Available online at www.sciencedirect.com

SCIENCE DIRECT®

Acta Psychologica 120 (2005) 35–54

acta psychologica

www.elsevier.com/locate/actpsy

Detection of change in shape and its relation to part structure

Marco Bertamini *, Tracy Farrant

Department of Psychology, University of Liverpool, Eleanor Rathbone Building, Bedford Street South, Liverpool L69 7ZA, UK

Received 17 November 2004; received in revised form 3 March 2005; accepted 4 March 2005
Available online 14 April 2005

Abstract

Using a change detection paradigm (Barenholtz, E., Cohen, E. H., Feldman, J., & Singh, M. (2003). Detection of change in shape: An advantage for concavities. *Cognition*, 89(1) 1–9), we measured sensitivity to the changes of shapes and in particular the difference between detecting a new convex or concave vertex. We conclude that concave vertices per se are not more salient, but changes in the sequence of convexities and concavities along a contour are salient. We argue that these changes are likely to signal a change in perceived part structure.

PsycINFO classification: 2323

Keywords: Visual perception; Shape; Signal detection; Object representation

1. Introduction

When curvature is measured along a contour, the sign of curvature can be used to distinguish convexities (positive curvature) from concavities (negative curvature).

* Corresponding author. Tel.: +44 151 794 2954; fax: +44 151 794 2945.
E-mail address: m.bertamini@liverpool.ac.uk (M. Bertamini).
URL: http://www.liv.ac.uk/vp/ (M. Bertamini).

0001-6918/$ - see front matter
doi:10.1016/j.actpsy.2005.03.002

Although what is positive and what is negative is arbitrary (conventionally convexities are positive), the difference between these two kinds of contours is important for how humans perceive shape. In this introduction, we will briefly review the literature, and we focus on the finding of greater sensitivity to concavities coming from visual search and change detection studies (Barenholtz, Cohen, Feldman, & Singh, 2003; Hulleman, te Winkel, & Boselie, 2000; Humphreys & Müller, 2000). Next, we will show in two new experiments that sensitivity is higher not for concave vertices per se, but for either convex or concave vertices depending on the context in which they are found. We argue that the best explanation for this context effect is that there is an early computation of part structure based on how convexities and concavities follow each other along a contour. We describe this argument in more detail in the introduction of Experiment 2.

The interest in contour curvature can be traced back as far as Alhazen (11th century).[1] More recently, Attneave (1954) has discussed the importance of curvature extrema for how shape is perceived. Perhaps the most important analyses, however, are those in Koenderink (1984), and Hoffman and Richards (1984) because they have shown how the sign of curvature of a contour relates to the Gaussian curvature of surfaces (see also Richards, Koenderink, & Hoffman, 1987). This link suggests that the visual system may exploit 2D curvature information to recover information about solid shape. In addition, peaks of negative curvature (concave cusps) in the image may also signal the presence of a transversal intersection between solid shapes. On this ground, Hoffman and Richards (1984) have proposed that observers should parse a shape into subparts at concave cusps (minima rule).

Much empirical evidence seems to support the importance of contour curvature for perception of part structure (for a review see, Singh & Hoffman, 2001). For instance, Bertamini (2001, see also Bertamini & Croucher, 2003, Bertamini & Mosca, 2004) has found that observers are faster at judging position of vertices perceived as convex. Bertamini (2001) has suggested that vertices that are perceived as parts have a position in a way that boundaries between parts (concave vertices) do not. In other words, there is an *explicit* (in the sense of Marr, 1982) representation of position for parts. For theoretical reasons as to why convexities should be seen as parts, see also Koenderink (1990), and, for a different approach, Leyton (1992).

Using a visual search task, Hulleman et al. (2000) and Humphreys and Müller (2000) found an asymmetry: Search was more efficient for targets with concavities among convex distractors than vice versa. As acknowledged in the papers, there

[1] Alhazen (Arabic name: Abu-'Ali Al-Hasan Ibn Al-Haytham, c. 965–1039) was born in Basra (now Iraq). Many of his books have survived and the most famous is a seven volume work on optics, which was translated into Latin in the 12th century and influenced the work of Roger Bacon and Kepler. Alhazen relied on experiment rather than on past authority, and among other things he argued that objects are seen by reflected light, and not by light emanating from the eye (the then popular extramission theory), and gave the first description of a camera obscura. He is often referred to as the father of modern optics. With respect to contour curvature, Norman, Phillips, and Ross (2001) cite this passage from Alhazen's Optics: "for sight will perceive the figure of the surfaces of objects whose parts have different positions by perceiving the convexity, concavity or flatness of those parts, and by perceiving their protuberance or depression".

are two possible interpretations of this effect. (i) Perhaps concavities are more salient and the visual system is particularly tuned to them; Humphreys and Müller go as far as hypothesising concavity "detectors". They argue that because of its importance to part structure, attention is drawn to concavities. In their recent review of attributes that guide the deployment of attention, Wolfe and Horowitz (2004) cite curvature as a likely attribute, with a possible preference for concavities. (ii) Alternatively, a closed region with a concavity is more complex than a strictly convex closed region in terms of perceived part structure. An extra part may have been perceived because of the concavity, but this is different than claiming that concavities per se are more salient than convexities. Bertamini and Lawson (2004a) have conducted a series of visual search experiments pitting strictly convex targets against strictly concave targets (circular holes), and found no evidence that concave contours per se are more salient. Note that circular figures and circular holes are congruent unlike the stimuli used by Hulleman et al. (2000) and Humphreys and Müller (2000), and therefore allow a more symmetrical design (Rosenholtz, 2001).

The paradigm used in this paper was introduced by Barenholtz et al. (2003). They presented irregular polygons to observers; after a mask the polygons appeared again but in half of the trials one vertex had been added or removed. The task was simply to detect the change; sensitivity was higher when a concave vertex was added or removed. Barenholtz et al. (2003) draw conclusions similar to those of Hulleman et al. (2000) and Humphreys and Müller (2000). Note that as for the visual search data, the authors acknowledge that there are two possible explanations for the findings. Was sensitivity higher to concave vertices per se, or was detection easier when the concave vertex affected perceived part structure? One may ask whether it is possible to separate the two, given that a new concavity is assumed to split the perceived shape into subparts. In this paper, we find a way around this problem, by comparing detection of concave and convex vertices in different contexts, and we conclude that concave vertices are not always more salient.

Our Experiment 1 is similar to the experiment by Barenholtz et al. (2003), but in half the trials the closed polygons are perceived as holes. For holes, what was a concavity is now a convexity and vice versa, therefore if there is higher sensitivity for concave vertices that would apply to vertices shifted *inwards* for the figure but *outwards* for the hole, as illustrated in Fig. 1. This experiment confirmed an advantage for concave vertices for figures, but not for holes. We discuss the limitations shared by our Experiment 1 and Barenholtz et al.'s experiment, and in Experiment 2 we modify the design to show more conclusively that it is not concave vertices per se that are more salient.

2. Experiment 1: Detection of change in figures and in holes

This experiment used the task introduced by Barenholtz et al. (2003). Observers viewed two shapes in succession, either identical or one a slightly altered version of the other, and reported whether they detected a change. The shapes were irregular polygons and the only possible change was the introduction of a new vertex (as

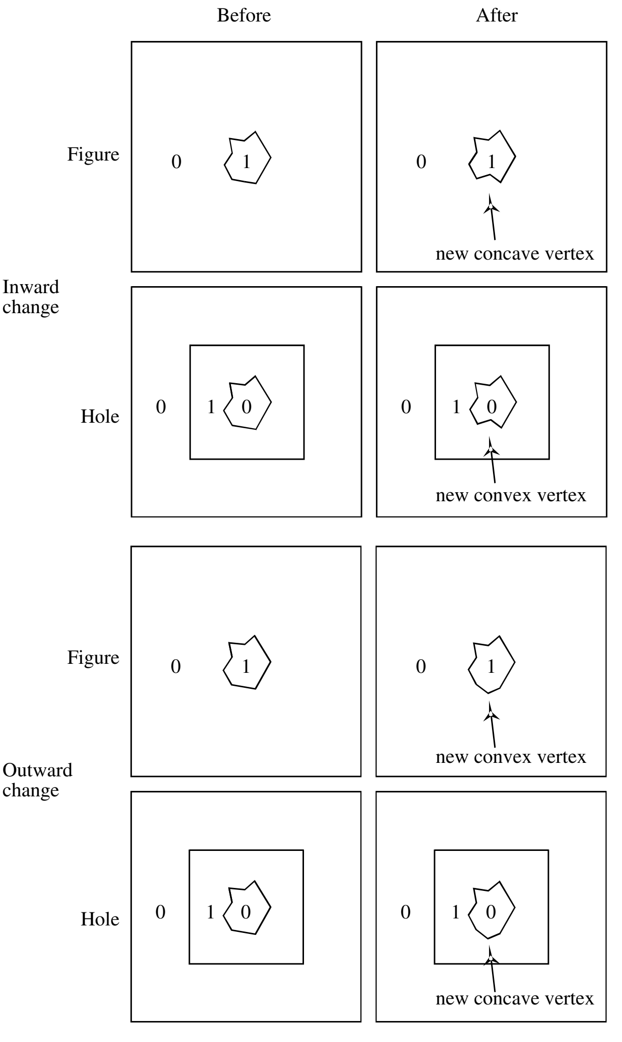

Fig. 1. Design of Experiment 1. The cells show the combination of Direction of change (inwards or outwards) and type of Object (figures or holes). The actual stimuli were random dot stereograms and the labels 0 and 1 are used here to indicate zero disparity and crossed disparity for the different regions. A new polygon was generated on each trial starting from locations on an ellipse, and the algorithm is described in more detail in the text.

opposed to the removal of a vertex, also possible in Barenholtz et al., 2003). In half of the trials, the new vertex was located inwards with respect to the straight edge that it replaced (towards the centre of the display) and in the other half it was located outwards (Fig. 1). For a polygon perceived as a figure the inward and outward conditions can be labelled concave and convex, respectively. The novelty of our experiment is that in half of the trials the polygon was perceived as a hole within a larger square. In the case of the hole, unlike the figure, the inward vertex is a convex vertex and the outward vertex is a concave vertex (Fig. 1). Therefore, if concavities are more salient we predict that performance will be higher for the inward condition in the case of figures, and the outward condition in the case of holes. Alternatively, if inward changes are always more salient we predict higher performance for inward changes for figures and holes.

There has been considerable interest in visual holes recently, and there is some disagreement about how the shape of holes is perceived. Some memory tasks reveal good recollection for the shape of a hole, even when the shape is later presented as an object (i.e., its complement) (Palmer, 1999). However, when the task requires a fast response to a property of the contour, holes seem to behave as if the contour belongs to the surrounding object (Bertamini & Croucher, 2003; Fantoni, Bertamini, & Gerbino, 2005). Although these different findings may create some confusion, we believe that they are not incompatible. A closed contour may be remembered well even when there is a figure-ground reversal (from study to test phase). This finding does not imply that the contour is shared by figure and ground, or that unidirectional border ownership does not apply to holes (Bertamini, 2005).

To make sure that there was no ambiguity about depth order, stimuli were generated as random dot stereograms (RDSs). We believe RDSs are particularly useful to compare convexity and concavity because they allow figure-ground reversal for a closed region without changing the contours themselves (Bertamini & Mosca, 2004; Bertamini & Lawson, 2004a, 2004b). In other words, the figures and holes that we are comparing are congruent. In addition, observers can select a response only after they have fused the two images, ensuring that the specified depth order was perceived.

2.1. Method

Participants. Twelve students of the University of Liverpool participated. They were screened for stereoacuity using the TNO stereotest. Acuity ranged between 15″ and 60″.

Stimuli and procedure. Stimuli were generated on a Macintosh computer, and presented on a Sony F500T9 monitor with a resolution of 1280×1024 pixels at 120 Hz. Two stereo images were presented with the use of a NuVision infrared emitter and stereoscopic glasses. The effect of interleaving left and right images was that effective vertical resolution and refresh rate were halved (512 pixels at 60 Hz).

In each trial, a new RDS was generated. There was a square background 9 cm wide, and an irregular polygon varying in size and complexity. When the polygon was presented as a hole, this hole belonged to a square region 4.5 cm wide. Stimuli and procedure are illustrated in Figs. 1, and 2. Note that to mask the first presentation, we used a new random dot background during the retention interval. The disparity between foreground and background was 101″.

To describe the stimulus, let us take the polygon perceived as a figure first. The polygon in the first interval had 5, 6, 7 or 8 vertices, of which either 1 or 2 were concave. The convex vertices were placed along an invisible ellipse whose axes were chosen to be between 35 and 48 pixels long. Therefore, the size and elongation of the

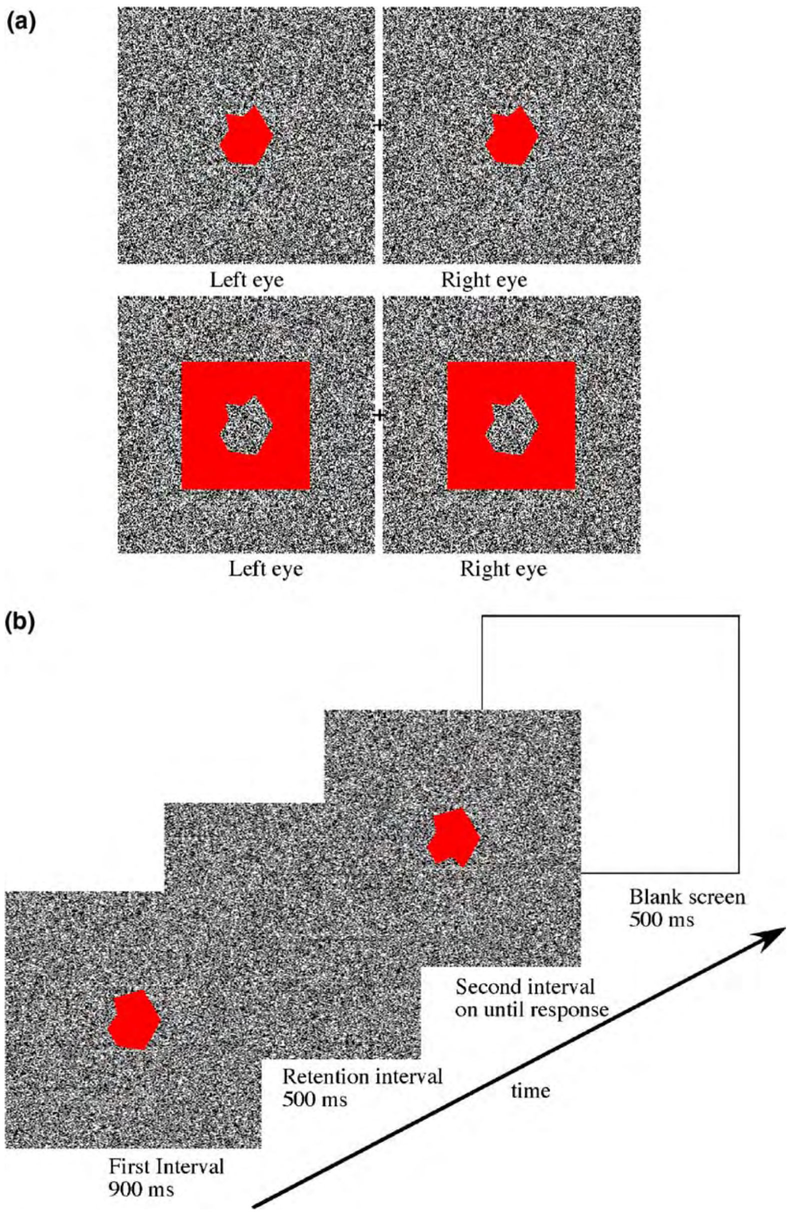

Fig. 2. (a) Two stereograms showing the same polygon presented as a figure and as a hole. In this illustration the figure is solid red, this type of stimuli was used in the first practice session to help observers become familiar with the task. In the second practice session and in the experiment, the contours were defined exclusively by binocular disparity. (b) Observers saw one stimulus in the first interval, and had to judge whether the shape had changed or had remained the same in the second interval. A change trial is illustrated. The random dot background during the retention interval is different from that in the first and second intervals, thus providing an effective mask.

polygon varied from trial to trial. When the shape changed from the first interval to the second, a new vertex was added splitting in half a straight edge between two convex vertices. Half of the time this new vertex was inside the edge (concave), and the other half it was outside the edge (convex). The displacement of this new vertex from the contour was chosen randomly to be either 10 or 12 pixels.

The logic of this algorithm is similar to that used by Barenholtz et al. (2003), although our polygons are simpler. One difference is that we only added vertices in the second interval, instead of adding or removing them. This was done because Barenholtz et al. (2003) found no difference between these two conditions. More importantly, there was a new condition in which the polygon was perceived as a hole. When the polygon was perceived as a hole, the terms convex and concave in the previous paragraph need to be swapped. Therefore, the hole stimulus had 5, 6, 7 or 8 vertices, of which either 1 or 2 were convex. Inward and outward changes are now convex and concave, respectively.

Each observer sat in a dimly illuminated room at a distance of approximately 57 cm from the monitor. Observers were instructed to respond with the right or left hand, to indicate that the shapes in the two intervals were the same or different. An acoustic feedback informed the participants when they responded incorrectly.

The practice was divided into two halves of 18 trials. In the first half, the stimuli (the polygon or the square with polygonal hole) were presented not as random dot surfaces but as solid red surfaces to make sure observers knew what type of shapes to look for. In the second half, the stimuli were RDSs. Throughout the practice when the participant made an incorrect response, after a warning beep, the stimuli were presented again but without a mask between first and second interval so that the change (or absence of a change) was clearly visible.

When the experiment proper started, each observer performed 384 trials in rapid succession. After every 64 trial, a block ended and the observer was allowed time to rest. The start of subsequent blocks was self-paced. The computer recorded response times and controlled the presentation of the stimuli using the VideoToolbox subroutines (Pelli, 1997).

Design. The factors were: Change from first to second interval (same or different shapes), Direction of change (inward or outward), type of Object (figures or holes) and Displacement (10 or 12 pixels). They were factorially combined in a within-subjects design.

2.2. Results and discussion

The following steps were taken in the analyses for this and Experiment 2. For each subject, we computed hit rates and false alarm rates for each condition. From these we computed d' values. Next we carried out a repeated measures analysis of variance (ANOVA) on d' values. The factors were: Direction of change (inwards or outwards), type of Object (figures or holes), and Displacement (10 or 12 pixels). Mean values are plotted in Fig. 3a.

There was no effect of Displacement, and a significant main effect of Object ($F(1,11) = 60.48, p < 0.001$), indicating that performance was higher for figures than

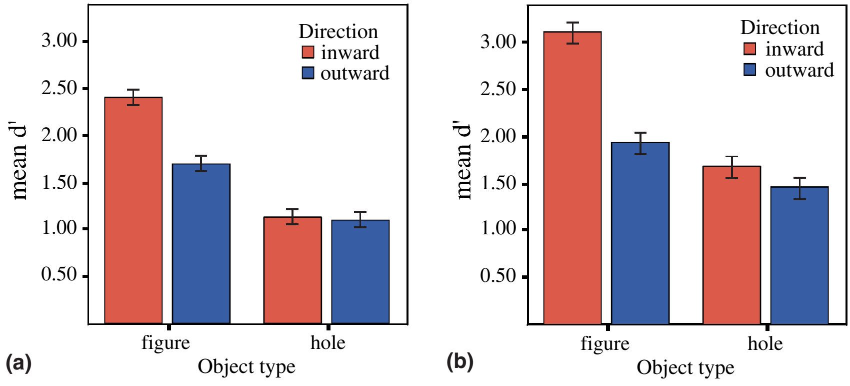

Fig. 3. Data from Experiment 1. (a) Mean d' values plotted for new vertices that were located inward and outward with respect to the centre of the polygon. For figures, inward vertices and outward vertices are concave and convex, respectively, for holes the opposite is true. (b) The sample was split into two groups based on performance. The bar graph for the group with higher performance shows the same pattern as the whole sample. Error bars are within-subjects standard errors.

for holes. There was a marginal effect of Direction ($F(1,11) = 6.75, p < 0.025$), indicating higher performance for the inside condition, and an interaction between Direction and Object ($F(1,11) = 11.99, p < 0.005$). This interaction is of critical importance: Is it the case that performance is higher for the inward condition for figures but outward condition for holes? As a direct test of the difference between these two means, we compared inward and outward conditions for holes using a paired t test. We found no significant difference ($t(11) = 0.27$). For figures the difference was significant ($t(11) = 3.27, p = 0.007$).

The results show that sensitivity was higher in the inward condition. For polygons perceived as figures this difference was significant and it is a replication of the finding in Barenholtz et al. (2003). The critical test is the case of holes. Even though the higher sensitivity for convexities was not significant, its direction is the opposite of what would be predicted by a concavity advantage. Because performance was lower for holes one may worry about a floor effect. Therefore, we split subjects into two groups based on performance for holes, and in Fig. 3b we plotted the mean d' values for the group with higher performance ($N = 6$). The means for inward and outward conditions were different for figures ($t(5) = 3.98, p = 0.011$) but not for holes ($t(5) = 1.42$). The graph shows the same pattern as for the whole sample. Moreover, sensitivity in our experiment was in general higher than in Barenholtz et al. (2003), making a floor effect unlikely.

At this point it is important not to throw the baby out with the bathwater. Perhaps Barenholtz et al. (2003) were correct in their conclusion about the fact that concavities are more salient but there were problems with their paradigm.

For a closed polygon the total turning angle of convex vertices must be higher than that of concave vertices (Feldman & Singh, 2005), and typically there are more convex than concave vertices (for instance, a closed polygon can have only convex vertices but it cannot have only concave vertices). In Barenholtz et al. (2003) as well as in Experiment 1, we always used figures (polygons) with more convex than concave vertices. This in itself may explain why the addition of an extra concave vertex in the figure was more easily detected than the addition of an extra convex vertex. To make this clearer, note that going from 5 to 6 is a 20% change but going from 2 to 3 is a 50% change. For holes, a similar argument applies but the roles are reversed: there are now more concave vertices than convex, therefore the addition of a convex vertex may be more salient, counterbalancing the effect of concavity. Whatever the outcome of Experiment 1, one would be left with this confound and therefore unable to draw firm conclusions.

The logical question, therefore, is how to avoid this problem. An experiment using polygons with an equal number of convex and concave vertices is difficult to design because the total turning angle cannot be the same. Polygons of this type would look either very regular (like a star) or with sides of very unequal length and distribution (e.g., when the concave vertices are inside a loop). Both of these solutions are problematic because a change from regular to irregular would be easy to spot, and because unequal length and distribution would create an inbalance in the very variable we are testing. A better solution is to isolate a section of an object, seen through an aperture. Within the visible section of the object, convex and concave vertices can be matched in number, size, and shape. The contours of the aperture itself would not be important because these contours do not belong to the object itself. In other words, we rely here on the principle of unidirectional border ownership (Nakayama, Shimojo, & Silverman, 1989). For a similar logic, see Bertamini and Lawson (2004b). In Experiment 2, we adopt such strategy, but before we explain the new stimuli we need to develop our predictions in more detail.

2.3. Concavities, convexities and part structure

Experiment 1 replicated Barenholtz et al.'s (2003) findings with respect to figures, but did not support the claim that sensitivity to concavities is higher in general (because of the different finding for holes). However, the stimuli of Experiment 1 suffer the same problem of the original study, as discussed above. We believe that a new type of classification of the stimuli is necessary.

As explained in Section 1, there are at least two possible explanations for what Barenholtz et al. (2003) found (and the same applies to Hulleman et al., 2000, and Humphreys & Müller, 2000). Either concave vertices per se are more salient or they are only more salient when they affect part structure. The problem with the part structure hypotheses is that the minima rule is not sufficient to determine the parsing. Several solutions to the general problem of parsing have been proposed: Hoffman and Singh (1997), Rosin (2000), Singh, Seyranian, and Hoffman (1999), Siddiqi, Tresness, and Kimia (1996); and Siddiqi, Kimia, Tannenbaum, and Zucker (2001). The predictions from these models overlap to some degree and more empirical evaluation is necessary (the most comprehensive empirical test so far has been carried out by de Winter & Wagemans, 2001). We will not discuss these models in detail because they deal with complex closed shapes, in which the symmetric axis plays a much greater role. In this paper, we are concentrating on more local information.

Barenholtz et al. (2003) follow the literature in saying that locally a new minimum (m−) will affect part structure. If so, minima will always be salient, independently from the overall part decomposition. The examples in Fig. 1 and all the stimuli in our experiments use simple polygons. Therefore, to translate the terminology from vertices to extrema, it seems reasonable to suggest that convex and concave vertices are positive maxima (M+) and negative minima (m−), respectively (as they would in a smoothed version of our stimuli). In addition, a straight segment between two convex vertices (M+) is equivalent to a positive minimum (m+), and a straight segment between two concave vertices (m−) is equivalent to a negative maximum (M−). Again, this would be the case in a smoothed version of our stimuli.

In Barenholtz et al. (2003) and in our Experiment 1, the contrast is between concave vertices (m−) and convex vertices (M+). It is true that M+ singularities do not play the same segmentation role as m− singularities.[2] However, we believe that the introduction of a new m− or M+ can affect part structure depending on the context. Locally, that is, when a complete partition of the object is not the issue, it is unlikely that an m− introduced between other two m− points (replacing an M−) will lead to the perception of a new part. Conversely, an M+ introduced between two m− points (replacing an M−) will introduce a new part. What we suggest is that minima (m−) determine part structure when they bracket a convexity, not simply because they are negative minima. Fig. 4 illustrates the new conditions necessary for the more complete design adopted in Experiment 2. For the purpose of this analysis let us assume that these are regions of larger surfaces. In Experiment 2 we isolate such segments by means of an aperture, and therefore avoid the difficult issue of how to achieve a complete partition of a complex closed shape.

The first two rows (A and B) illustrate the case of a convex/concave vertex added between two convex vertices. These are the only two cases used in Barenholtz et al. (2003). In the next two rows (C and D), convex/concave vertices are added between two concave vertices. This is the case of the holes in Experiment 1. We believe that what is salient does not depend on whether the vertex is convex or concave but whether it leads to a change in perceived part structure. When a concave vertex is added between convex vertices, this concavity splits a region previously perceived as a single part into two parts. However, when a concave vertex is placed between concave vertices this is not the case. Importantly, and perhaps counter-intuitively, when a convex vertex is added between two concave vertices this is again a change in perceived part structure: a new part has now appeared that was not present before. As both concave vertices are there both before and after, in themselves they cannot be responsible for the perceived change in part structure. For completeness we also show (E and F) the case of new vertices inserted between a convex vertex on one side

[2] Richards and Hoffman (1985) have suggested five primitive elemental shapes, called *codons*. The names, 0+, 0−, 1+, 1−, and 2, refer to the number of inflections inside the codon segment, and the sign refers to the convex or concave versions when two exist. The codon representation has many important properties, which cannot all be listed here, but it relies specifically not on the sign of contour curvature but on curvature extrema. The boundaries between codons are positive minima (m+) and negative minima (m−) of curvature.

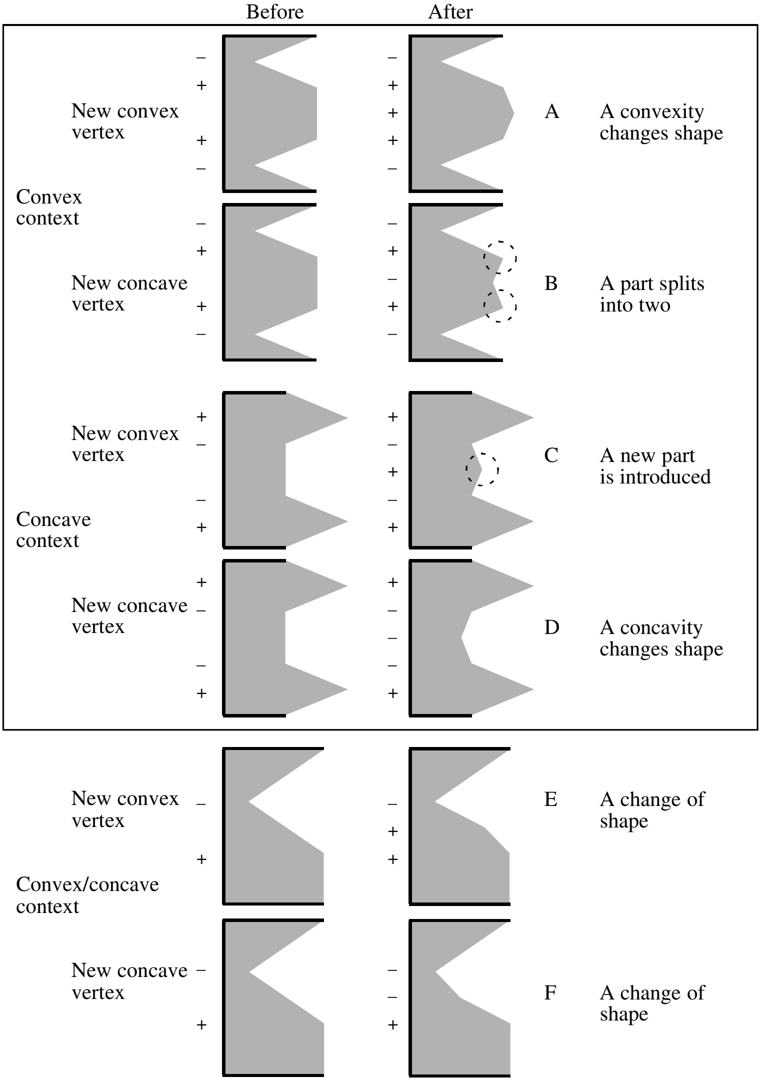

Fig. 4. The effect of adding a new concave and convex vertex to a basic shape. Vertices are added within different contexts. There are three types of contexts: a pair of convex vertices, a pair of concave vertices, or a pair formed by one convex and one concave vertices. This third type is presented for completeness but was not studied in Experiment 2, and is therefore outside the box. Within each context either a convex or a concave vertex can be added. The new convex vertex is added between two convex vertices (A), or between two concave vertices (C); the new concave vertex is added between two convex vertices (B), or between two concave vertices (D). On the right, next to the labels, we have noted what we believe is the likely interpretation in terms of part structure. This classification is based on the idea that when a convexity is bracketed by two concavities it is segmented as a visual part.

and a concave vertex on the other. In neither of these two cases there is a change of perceived part structure.

In all these conditions, there is also a change of codon structure (Richards & Hoffman, 1985). The change in terms of new parts is therefore not the same as a change in codon coding, unless one makes the assumption that codon 2 is special and more important than the others.

In deciding whether a change of perceived part structure was present or not in Fig. 4 we have followed a simple rule. Part segmentation does not always happen at negative minima, rather segmentation happens when convexities and concavities follow each other. Part structure is based on the string of changes in curvature sign along one contour, not just the string of extrema. We call this criterion the *bracketing* hypothesis.

We do not claim originality for the bracketing hypothesis. Indeed, the idea is clearly present in the literature. Take this quote: "natural parts ... lie between concave cusps" (Richards & Hoffman, 1985 p. 266). Even more explicit is the point made by Koenderink (1990) about convex regions being typically parts (p. 251). However, in the very same paper, Richards and Hoffman (1985) state the following rule: "Segment a curve at concave cusps ... in order to break the shape into its parts" (p. 266). Note that the cusps (m−) are the critical points along a contour without reference to their context. This is confirmed in the examples, because the authors predict that an elliptical hole will split the shape at the two negative minima points. For holes they provide an ad hoc solution, claiming that such holes would not be perceived as holes. This is problematic because it works well for pictorial stimuli, but is unclear what would happen to unambiguous holes (for more on holes see Bertamini, 2005). More importantly, for a local analysis like that in Fig. 4, whether the contour segment belongs to a hole or not is unspecified. Should we expect the rule of segmenting a curve at concave cusps (m−) to hold? If so, the prediction is different from the prediction based on the bracketing hypothesis. Specifically, in Fig. 4 should a new concave vertex in D lead to a change in perceived part structure?

Most observers will agree by simply looking at Fig. 4 that cases B and C are more salient, but they may argue that in A the change does not *add* a convexity, unlike C, and that in D the change does not *add* a concavity, unlike B. This is true, what we are comparing is the salience of a new m− and M+, not concavities and convexities in general. In other words, we use vertices to test the role of extrema and not just the sign of contour curvature. However, note that, disregarding figure-ground differences, all changes are matched in terms of how much the contour changes orientation and in terms of area of the new region. More importantly, it is also possible to directly compare B with C, and A with D. The former is a comparison between a new convexity and a new concavity. The latter is a comparison between a new maximum within a region with positive curvature and a new minimum within a region with negative curvature. If concavities are more salient per se, performance should be higher in B than in C, and perhaps also in D than in A since it is a change within a more salient region. We predict that this will not be the case.

Another interesting comparison is between the two contexts, because one context has positive curvature (A + B) and the other negative curvature (C + D). If attention is directed toward the context with negative curvature one would predict a main effect of context, with higher sensitivity for changes taking place within the context with negative curvature. Again, we predict that this will not be the case because we do not believe that concavities attract attention simply by virtue of being concavities.

3. Experiment 2: Concavities in context

We are now ready to introduce Experiment 2. The stimuli are similar to those in Experiment 1, except that we never present closed polygons. Instead we present only one portion of a larger polygon, occluded by a frame on three sides. This allows us to concentrate our attention (and the attention of the observers) on one unbound contour. There are four conditions, corresponding to cases A–D of Fig. 4. Fig. 5 shows the shape of the stimuli. Note that the basic shape has a fixed number of vertices, and that the figural side of the contour can be either to the left or the right. By orienting the surfaces in the vertical dimension we avoid a possible interference from the known ''lower region'' cue (Vecera, Vogel, & Woodman, 2002).

We predict higher d' values for conditions where the sequence of convexities and concavities changes (B and C) because we believe this to affect part structure based on the bracketing hypothesis. The idea that minima are more salient in themselves would predict higher d' values for conditions B and D. Because we will analyse the type of vertex (convex, A + C, and concave, B + D) as a factor crossed with the factor context (between convex vertices, A + B, or between concave vertices, C + D), we predict a significant interaction.

3.1. Method

Participants. Twelve students of the University of Liverpool participated. They were screened for stereoacuity using the TNO stereotest. Acuity ranged between 15″ and 120″.

Stimuli and procedure. Equipment and procedure were similar to those of Experiment 1, but the stimuli were different. The new stimuli are based on the logic of Fig. 4 and are illustrated in Fig. 5. There was a basic shape with seven straight edges; the contour turned either to the left and then to the right or vice versa. Independent of whether it turned left or right first, the figural side could be either to the left or to the right. Note that, unlike Experiment 1, there are now always two surfaces, one in front of the other (we refer to the one in front as figure) and they are seen through a square aperture. This means that only a portion of the figure as well as a portion of the ground are visible and that the shape analysis must be more local because the global shape of the figure is unspecified.

If one were to number the segments, the straight edge of numbers 3 and 5 were always the location of the change: one straight segment was replaced by a pair of segments with either a new convex (outward) or a new concave (inward) vertex (see Fig. 5). Observers were instructed about the possible changes and about their location. Moreover, as in Experiment 1, there was a practice session. The practice was divided into two halves of 18 trials. In the first half, the figure (with crossed disparity and

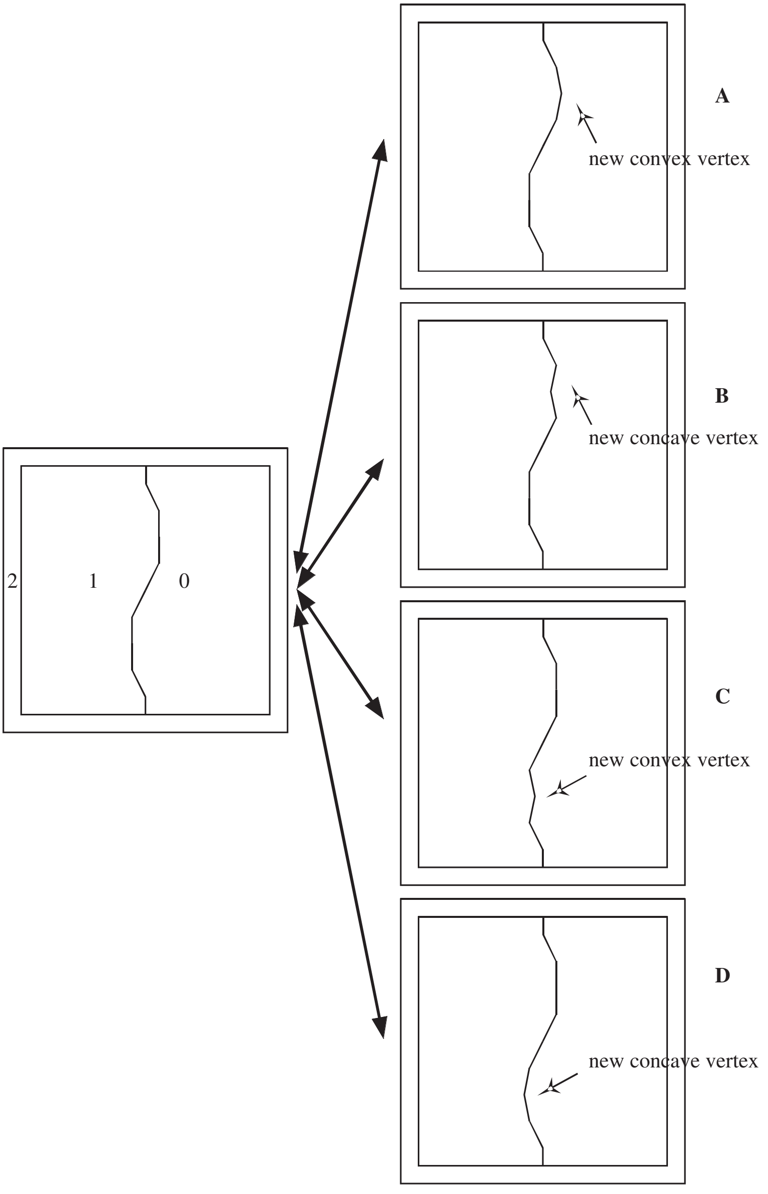

Fig. 5. Illustration of stimuli used in Experiment 2. In the experiment, the stimuli were random dot stereograms. The image on the left shows the basic shape, with seven straight edges, and numbers 0, 1, and 2 to indicate increasing levels of disparity. In this example, the contour turns right first and then left but the opposite was equally frequent. Moreover, the surface in front (1) could appear either to the left or to the right with equal frequency. The right column shows four different types of changes. These changes could appear in the second interval or disappear in the second interval as indicated by the bidirectional arrows. The labels on the right are the same as those used in Fig. 4, which should be consulted for a more detailed description of the design.

therefore perceived in front) was presented not as a random dot surface but as a solid red surface to make sure observers knew what type of shapes to look for. In the second half of the practice, the stimuli were RDSs. Throughout the practice when the participant made an incorrect response, after a warning beep, the stimuli were presented again but without the mask between first and second interval so that the change (or absence of a change) was clearly visible.

For trials in which there was no change, the same set of shapes from the change trials was presented. Moreover, in half of the no change trials the first interval had the basic shape. This ensured that whatever shape was visible in the first interval provided no information as to whether this was a change or a no change trial.

When the experiment proper started, each observer performed 384 trials in rapid succession. After every 64 trials, a block ended and the observer was allowed time to rest. The start of subsequent blocks was self-paced.

Design. The factors were: Change from first to second interval (same or different shapes), Appearance (change introduced in the second interval, appearance trial, or removed in the second interval, disappearance trial), Type of change (convex or concave), and Context (between convex or concave vertices). They were factorially combined in a within-subjects design.

3.2. Results and discussion

As for Experiment 1, for each subject we computed hit rates and false alarm rates for each condition. From these we computed d' values. Next we carried out a repeated measures analysis of variance (ANOVA) on d' values. The factors were: Appearance (whether the change appeared or disappeared in the second interval), Type (convex or concave), and Context (between convex or concave vertices). Mean values are plotted in Fig. 6. There were no significant main effects for Appearance and Type, and a marginal main effect of Context ($F(1,11) = 7.47, p < 0.019$), suggesting that sensitivity was slightly higher for changes between convex vertices (A + B). As predicted, there was a significant interaction between Type and Context ($F(1,11) = 13.62, p < 0.004$).

As can be seen in Fig. 6, the pattern of the interaction between Type and Context is the same for both appearance and disappearance conditions. The lack of a significant effect of Appearance is interesting as it suggests that, other things being equal, the introduction of a new vertex is not more salient than the removal of a vertex or vice versa.

In the graph of Fig. 6, we have also written the four labels used in Fig. 4: A–D. As predicted, efficient change detection for a convex or concave vertex depends upon the context, and more specifically performance is higher when there is a change of curvature sign, which we take to lead to a perceived change in part structure. As discussed in Section 1, it is also interesting to compare directly B with C, and A with D. Paired t tests (Bonferroni adjusted) did not confirm any significant difference ($t(11) = 1.97$, and $t(11) = 0.33$, respectively). What this means is that there is neither an advantage for convex nor for concave vertices when convexity and concavity are compared within the same context.

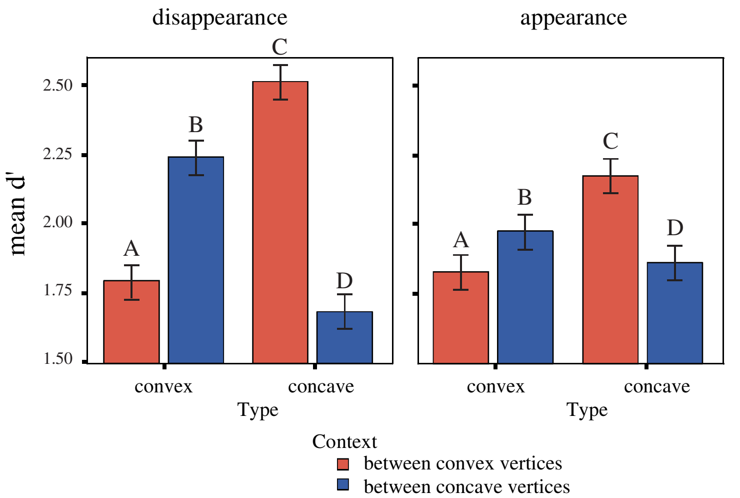

Fig. 6. Data from Experiment 2. Mean d' values plotted for new convex and concave vertices that were placed within two different contexts: between convex vertices or between concave vertices. In half of the trials, the change appeared in the second interval and in the other half it appeared in the first interval and the two bar graphs show results for disappearance and appearance. Error bars are within-subjects standard errors.

Finally, although there was no time pressure on participants to respond, the computer also recorded response time. When we plotted mean response time we found a pattern entirely consistent with the d' graph of Fig. 6.

4. Conclusion

Experiments 1 and 2 show consistent evidence that sensitivity to change of shape is not explained only by a concavity advantage. We list here all the various pieces of evidence. (a) In Experiment 1 we could not confirm an advantage for detection of new concave vertices when they were located on the contour of a hole. In Experiment 2 (b) we found no overall difference between detection changes involving convex or concave vertices, either when the vertex was added or deleted. Also in Experiment 2 (c) we found a small advantage for detection of changes taking place within a convex context, a finding we would like to study further. Finally, we found (d) that detection was good for a change in the sequence of convexities and concavities, whether this was due to the addition (or deletion) of a convex or a concave vertex.

These results contradict the proposal that concavities are inherently more salient, as in some cases new convex vertices become more salient than new concave vertices depending on the context (Experiment 2). Our stimuli are superior to those used by Barenholtz et al. (2003). In random dot stereograms the figure-ground relations are unambiguous and therefore convexity and concavity coding is also unambiguous, but more importantly in Experiment 2 we use an aperture technique that allows a more symmetric design (i.e., same number of convex and concave vertices in the basic shape). We interpret our results in terms of perceived part structure: sensitivity was higher when the change of sign along the contour affected perceived part structure. This is consistent with the existing evidence that part structure is perceived early (e.g. Xu & Singh, 2002). There is also evidence of good detection of changes of part structure for three dimensional objects (Keane, Hayward, & Burke, 2003), although in that study the changes did not alter the total number of parts.

Although we interpret our findings in terms of perceived parts, alternatively one may claim that changes of sign are salient in themselves or because they imply a change of direction of the contour (as opposed to a continuation in the same general direction). We see these are different wordings for a basic effect related to the sign of curvature, but we believe that the part structure interpretation is the one that provides the best explanation for why changes of sign are salient.

To classify our stimuli we have used the bracketing hypothesis (Fig. 4). This is not a radical departure from the previous literature on early computation of part structure, and is a relation of the *minima rule*, but it is nevertheless useful for a few reasons that we list here.

(a) One purpose of the bracketing hypothesis is to make more explicit how concavities affect perceived parts. This is an alternative to the idea that concavities per se are more salient (Barenholtz et al., 2003; Humphreys & Müller, 2000). In other words, one possibility is that because concavities affect part structure they are always more salient, and the other is that concavities are more salient only when they affect part structure. We believe the evidence is strongly in favour of the second (see Bertamini & Lawson, 2004a, 2004b).

(b) The bracketing hypothesis is weaker than the minima rule because it implies that the sequence of extrema within a contour segment with negative curvature are not important. That is, negative minima (m−) and maxima (M−) can follow each other without a significant role to play in determining part structure. The literature currently does a poor job in separating what empirical evidence supports the role of concavities as such and what goes further and supports the minima rule (a more demanding task). Elegant analyses based on differential geometry show how important minima (m−) *should* be for the visual system (Hoffman & Richards, 1984; Richards et al., 1987). Yet, we know of little empirical evidence to suggest a qualitative difference in how concavities are treated when (a) a negative minimum is present, (b) two minima are present with a negative maximum in between, or (c) a concavity with constant curvature (no minimum) is present. For instance, when Hulleman et al. (2000) tested this difference empirically they failed to find any support for the minima rule (whilst their data are consistent with the bracketing hypothesis).

Perhaps some may find the distinction between concavities and concave cusps (minima) unnecessary. After all if the amount of curvature is important, a cusp is simply the place where the concavity is most salient. However, as mentioned above, the emphasis on the extrema in the literature leads to problems in special cases, like when negative minima and maxima alternate. Confusingly, evidence that the sign of curvature affects perceived part structure has often been described as supporting the minima rule (e.g., Baylis & Driver, 1995), whilst the critical role of minima would require a more direct test. In our Experiment 2, we have shown that salience of the same type of extrema depends greatly on the context.

It is also interesting that negative maxima (M−) are not very common in natural images. de Winter and Wagemans (2001) have analysed line drawings of everyday objects. Of a total of 3788 singularities only 87 (2.2%) were M− compared to 879 (23.2%) for m−. It seems reasonable that although the system should treat m– as extremely important extrema, it may not be worth the extra effort if all concavities are similar once their salience is weighted by the amount of curvature.

The approach in this paper and the literature discussed relies on contour information to partition a visual shape. This approach has been contrasted with the approach that relies on a fixed set of volume primitives (see Singh & Hoffman, 2001, for a review). However, even starting from contour information some authors have stressed the importance of the resulting parts. For instance, Hoffman and Singh's part salience factors (1997) consist of the size of the part relative to the whole object, the degree of protrusion, and the strength of the part boundaries. This is relevant here because a part obtained by segmenting at minima that do not bracket a convexity is likely to fail such criteria (no protrusion). For a different proposal that also relies on the convexity of the resulting part, see Rosin (2000).

(c) The special situations where negative minima (m−) are present with a maximum (M−) in between are interesting because there are no convexities to provide the *parts* along that same contour. Recent models of part segmentation do seem to deal well with these situations because part cuttings are required to cross a medial axis (Siddiqi et al., 1996; Singh et al., 1999). However, let us take again the example of a hole. Based on the minima rule an elliptical figure with a concentric elliptical hole should split into two parts. This prediction seems to follow also from the short-cut rule (Singh et al., 1999) because all three criteria are met: the part cut is a straight line, it crosses a local axis of symmetry, and it involves at least one minimum point. This part structure may or may not be perceived (we have no data to resolve the issue), but the bracketing hypothesis makes the prediction that it should not because on each continuous contour there is never a change of sign.

In summary, we present evidence that changes of contour curvature that change curvature sign are detected more efficiently than equivalent changes not leading to a change of sign. In line with the literature we believe that curvature sign is important because it affects perceived part structure. We define a change of part structure on the basis of the bracketing hypothesis: convexities next to concavities are perceived as parts. This rule is simple and its motivation comes from the work by Hoffman and Richards (1984), Richards and Hoffman (1985), and Koenderink (1984) and can be seen as a relaxation of the minima rule, but it is novel in that it makes some new predictions, some of which need to be tested in future research.

References

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review, 61*(3), 183–193.

Barenholtz, E., Cohen, E. H., Feldman, J., & Singh, M. (2003). Detection of change in shape: An advantage for concavities. *Cognition, 89*(1), 1–9.

Baylis, G. C., & Driver, J. (1995). Obligatory edge assignment in vision: The role of figure and part segmentation in symmetry detection. *Journal of Experimental Psychology: Human Perception and Performance, 21*(6), 1323–1342.

Bertamini, M. (2001). The importance of being convex: An advantage for convexity when judging position. *Perception, 30*, 1295–1310.

Bertamini, M. (2005). Who owns the contour of a hole? (submitted for publication).

Bertamini, M., & Croucher, C. J. (2003). The shape of holes. *Cognition, 87*(1), 33–54.

Bertamini, M., & Mosca, F. (2004). Early computation of contour curvature and part structure: Evidence from holes. *Perception, 33*, 35–48.

Bertamini, M., & Lawson, R. (2004a). Visual search for a figure among holes and for a hole among figures. (submitted for publication).

Bertamini, M., & Lawson, R. (2004b). Testing inbuilt assumptions about convexity using a depth discrimination task. (submitted for publication).

de Winter, J., & Wagemans, J. (2001). *Segmentation of object outlines into parts: From a large-scale study to a qualitative model* (Psychological Reports No. 277). University of Leuven.

Fantoni, C., Bertamini, M., & Gerbino, W. (2005). Contour curvature polarity and visual interpolation. *Vision Research, 45*, 1047–1062.

Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review, 112*, 243–252.

Hoffman, D. D., & Richards, W. (1984). Parts of recognition. *Cognition, 18*, 65–96.

Hoffman, D. D., & Singh, M. (1997). Salience of visual parts. *Cognition, 63*, 29–78.

Hulleman, J., te Winkel, W., & Boselie, F. (2000). Concavities as basic features in visual search: Evidence from search asymmetries. *Perception and Psychophysics, 62*(1), 162–174.

Humphreys, G., & Müller, H. (2000). A search asymmetry reversed by figure-ground assignment. *Psychological Science, 11*(3), 196–201.

Keane, S. K., Hayward, W. G., & Burke, D. (2003). Detection of three types of changes to novel objects. *Visual Cognition, 10*(1), 101–127.

Koenderink, J. (1984). What does the occluding contour tell us about solid shape? *Perception, 13*, 321–330.

Koenderink, J. (1990). *Solid shape*. Cambridge, MA: MIT Press.

Leyton, M. (1992). *Symmetry, causality, mind*. Cambridge, MA: Bradford Books.

Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.

Nakayama, K., Shimojo, S., & Silverman, G. H. (1989). Stereoscopic depth: Its relation to image segmentation, grouping, and the recognition of occluded objects. *Perception, 18*, 55–68.

Norman, J. F., Phillips, F., & Ross, H. E. (2001). Information concentration along the boundary contours of naturally shaped solid objects. *Perception, 30*, 1285–1294.

Palmer, S. E. (1999). *Vision Science: Photons to phenomenology*. Cambridge, MA: MIT Press.

Pelli, D. G. (1997). The video toolbox software for visual psychophysics transforming numbers into movies. *Spatial Vision, 10*, 437–442.

Richards, W., & Hoffman, D. D. (1985). Codon constraints on closed 2D shapes. *Computer Vision, Graphics and Image Processing, 31*, 265–281.

Richards, W., Koenderink, J. J., & Hoffman, D. D. (1987). Inferring three-dimensional shapes from two-dimensional silhouettes. *Journal of the Optical Society of America A, 4*, 1168–1175.

Rosenholtz, R. (2001). Search asymmetries? What search asymmetries? *Perception and Psychophysics, 63*(3), 476–489.

Rosin, P. L. (2000). Shape partitioning by convexity. *IEEE Transactions on Systems Man and Cybernetics, Part A, 30*, 202–210.

Siddiqi, K., Kimia, B. B., Tannenbaum, A., & Zucker, S. W. (2001). On the psychophysics of the shape triangle. *Vision Research, 41*, 1153–1178.

Siddiqi, K., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception, 25*, 399–424.

Singh, M., & Hoffman, D. D. (2001). Part-based representations of visual shape and implications for visual cognition. In T. F. Shipley & P. J. Kellman (Eds.), *From fragments to objects: Grouping and segmentation in vision* (pp. 401–459). Oxford: Elsevier.

Singh, M., Seyranian, G. D., & Hoffman, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception and Psychophysics, 61*, 636–660.

Vecera, S. P., Vogel, E. K., & Woodman, G. F. (2002). Lower region: A new cue for figure-ground assignment. *Journal of Experimental Psychology: General, 131*, 194–205.

Wolfe, J., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience, 5*, 1–7.

Xu, Y., & Singh, M. (2002). Early computation of part structure: Evidence from visual search. *Perception and Psychophysics, 64*, 1039–1054.