O}}}{100}\right) \tag{2}$$

$$d'_{\mathrm{M}} = 2 \times \Phi^{-1}\left(\frac{P_{\mathrm{M}}}{100}\right) \tag{3}$$

where P_{O} and P_{M} are percentage correct detection obtained when the micropattern elements in the display contained only either orientation (Experiment 1) or motion (Experiment 2) cues, respectively, and Φ^{-1} is the inverse of the cumulative Gaussian distribution function (number of standard deviations corresponding to each percentage correct score).

For probability summation of two independent signals when both are simultaneously present (i.e., the outputs of two independent mechanisms) the predicted resultant d' value (d'_{OM}) is

$$d'_{\mathrm{OM}} = \sqrt{(d'_{\mathrm{O}})^2 + (d'_{\mathrm{M}})^2} \tag{4}$$

The corresponding percentage correct detection (P_{OM}) for the composite stimulus when the contours are defined by both orientation and motion (Experiment 3) is then given by

$$P_{\mathrm{OM}} = 100 \times \Phi\left(\frac{d'_{\mathrm{OM}}}{2}\right) \tag{5}$$

References

Bex, P. J., Simmers, A. J., & Dakin, S. C. (2001). Snakes and Ladders: The role of temporal modulation in visual contour integration. *Vision Research, 41*, 3775–3782.

Bex, P. J., Simmers, A. J., & Dakin, S. C. (2003). Grouping local directional signals into moving contours. *Vision Research, 43*, 2141–2153.

Chang, J. J., & Julesz, B. (1983). Displacement limits, directional anisotropy and direction versus from discrimination in random-dot cinematograms. *Vision Research, 23*, 639–646.

Das, A., & Gilbert, C. D. (1995). Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature, 375*, 780–784.

Elder, J. H., & Goldberg, R. M. (2002). Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision, 2*, 324–353, Available from http://journalofvision.org/2/4/5/. DOI 10.1167/2.4.5.

Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local "association field". *Vision Research, 33*, 173–193.

Fredericksen, R. E., Verstraten, F. A. J., & van de Grind, W. A. (1994). Spatial summation and its interaction with the temporal

integration mechanism in human motion perception. *Vision Research, 34*, 3171–3188.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research, 41*, 711–724.

Gilbert, C. D., & Wiesel, T. N. (1979). Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature, 280*, 120–125.

Gilbert, C. D., & Wiesel, T. N. (1983). Clustered intrinsic connections in cat visual cortex. *The Journal of Neuroscience, 3*, 1116–1133.

Gilbert, C. D., & Wiesel, T. N. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *The Journal of Neuroscience, 9*, 2432–2442.

Graham, N. V. S. (1989). *Visual pattern analyzers*. New York: Oxford University Press.

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

Hess, R., & Field, D. (1999). Integration of contours: New insights. *Trends in Cognitive Sciences, 3*, 480–486.

Hess, R. F., & Ledgeway, T. (2003). The detection of direction-defined and speed-defined spatial contours: One mechanism or two? *Vision Research, 43*, 597–606.

Hess, R. F., Ledgeway, T., & Dakin, S. (2000). Impoverished second-order input to global linking in human vision. *Vision Research, 40*, 3309–3318.

Li, Z. (1998). A neural model of contour integration in the primary visual cortex. *Neural Computation, 10*, 903–940.

Kovacs, I. (1996). Gestalten of today: Early processing of visual contours and surfaces. *Behavioural Brain Research, 82*, 1–11.

Kruger, N. (1998). Collinearity and parallelism are statistically significant second-order relations of complex cell responses. *Neural Processing Letters, 8*, 117–129.

Ledgeway, T., & Hess, R. F. (2002). Rules for combining the outputs of local motion detectors to define simple contours. *Vision Research, 42*, 653–659.

Li, W., & Gilbert, C. D. (2002). Global contour saliency and local colinear interactions. *Journal of Neurophysiology, 88*, 2846–2856.

Malach, R., Amir, Y., Harel, H., & Grinvald, A. (1993). Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primary striate cortex. *Proceedings of the National Academy of Sciences of the USA, 90*, 10469–10473.

Morgan, M. J., & Solomon, J. A. (2000). Facilitation from collinear flanks is cancelled by non-collinear flanks. *Vision Research, 40*, 279–286.

Pelli, D. G., & Zhang, L. (1991). Accurate control of contrast on microcomputer displays. *Vision Research, 31*, 1337–1350.

Polat, U., & Sagi, D. (1993). Lateral interactions between spatial channels: Suppression and facilitation revealed by lateral masking experiments. *Vision Research, 33*, 993–999.

Rockland, K. S., & Lund, J. S. (1982). Widespread periodic intrinsic connections in the tree shrew visual cortex. *Science, 215*, 1532–1534.

Rockland, K. S., & Lund, J. S. (1983). Intrinsic laminar lattice connections in primate visual cortex. *Journal of Comparative Neurology, 216*, 303–318.

Schmidt, K. E., Goebel, R., Lowel, S., & Singer, W. (1997). The perceptual grouping criterion of collinearity is reflected by anisotropies of connections in the primary visual cortex. *European Journal of Neuroscience, 9*, 1083–1089.

Sigman, M., Cecchi, G. A., Gilbert, C. D., & Magnasco, M. O. (2001). On a common circle: Natural scenes and Gestalt rules. *Proceedings of the National Academy of Sciences of the USA, 98*, 1935–1940.

Ursino, M., & La Cara, G. E. (2004). A model of contextual interactions and contour detection in primary visual cortex. *Neural Networks, 17*, 719–735.

Verghese, P., McKee, S. P., & Grzywacz, N. M. (2000). Stimulus configuration determines the detectability of motion signals in noise. *Journal of the Optical Society of America A, 17*, 1525–1534.

Verghese, P., Watamaniuk, S. N. J., McKee, S. P., & Grzywacz, N. M. (1999). Local motion detectors cannot account for the detectability of an extended trajectory in noise. *Vision Research, 39*, 19–30.

Watamaniuk, S. N. J., McKee, S. P., & Grzywacz, N. M. (1995). Detecting a trajectory embedded in random direction motion noise. *Vision Research, 35*, 65–77.

Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt. *Psychologische Forschung, 4*, 301–350.

Yen, S., & Finkel, L. H. (1998). Extraction of perceptually salient contours by striate cortical networks. *Vision Research, 38*, 719–741.