

Processing convexity and concavity along a 2-D contour: figure–ground, structural shape, and attention

Marco Bertamini · Johan Wagemans

© Psychonomic Society, Inc. 2012

Abstract Interest in convexity has a long history in vision science. For smooth contours in an image, it is possible to code regions of positive (convex) and negative (concave) curvature, and this provides useful information about solid shape. We review a large body of evidence on the role of this information in perception of shape and in attention. This includes evidence from behavioral, neurophysiological, imaging, and developmental studies. A review is necessary to analyze the evidence on how convexity affects (1) separation between figure and ground, (2) part structure, and (3) attention allocation. Despite some broad agreement on the importance of convexity in these areas, there is a lack of consensus on the interpretation of specific claims—for example, on the contribution of convexity to metric depth and on the automatic directing of attention to convexities or to concavities. The focus is on convexity and concavity along a 2-D contour, not convexity and concavity in 3-D, but the important link between the two is discussed. We conclude that there is good evidence for the role of convexity information in figure–ground organization and in parsing, but other, more specific claims are not (yet) well supported.

Keywords Visual perception · Shape · Object representation · Convexity

M. Bertamini (✉)

Department of Experimental Psychology, Eleanor Rathbone Building, University of Liverpool, Liverpool L69 7ZA, UK
e-mail: M.Bertamini@liv.ac.uk
URL: <http://www.liv.ac.uk/VP/>

J. Wagemans

Laboratory of Experimental Psychology, University of Leuven (KU Leuven), Leuven, Belgium
e-mail: johan.wagemans@psy.kuleuven.be
URL: <http://www.gestaltrevision.be>

For the human visual system, contours are an important source of information about shape—and solid shape, in particular. This can be illustrated by considering that a paper and a pencil are sufficient to produce images that human observers interpret as meaningful. Even people without artistic training can draw objects, such as a house or a giraffe, that other observers can recognize without contextual help. This idea that contours—and curved contours, in particular—easily elicit the perception of solid shapes can be found, for instance, in Mach (1959).

From contours in an image, the visual system can extract information about convexities and concavities along that contour, and this specific process is the focus of this review. The visual system can extract contours from pictures, but also from projections of rich natural scenes (Kogo & Wagemans, 2012). Although scientific interest in this topic has a long history, the body of evidence has increased greatly in the last 20 years.

There are two main areas of perception in which convexity features prominently. The first is the process of separating foreground from background, and the second is structural description of shapes in terms of parts. We will consider these two aspects in turn, before considering also their implications for how attention is directed to contours. The structure of the review will rely on asking and trying to answer a series of questions (see Table 1): (1) Does convexity affect figure–ground organization? (2) Does convexity affect perceived metric depth? (3) Are convexities and concavities of a visual hole coded as in its complement (figure)? (4) Are convexities perceived as “parts”? (5) Are convex and concave extrema (maxima and minima) critical for parsing? (6) Is attention directed to convexities, concavities, or neither? (7) Is convexity coded explicitly in the brain?

Because of the historical interest in convexity in the Gestalt tradition, it seems appropriate to start with figure–ground organization, and this is the topic of the first three questions. We start with the role of convexity in figure–

Table 1 Seven questions used to consider what is the evidence in relation to the role of convexity in figure–ground organization (the first three questions) part parsing (the next two), and attention

Question	Our answer
(1) Does convexity affect figure–ground organization?	Yes.
(2) Does convexity affect perceived metric depth?	Not proven yet
(3) Are convexities and concavities of a visual hole coded as in its complement (figure)?	Not for perceptual tasks where contour ownership is important; perhaps in memory
(4) Are convexities perceived as “parts”?	Yes, or at least they play an important role in part parsing
(5) Are convex and concave extrema (maxima and minima) critical for parsing?	Not proven; at least the presence of extrema is not as important as the sign of curvature
(6) Is attention directed to convexities, concavities, or neither?	Neither; not in any automatic sense that is independent of the context
(7) Is convexity coded explicitly in the brain?	Yes

ground and ordinal depth (1) and the perceived separation (metric depth) between figure and ground (2) and consider whether visual holes are special (3). The next two questions are closely linked to the role of convexity in parsing and the highly influential “minima rule” for parsing (4 and 5). The role of convexity within these two areas (figure–ground organization and parsing) has been used to explain some effects of convexity in terms of how attention is directed to shapes. This is the topic of the sixth question (6). We finish with a question concerning what is known about how convexity is coded in the brain (7).

We will deal with this list of questions after an introduction in which we provide the definition of the terms *convexity* and *concavity* and a brief discussion of the relationship between convexity in 2-D and in 3-D. This link is fundamental to understanding why 2-D information is important for perception of solid shape, and clarity on this distinction between 2-D and 3-D is also necessary to avoid blurring together specific findings in the literature. The introduction will also include a section on the relationship between convexity and concavity information and the process of parsing of shapes.

Definitions

For any smooth contour in the image, it is possible to measure curvature at any location along the contour: the change in tangent direction as we move along the curve. For example, curvature along a circle is constant and is equal to the reciprocal of its radius and, therefore, decreases as the radius increases with scale. What is important to note is that in addition to a magnitude, curvature needs a sign.

To understand the meaning of *convexity* and the need for a sign, we start by thinking of a segment of a contour as a line that divides two adjacent portions of the plane, within a local aperture. We select a short segment so that it curves smoothly and does not include any inflections. For this segment of a contour, the region within the aperture that

includes all chords connecting any pair of points on the contour is labeled as convex, and the complementary region is labeled as concave. This sign allows us to distinguish the convex and concave regions, and by convention, a positive sign refers to convexity and a negative one to concavity.

If we were to label both regions, we would have a redundant and complementary sequence of convexities and concavities. That is, for every convexity, we have a corresponding concavity on the opposite side of the contour, and vice versa. To uniquely label a region as convex or concave, we need to choose only one of the two sides as we move along the contour. This can be done by choosing to label contours with respect to the “inside,” but this convention requires that we know which is the inside and which is the outside of the region of interest. For example, we can say that a circle has positive (convex) curvature only because we choose the sign with respect to the inside area of the circle. Without a rule to select the side, there is ambiguity about convexity and concavity. We will come back to this when talking about figure–ground organization.

A flat line is a special case of contour; it has zero curvature, but as long as we can measure some nonzero curvature, we can also assign both a magnitude and a sign to it. Furthermore, a point of zero curvature along the contour will also mark any smooth transition from convexity to concavity (inflection).

Figure 1 provides a visual example of contour curvature along a shape. The example is that of a large letter C, but a segment of the contour is also shown in isolation. The short segment within the aperture (circle) separates two regions. Although there are always two regions, one convex and one concave, the ambiguity is resolved on the basis of which is the inside. In the example in Fig. 1, the contour segment is convex if the region on the left (darker gray) is inside, but it is concave if the region on the right (lighter gray) is inside. If we think of the circular window as a receptive field of a neuron, then there is an ambiguity; this has to be resolved by some additional criterion that selects one or the other of the two sides to assign a unique coding to the contour segment. A trivial example would be to always code local contour




Fig. 1 The C-shape on the left has a contour that changes in curvature, positive on the outside and negative on the inside. Where curvature is high, this can be a positive maximum (M+) or a negative minimum (m-). Where curvature is low, this can be a positive minimum (m+) or a negative maximum (M-). A segment of contour is also shown in isolation in the circular aperture, and this shows the local ambiguity: It would be coded as convex when the darker gray region is perceived as the inside (foreground) or as concave if the lighter gray region is perceived as the inside. Because of the principle of unidirectional contour ownership, a contour belongs to only one side and, therefore, cannot be both convex and concave at the same time

curvature as convex. More plausible mechanisms would allow coding of both convexity and concavity on the basis of local as well as global properties of the shape, such as closure. A process based on closure would combine information across space, because closure is a global property.

The sign of curvature, positive or negative, corresponds to convexities and concavities, respectively, but in both cases, there are also important singularities along the contour when curvature reaches a local maximum or a local minimum. With respect to these singularities, we will use the standard convention of labels: M+ and m+ for positive maxima and minima, respectively, and M- and m- for negative maxima and minima. They are illustrated along the contour of the large letter C in Fig. 1. Note that M+ and m- are related to each other in having a local peak of unsigned curvature, and a figure-ground reversal would change one into the other. These local peaks are also referred to as *extrema* of curvature, and it has been suggested that they are especially informative about shape (Attneave, 1954).

There is an interesting asymmetry between convex and concave regions of a closed contour. This is best explained by considering a polygon with an arbitrary number of vertices. The exterior angle is the change of direction that takes place at each vertex as we follow the polygon in one direction. For example, a square has four exterior angles of 90° each. For squares, but also for all polygons, the sum of all exterior angles, taking concave ones as negative, is 360° , because to have closure we need to return back to the starting point. For example, if we use three squares to make an L shape, we have five angles of 90° and one angle of -90° ,

which sum to 360° . Since this is a positive number, we can say that there is always more convexity than concavity. The statement is not about the number of vertices of each type; there can be more concave vertices than convex vertices in the polygon, but they must differ in angle. In a hexagram (star of David), for instance, we have six convex and six concave vertices, but the convex ones are acute and, thus, their exterior angle is large.

What is true for polygons is also true for smooth contours, where, instead of a sum, we have a total turning angle. Feldman and Singh (2005) have provided a formal treatment of this asymmetry (for an extended discussion, see also Lim & Leek, 2012; Singh & Feldman, 2012). Because local curvature for closed contours is not equally likely to be positive or negative, it follows that concave regions are more informative. This could have important implications for how attention is directed to regions that carry more information.

A finite (closed) figure can be strictly convex, as in the case of a circular object, but it cannot be strictly concave. Terminology can be confusing because, for polygons, the term *convex* is used for strictly convex polygons but the term *concave* is used for a polygon with one or more concave angles (all others). However, a strictly concave contour can be defined, but by definition, it would have to be a hole within a larger object. We will consider holes as a special case and the related literature in the section on visual holes.

The idea that shape perception relies on convexity and concavity information can be traced back to what is arguably the first book on vision science: the book of optics written by the Arab scholar Alhazen (Ibn al-Haytham) (965–c. 1040 AD) (Howard, 1996). One of his claims was that convex surfaces, but not concave ones, lead to the perception of a solid object (“if the body has a convex surface that bulges towards the eye . . . then if sight perceives the convexity of the surface it will perceive the body’s solidity”; Sabra, 1989, p. 169). This raises the issue of how to define convexity and concavity for a surface, because so far we have given a definition only for a 2-D contour. The relationship between convexity in 2-D and 3-D is important and a source of potential confusion in the literature.

As we have seen, the definition of 2-D convexity on the basis of the sign of contour curvature is simple, and it matches the intuitive meaning of convexity for a 2-D shape. But on a surface, curvature can be measured in more than one direction, and for a given location, there are two orthogonal directions that have maximal and minimal curvature, as was proven by Euler in 1760. These are known as principal directions. It follows that, if neither is zero, they can both be positive (the surface is locally convex), both negative (locally concave), or one positive and one negative (saddle shaped). Koenderink (1984; see also Koenderink & van Doorn, 1982; Koenderink, van Doorn, Kappers, & Todd,

2001; Lappin & Craft, 2000; Lappin, Norman, & Phillips, 2011; Richards, Koenderink, & Hoffman, 1987) has shown that convexities and concavities of contours in a projected image allow the observer to draw inferences about local surface geometry. Convexities in the image are projections from a surface that is locally convex, and concavities are projections from a local saddle shape. This is true independently of distance, and it relies only on the assumption that a contour in the image corresponds to the rim of a smooth surface in 3-D (the rim is defined as the boundary between the visible and invisible parts of the object).

There is strong evidence that the visual system relies on a convexity assumption for 3-D shapes (Hill & Bruce, 1993; Johnston, Hill, & Carman, 1992), in agreement with Alhazen's intuition. A good example of a convexity assumption is provided by the hollow-face illusion where the inside of a mask is misperceived as convex. Recent evidence suggests that this 3-D convexity assumption is present already in 6-month-old infants (Corrow, Granrud, Mathison, & Yonas, 2011), although the role of experience in this illusion is still a topic of active research. It is interesting, for example, that the illusion is stronger for upright than for inverted faces (Hill & Bruce, 1993). Moving away from the special case of faces, Langer and Bühlhoff (2001) tried to compare different assumptions about 3-D shape and concluded that the convexity assumption is as strong as two other assumptions: the viewpoint from above and the light from above (but see also van Doorn, Koenderink, & Wagemans, 2011; Wagemans, Van Doorn, & Koenderink, 2010). What is important for this review is that a bias for 3-D convexity over 3-D concavity does not imply an equivalent bias for 2-D convexity over 2-D concavity, because of the difference discussed earlier in how 2-D information is information about 3-D surface curvature—namely, the fact that a concave contour is not the projection of a concave surface.

The rest of this article will focus on 2-D contour information, but in the next section, we will cover the link between 2-D convexities and concavities and perceived part structure.

Parsing

We have seen that contour curvature provides information about solid shape (Koenderink, 1984, 1990). However, once information about surface curvature is available to the system, it can also be used to parse objects, leading to a structural description based on parts.

Hoffman and Richards (1984) have proposed a parsing rule based on the principle of transversality: When two solid shapes overlap, their intersection creates a concave crease. On a smooth surface, this crease is a saddle region (Fig. 2). Therefore, it makes sense to divide solid shapes into parts at concavities along the image contour, and in addition, the

precise transition between the two parts will be located at the point with the greatest negative curvature (i.e., minima of curvature). Because of this role of minima, Hoffman and Richards's proposal goes beyond the role of the sign of curvature and is known as the "minima" rule.


Although the minima rule adds to the general idea of a link between convexity and parts, in itself it is not enough to exactly define parts; therefore, further elaborations have been proposed. Hoffman and Singh (1997) have proposed a combination of salience factors—size of the part relative to the whole object, degree of protrusion, and strength of the part boundaries—and Singh, Sahrawian, and Hoffman (1999) have proposed a short-cut rule based on length of cut. A related but different approach has been taken by Siddiqi, Tresness, and Kimia (1996) and Dhandapani and Kimia (2002) on the basis of necks and limbs. Other models of segmentation are based on convexity in general (Rosin, 2000) or look for the largest convex patches (Vaina & Zlateva, 1990). In some approaches, all of these factors have been integrated into one coherent framework (De Winter & Wagemans, 2006). Later in the article, we will consider the empirical evidence for a role of concavities in perceived part structure.

Local and global definitions of convexity

One final issue of definition must be discussed. We have introduced convexity as a local property of contour curvature: the local sign of curvature. A different use of the term refers, instead, to the global degree of convexity of complex 2-D shapes. The convex hull or envelope of a set of points or vertices of a polygon is defined as the minimal convex set containing all those points. However, very different shapes can share the same convex hull. Using the convex hull, we can define a global measure of shape convexity as the proportion of the area of a polygon over the area of its convex hull (Preparata & Shamos, 1985). More sophisticated measures of global convexity have been put forward, for instance, by Pao, Geiger, and Rubin (1999) and by Rosin (2000). So far, these measures have not been fully validated in terms of their role in human perception. It can clearly be useful to have a formula that assigns a single value of convexity to any complex 2-D shape. In image analysis, this type of measure is known as a shape factor: a value that is affected by an object's shape but is independent of its dimensions. An example that is closely related to that of convexity is that of compactness, defined as the square of the perimeter over the area or, alternatively, as the area of a shape over the area of a circle having the same perimeter. Another well-known shape factor is the aspect ratio.

In the previous paragraph, we contrasted a local meaning of convexity (the sign of curvature at one location along a

Fig. 2 Top: Concave creases (and therefore, minima of curvature along the projected contour) are informative about the transversal intersection of solid shapes (Hoffman & Richards, 1984). Note the pair of minima along the contour that mark the boundary between the two intersecting objects. Bottom: A silhouette of a horse with lines showing how observers parsed the object (from De Winter & Wagemans, 2006). Note that in many cases, local concavities correspond to the part cuts



contour) and a global meaning of convexity. An intermediate definition is also possible. For example, Fowlkes, Martin, and Malik (2007) computed convexity within a region from the fraction of point pairs for which a straight line connecting the two points lies completely within the region. This value can be computed for both sides of the contour, and convexity is then given by the log ratio of the two values. Changing the size of the region would make this measure more local or more global.

In the rest of this article, we will refer to convexity in the sense of the local sign of a contour or, for a more complex contour, as the piecewise property of the segments that make up that contour. Some examples are discussed in the next section.

Does convexity affect figure–ground?


Convexity was discussed within the Gestalt tradition, in particular in relation to figure–ground organization (Koffka, 1935; see also Wagemans et al., 2012). Arnheim (1954) described a simple demonstration in which a symmetrical shape is perceived as a figure when convex and as a hole when concave (Fig. 3). This effect was not tested empirically by Arnheim, but it was later confirmed empirically (Bertamini, 2006). In the terminology used by Arnheim, the convexity of a closed contour can be interpreted as a piecewise description of the convexities along the contour that form the shape, and similarly for concavity. Alternatively, a global measure of convexity could be applied, and

Pao and Geiger's (2001) formula would give a result consistent with Arnheim's labels.

The first experimental test of the role of convexity in figure–ground organization was carried out by Kanizsa and Gerbino (1976). They used a task in which observers had to choose which region appeared as foreground. As is shown in Fig. 3, the black and white complex regions could be piecewise convex or concave. Kanizsa and Gerbino concluded that convexity was a powerful factor and that it was stronger than symmetry when one was pitted against the other, as is the case in the examples shown in Fig. 3. Similar stimuli have been used by Hoffman and Singh (1997; see their face–vase demonstration) and, more recently, by Vecera et al. (2004) and by Peterson and Salvagio (2008). These studies have confirmed the role of convexity in figure–ground assignment. However, Peterson and Salvagio found that the convexity effect was weak for very simple displays with just two regions but that it became stronger as the number of edges delimiting alternating convex and concave regions increased. They argued that this is evidence of a spreading of the convexity bias from region to region.

The association between convexity and figure–ground was already discussed earlier in relation to the definition of convexity. Deciding what is the inside of a region with a closed contour will determine where the contour is convex and where it is concave. For example, for a surface specified within a random dot stereogram, there is no ambiguity about what is foreground and, therefore, no ambiguity on where the inside of the surface is. For such a surface, there is also no ambiguity on where convexities and concavities are. In

Fig. 3 The piecewise convexity of the shape in the examples determines the perception of that region as figure. In the example from Arnheim (1954), the concave shape tends to be seen as a hole (Bertamini, 2006). In the example from Kanizsa and Gerbino (1976), convexity wins over symmetry. The effect of convexity increases with number of regions (Peterson & Salvagio, 2008), and it increases after an exogenous precue of the region (Vecera, Flevaris, & Filapek, 2004)



other cases, the inside is less obvious, as in all the examples in Fig. 3, and the degree of local convexity is one of the factors that determine figure–ground. However, this convexity bias is only one of many figure–ground factors, and therefore other factors contribute to the interpretation of the contour (i.e., size, closure, shading, etc.). The competition between these factors has been described as a form of reciprocal inhibition across the two opposite sides of a contour (Peterson & Skow, 2008). In some models, this inhibition is also at the heart of the dynamics of switching between two alternative figure–ground organizations with identical inputs (Kogo, Galli, & Wagemans, 2011).

Although an extensive review of the computer vision literature is beyond the scope of this article, it is interesting to note that convexity is also used as a cue to identify regions of interest in computer vision models (e.g., Asada & Brady, 1984; Jacobs, 1996; Zunic & Rosin, 2002).

Kim and Feldman (2009) have reported an interesting extension of the effect of convexity in figure–ground that we have just described. They found that even for relatively simple shapes with local concavities, these are sometimes perceived as convexities despite the fact that, in terms of the global interpretation of the object, this creates an inconsistency. They concluded that contour ownership assignment can act as a local process and is not always constrained by global information.

The idea that convexity is linked to foreground is appealing because we live in a world of finite opaque objects. This aspect of the natural environment has recently been documented and quantified by Burge, Fowlkes, and Banks (2010). Using a range finder, Burge et al. (2010) confirmed that contours sampled from natural images tend to separate a farther region on the concave side and a nearer region on the convex side. Therefore, other things being equal, it is

rational to assume that the convex side of a contour is more likely to be the foreground.

In addition to the studies that have shown a preference for identifying the convex region as foreground (e.g., Kanizsa & Gerbino, 1976), there is also evidence that the link affects tasks in which convexity is task irrelevant. Bertamini and Lawson (2008) presented observers with random dot stereograms in which a square region seen within an aperture was split into two surfaces. Observers were asked to report whether the foreground was the surface on the left or that on the right. Response times were longer when the concave side was in front, despite the fact that shape information was not available to the visual system of the participants prior to depth information and despite the fact that the shape of the foreground was irrelevant for the task. We can link this finding to Feldman and Singh's (2005) analysis: Other things being equal, a contour segment in isolation is more likely to correspond to a convexity (see also Burge et al., 2010).

The issue of figure–ground organization can be separated from that of grouping (see also Wagemans et al., 2012). Are regions grouped together on the basis of convexity? Liu, Jacobs, and Basri (1999) made the assumption that the stronger the grouping between two regions, the harder it will be to resolve relative stereoscopic depth. They used a display in which two regions were presented one above the other and were partly occluded in the middle (Fig. 4). For regions that formed a convex object (as compared with regions connected by concavities), it was harder to resolve their relative stereoscopic depth. Using this technique, they concluded that convexity is a strong factor that groups two regions together.

Fantoni, Bertamini, and Gerbino (2005) found that the extent of amodal completion of a partly hidden surface




Fig. 4 Top: In the study by Liu et al. (1999), an occluder (black) was present in front of a pair of regions within a stereogram. For convex contours (left), grouping was stronger, and consequently, there was a tendency to perceive the two figural elements as coplanar, making the task of adjusting relative depth harder. Bottom: Illustration of the stimuli used in one of the experiments by Fantoni et al. (2005) to compare amodal completions of shapes. The small hexagon in front was the occluder, and the partly occluded contours had convex or concave vertices. The task was to estimate the vertical distance (gap) between the vertices. Greater extrapolation behind the occluder was found for concave vertices

depends on whether the contours that are occluded are convex or concave. In a series of studies, they asked observers to judge the perceived extent of a vertex that was partly occluded. The stimuli used in one experiment with random dot stereograms are described in Fig. 4, using shading to show depth stratification. Observers had to adjust a probe to report the location of the hidden vertices of the object (or of the hole). In the condition where the vertices were concave, the extrapolation was greater. One possibility is that extending a concave region minimizes the surface that is amodally created, but more research is necessary to test this hypothesis. Fulvio and Singh (2006) reported a similar difference between convex and concave contours, using illusory surfaces instead of amodal completion: When the contour was perceived as convex, it was smoother (less extension) than when it was perceived as concave.

Does convexity affect metric depth?

Because convexity affects figure–ground organization, convexity can be described as an ordinal depth cue.

An important question is how ordinal and metric information about depth is combined. Burge, Peterson, and Palmer (2005) presented observers with a pair of stereograms and asked them to select the one that appeared to have greater depth between foreground and background (see Fig. 5). In one condition, the foreground had a familiar shape (a face profile, mainly convex); in another condition, the foreground was the complement of the face profile (mainly concave). Burge et al. (2005) found that greater disparity was required in a concave stimulus, as compared with a convex stimulus, to achieve a subjective equality of depth. That is, there was a shift in the point of subjective equality (PSE). In a later paper, Burge et al. (2010) found a similar effect for stimuli that separated convexity from familiarity.

The idea that convexity leads to an increase in perceived metric depth has been criticized. Bertamini, Martinovic, and Wuerger (2008) used random dot stereograms and replicated the shift in PSE. But the fact that luminance was not necessary is problematic because, in previous work on figure–ground organization, the role of factors such as familiarity was found only when shape information was carried by luminance contours (Peterson & Gibson, 1993). An alternative explanation is that the shape information created a response bias. Instead of a forced choice task, Gillam, Anderson, and Rizwi (2009) used a different technique (a depth probe) and concluded that the effect found by Burge et al. (2005) was probably a response bias. Note that this paper was published before Burge et al. (2010), and the debate continues.

Are convexities and concavities of a visual hole coded as in its complement (figure)?

Visual holes are regions of ground that are completely enclosed. An interesting question is whether coding of convexity and concavity is ambiguous in this case or whether it proceeds as if the hole region were a foreground. Palmer raised the issue of holes in 1999 and suggested that perhaps they are figures for purposes of describing shape (Palmer, 1999, p. 287). Some support for this idea comes from studies in which the memory for the shape of a hole was similar to the memory for the shape of an object (Nelson, Thierman, & Palmer, 2009; Palmer, Davis, Nelson, & Rock, 2008). In a paper that was more specifically concerned with convexity coding, Feldman and Singh (2005) suggested that holes have “a quasi-figural status, as far as shape analysis is concerned” (p. 248) and, therefore, regions coded as convex for a closed contour would be equally convex for a congruent contour perceived as a hole. A different interpretation has been put forward by




Fig. 5 The stereograms used by Burge et al. (2005) had two surfaces defined by a difference in luminance as well as disparity. One region had the shape of a profile, thus having both a high degree of familiarity and convexity. In this example, we assume that the lighter region is the foreground as specified by disparity. On the right are examples that remove the role of familiarity

Bertamini (2006), who suggested that memory tasks may rely on strategic use of shape information derived from the object-with-hole. In tasks where responses have to be produced quickly, the evidence is that holes (ground) and objects (figure) behave differently. A series of papers (Bertamini & Croucher, 2003; Bertamini & Farrant, 2006; Bertamini & Mosca, 2004) have started from the phenomenon that vertices perceived as convex are responded to faster and have found that, for the same outline, the response speed changed when that region was perceived as a hole or as an object, suggesting that convexity for a hole is coded on the basis of the inside being a ground region. These papers are discussed in more detail later in relation to attention allocation.

Holes have been used in other studies as well, although these are less relevant with respect to the study of convexity and concavity. Albrecht, List, and Robertson (2008) found that the region cued inside a hole is the background surface; Bertamini and Hulleman (2006) studied moving holes and found resistance to perceiving holes moving independently from the object-with-hole; Horowitz and Kuzmova (2011) used the multiple object tracking paradigm and found that holes can be tracked as well as objects, and Zhou, Luo, Zhou, Zhou, and Chen

(2010) found that tracking was significantly impaired when the items changed their topological properties (holes were added or removed).

Are convexities perceived as “parts”?

Some influential authors have argued that perception of visual shape relies on a representation of constituent parts and the spatial relations among them (Biederman, 1987; Marr & Nishihara, 1978). As was discussed in the introduction, the link between convexities and part structure has a long history, going back to Alhazen. More recently, Koenderink (1990) has pointed out that “things growing or blown up from the inside tend to possess positive Gaussian curvature. . . . An extreme example is the famous *Venus of Willendorf* from Palaeolithic origin. Clearly, this fact that we all seem to know from introspection is not a short-lived cultural whim” (p. 251).

Hoffman and Richards’s (1984) minima rule provides a specific mechanism that links concavities with part parsing: The visual system parses contours at extrema of negative curvature (i.e., minima). Many empirical studies have tested this proposal or, at least, interpreted their findings in the light of the minima rule. Braunstein, Hoffman, and Saidpour (1989) presented participants with a surface defined by dots and rotating about a vertical axis. This was followed by a display of four alternative parts. Most of the correct responses in this memory task were generated when the part presented at test was defined by minima. In their Experiment 3, participants were asked to mark apparent part boundaries: 81 % of participants marked part boundaries indicating minima parts, 10 % indicating maxima parts, and 9 % other positions. Another study that attempted to test the minima rule was by Lamote and Wagemans (1999). They used closed contours and found that convex (M+) deletions were easier to detect than the deletions around the other two singularities (m- and inflections). Although both singularities involve regions of high curvature changes, they differ in that concave (m-) deletions create gaps that concur with spontaneous segmentation. Xu and Singh (2002) have made a very similar argument to explain their visual search results. When a silhouette is segmented by introducing a cut, this segmentation changes the perception of that object drastically, making it more salient when placed among silhouettes without a cut, but only if this cut does not correspond to concave cusps. The rationale is that the concave cusps induce a part parsing similar to that of the cut, making the items similar in terms of part structure. When items have similar part structure, visual search becomes more difficult.

It is interesting to note that parts have many properties in common with visual objects. For example, Sakai and Inui (2002) have found that the capacity of visual short-term memory (VSTM), defined by pattern complexity (the degree that a pattern can be handled for several seconds with little loss), is four convex parts (see also Bertamini, Helmy, & Hulleman, 2012, for more evidence about VSTM).

A memory task for silhouettes was used by Cohen and Singh (2007). The sequence was the reverse of the memory task used by Braunstein et al. (1989). Instead of seeing the object and then one part that may have come from it, observers first saw a segment and then had to decide whether the segment was present in the silhouette. Cohen and Singh found that identification of contour segments was modulated by curvature: Segments with minima boundaries were identified with greater accuracy than segments with positive maxima or inflections of the same length.

There is also work with young children that supports the claim that convexities are perceived as parts. These studies built on previous work that had demonstrated that infants focus their attention on edges and line junctions (e.g., Bhatt & Bertin, 2001; Salapatek, 1968). Bhatt, Hayden, Reed, Bertin, and Joseph (2006) found that 5- and 6-month-old infants detected concave elements amid convex distractors, but not convex elements among concave distractors (but see below for our discussion of this type of search asymmetry). More recently, Hayden, Bhatt, Kangas, and Zieber (2011) habituated 5-month-olds to a silhouette that contained a segment presented either as convex or as concave (i.e., a cavity). In a novelty preference test, children exhibited a preference for the concave shape, indicating that convex shapes were more salient than concave shapes during habituation, or at least that changing a concave shape into an individual closed object is more surprising than changing a convex shape into a closed object.

Another phenomenon supports the idea that a convexity is perceived as a part, whereas a concavity is perceived as a boundary between parts. When observers are asked to judge the position of a vertex, they are faster if the vertex is perceived as convex. This effect was studied using a vertex, as compared with a reference line, by Bertamini (2001), but it had already been noted, using simple polygons, by Gibson (1994). For a given polygon, this effect reverses when figure-ground relationships are reversed (Bertamini & Croucher, 2003), and this reversal is particularly clear when random dot stereograms are used so that there is no ambiguity about which region is the figure and which is the ground. Using stereograms, a specific closed region can be perceived as either a surface (object) in front of a background or a hole (ground) (Bertamini & Farrant, 2006; Bertamini & Mosca, 2004). Specifically, an interaction is present between type of object (object or hole) and shape that can be explained by faster responses to vertices perceived as

convex because these are convex surfaces that define structural parts. This role of convex regions on the contour is predicated on the fact that contours are interpreted as rims of surfaces. It follows that thin (wire-like) objects should not behave as silhouettes, because their contour cannot be interpreted as a rim. Bertamini and Farrant (2006) confirmed that thin objects were different from both objects and holes.

In a series of studies, Baylis and Driver (1993, 1994, 1995, 2001) found evidence of a critical role of convexities and concavities in perception of shape and symmetry. They concluded that symmetry detection across an axis was easy because convexities and concavities were matched on the two sides. This is true for a reflection but not for a translation of a pair of contours closed to form a single object and can, therefore, explain the large difference in performance for reflection and translation. By manipulating figure-ground organization (by changing the way that the contours were closed to form objects), Baylis and Driver demonstrated that repetition can be made easier to detect and reflection harder (see Fig. 6).

Although it seems intuitively appealing that convexities are matched with concavities when translation is present in two facing objects (like the pieces of a jigsaw puzzle), evidence using parts that are too complex to fit after a simple translation in the plane (Bertamini, Friedenberg, & Kubovy, 1997; Friedenberg & Bertamini, 2000) or similar findings for translation and for rotation within a single object (Bertamini, Friedenberg, & Argyle, 2002) make the jigsaw explanation less satisfactory. Treder and van der Helm (2007) used stereograms to place the corresponding parts of the object on the same or on different planes. They found that efficient detection of symmetry depends on structural correspondence of parts within a depth plane, but the same was not true of detection of translation. This supports a close link between symmetry and the presence of a single




Fig. 6 Detection of reflection and repetition is greatly affected by convexity coding (Baylis & Driver, 1995). The task is to compare the two sides and detect whether the contours match. The match can be a reflection (top) or a translation (bottom)

object. Koning and Wagemans (2009) used projected 3-D objects and concluded that structural coding in terms of perceived parts, not a matching strategy of the images, underlies the one-object advantage for symmetry and the two-object advantage for translation.

Given the convexity advantage in a task about position, it is perhaps not surprising that convexities also play a greater role in perception of symmetry. Hulleman and Olivers (2007) asked observers which of a pair of silhouettes was perfectly symmetrical, and the deviations from symmetry could be located at convexities or concavities. Detecting asymmetry was easier when there was a mismatch between the convexities on either side of the axis of symmetry. However, a later study by Bertamini et al. (2012) found that the effect was mediated by a strategic preference by observers for monitoring the region with convexities, rather than the region with concavities.

Barenholtz and Feldman (2003) studied a task in which observers had to compare a pair of small probes located along a sinusoidal contour. Performance was better when both probes were located along a convex contour, rather than a concave contour. They interpreted this in terms of part structure because, when the two probes were on a concave contour, they may have been perceived as belonging to two separate objects, thus making the comparison more difficult (Fig. 7). A similarity with the positional judgment studied by Bertamini and Farrant (2006) is the fact that neither effect is retained when snake-like contours are used, presumably because only for closed shapes do we have unambiguous convex and concave regions. However, Barenholtz and Feldman (2003) used only a probe discrimination task, and it would be interesting to test performance when observers have to respond to probes located along a convex or a concave contour so that no comparison is necessary.

When contours move, the motion is likely to be organized so that what moves are the convex parts (Barenholtz & Feldman, 2006). Although Barenholtz and Feldman (2006) studied several configurations, the simplest to imagine is that of a single vertex seen within a circular aperture separating a concave and a convex side. With a change in the angle, the figure-ground ambiguity is reduced, and observers report two parts moving around a concave “hinge.” This interesting phenomenon can be interpreted as indirect evidence of parsing, since the concavity provides a better articulation due to its role as a boundary between parts (Barenholtz & Tarr, 2008).

Are extrema (max and min) critical?

There are two aspects to this question. The first problem is whether extrema of curvature play a role in describing and storing shape information. The second aspect is specific to parsing; we have seen that convexities are perceived as parts and concavities as the boundaries between parts, but how critical is the presence of extrema in parsing?

With respect to the first, more general question, in 1954, Attneave provided an explicit claim that extrema are the most informative points for shape perception. In a now famous demonstration, he created a drawing of a sleeping cat made by selecting 38 points of maximum curvature (both convex and concave) from the original image contour and connecting these with a straightedge. Despite the loss of information, the cat is easily recognizable. Attneave made no suggestion of a fundamental difference between minima and maxima. The recent paper by Carlson, Rasquinha, Zhang, and Connor (2011), cited earlier, makes direct reference to Attneave. They made recordings from macaque monkey's area V4 and concluded that object representation




Fig. 7 Observers were faster at responding to information located on a convexity (B) rather than a concavity (A). Because the task involved comparing two small probes within the highlighted region, Barenholtz and Feldman (2003) concluded that there is a cost in comparing probes

that cross a boundary between parts. An interesting control condition used “snakes” instead of “peanuts.” Here, convexity and concavity are the same, but for snakes, there is no obvious part segmentation; therefore, the effect did not extend to these shapes

is based on uncommon but diagnostic regions of “acute contour curvature” (both convex and concave).

Attneave's focus on extrema can be linked with another idea, also present in the literature, that corners are highly informative. This was suggested for example by Hebb (1949) in his classic book on *The Organization of Behavior*, in which he claimed that corners are locations of greatest activity in the visual field.

Contrary to Attneave's (1954) idea, Kennedy and Domander (1985) suggested that the most informative regions are those with least curvature and found empirical evidence for this proposal, later supported also by findings by Panis, De Winter, Vandekerckhove, and Wagemans (2008). Kennedy and Domander's study, however, has been criticized by Norman, Phillips, and Ross (2001) because they used man-made objects (e.g., chest of drawers, window, etc.) with mainly straight edges. Norman et al., instead, used more natural objects (silhouettes of potatoes) and asked participants to select a subset of locations to best describe the shape. They found that there was a correspondence between the chosen positions of the points and the locations along the boundary contour that were local curvature extrema (convexities or concavities). This has been confirmed with the same method on a larger sample of outlines of everyday objects in a study by De Winter and Wagemans (2008). Subsequently, Panis and Wagemans (2009) have tried to resolve these apparent contrasting results by examining the time course of object identification. They concluded that straight edges were more important for complex outlines early in the matching process and that curved edges were more important for simple outlines in later (top-down) matching processes.

We now turn to the question of the specific role of extrema in parsing. Great impetus to this line of research came from the minima rule (Hoffman & Richards, 1984). Braunstein et al. (1989) presented observers with a surface rotating about a vertical axis, followed by a display of four alternative parts. These parts were extracted from the object by segmenting the outline of the shape at concavities or at convexities. When both a part defined by concavities and a part defined by convexities were presented, most of the correct responses were made by choosing the part defined by the concavities. Moreover, in another experiment, participants were asked to explicitly mark part boundaries, and most of them used concavities for parsing (81 %), a minority used convexities (10 %), and a few used other positions (9 %) (Braunstein et al., 1989).

Another well-known series of studies by Baylis and Driver (1993, 1994, 1995) found evidence of a critical role of convexities and concavities in memory for shapes and perception of symmetry. Baylis and Driver concluded that part parsing plays an important role in perception of shape and symmetry and that parts are

created by convexities and concavities. They related this to Hoffman and Richards's minima rule, and for instance, the minima rule is cited within the abstract of their 1995 paper. Note, however, that these findings are about convexity and concavity, and not the more specific role of minima. Baylis and Driver (1995) chose their words carefully and concluded that “our present results show that these parts *may be defined* in accordance with Hoffman and Richards's (1984) minima rule” (p. 1340; italics ours). Caution, therefore, has to be taken not to use this evidence as directly supporting the minima rule.

The role of extrema has to be evaluated in the context of other factors as well. For example, Singh and Hoffman (2001) noted that the curvature at a negative minimum and the turning angle can vary independently. This is because the minimum is a local value, while the turning angle is a more global aspect of the shape (the angle through which the contour turns between the inflection points). Singh and Hoffman (2001) suggested that parsing for a given minimum will depend on the turning angle because both affect the salience of a part cut. In their study of perceived transparency, Singh and Hoffman (1998) found evidence of the importance of the turning angle in parsing. This combination of various factors in determining a part cut has also been noted by De Winter and Wagemans (2006; see below). Inevitably, the simplicity of the original idea that local minima are the key to parsing is lost with the recognition that minima are just one of several factors and that the same minimum in different contexts is more or less salient.

An even more critical comparison for the evaluation of the minima rule is that between concavities that have a minimum and concavities that do not have a minimum. Hulleman, te Winkel, and Boselie (2000) found that shapes with concavities are easier to find in visual search, as compared with shapes without concavities. But the presence or absence of a minimum within the concavity made no difference. We will consider this study again in the discussion of how convexity and concavity drive attention.

To evaluate the various factors that play a role in parsing 2-D silhouettes, De Winter and Wagemans (2006) conducted a large-scale study in which observers were asked to segment outlines of everyday objects. All previously proposed rules of segmentation were supported. Minima of curvature were often used for segmentation, and segmentation points were often near extrema. However, segmentation also depended on global characteristics such as proximity, collinearity, symmetry, and elongation. From their extensive analysis, De Winter and Wagemans (2006, 2008) concluded that locations are most salient when the connecting segments to the next salient point create a sharp turning angle and when the resulting part sticks out from the rest of the object.

Is attention directed to concavities?

Much research has been carried out on what guides visual attention. In their review, Wolfe and Horowitz (2004) listed the direction of curvature as one factor that guides attention (for curved lines that have curvature to the left vs. to the right; Treisman & Gormican, 1988). When the contours are part of a closed contour, convexity and concavity are features of a new dimension, and there is a “possible preference for concavities” (p. 5). To support this conclusion, Wolfe and Horowitz cited Barenholtz, Cohen, Feldman, and Singh (2003). We now turn to this experimental evidence.

Barenholtz et al. (2003) used a change detection task. A simple polygon was presented for 250 ms, then a mask was shown for 200 ms, and then a new polygon was presented that could be the same or could have a vertex added or removed. Participants had to report whether the shape was the same or had changed. Barenholtz et al. found that changes were more easily detected if they involved a concavity, rather than a convexity. The difference in performance was large: Percent correct for concave and convex changes was 70.83 % and 36.94 %, respectively, and sensitivity (d') was above one in the concave condition and was close to zero in the convexity condition. This was true when the new vertex was added or removed and was not affected by the magnitude of the vertex displacement. The title of this paper was “Detection of Change in Shape: An Advantage for Concavities,” and this suggests a basic difference between convex and concave vertices. However, the manipulation was one that compared changes in polygons that were not different only in sign of curvature.

This change detection paradigm was later used in other studies: Cohen, Barenholtz, Singh, and Feldman (2005), Bertamini and Farrant (2005), and Vandekerckhove, Panis, and Wagemans (2007). These studies replicated the effect, showing higher sensitivity for concave changes. However, there are differences of interpretation. Vandekerckhove et al. distinguished a local and a global account. In the local sense, the difference is between any local concavity and any local convexity. In the global sense, the difference is due to the fact that a new concavity in a simple object changes the way the object is parsed, introducing extra parts. Vandekerckhove et al. extended the original finding to filled polygons, smoothly curving shapes, and quantitative changes to the angles that did not change part structure. Even a change in the size of a concave vertex, which should not change perceived part structure, was more easily detected than a change in a convex vertex. On the basis of these results, they argued for both local and global effects. Cohen et al. also adopted the strategy of changing only the local angle without introducing new vertices, and they concluded in favor of the local explanation.

In addition to the problem of distinguishing a local and a global interpretation of this effect, there are two additional methodological problems. Any new concave vertex tends to be located nearer the center of a polygon and nearer fixation, as compared with a new convex vertex. Distance from fixation may affect performance, and independently of the location in the visual field, the center of an object can provide a frame of reference for judgments about position (Wang & Burbeck, 1998). Another problem is that the polygons used in Barenholtz et al. (2003) and Vandekerckhove et al. (2007) had, on average, more convex vertices than concave vertices. As a consequence, a new convex vertex was added along a contour already containing several convex vertices, but a new concave vertex was added along a contour that contained few if any such vertices. Some, but not all, of these methodological issues with the study by Barenholtz et al. (2003) were addressed in later experiments by the same authors (Cohen et al., 2005).

To resolve these issues, Bertamini and Farrant (2005) and Bertamini (2008) used a similar change detection procedure but isolated a segment of contour within an aperture. To ensure that, within the aperture, there was a foreground and a background, they used stereograms. The design of Experiment 2 in Bertamini and Farrant (2005) made it possible to compare the detection within a convex/concave context and the detection when a change of sign is present or absent (Fig. 8). They found that the only significant factor was change of sign: a new concave vertex along a convex contour or a new convex vertex along a concave contour. Furthermore, in Bertamini (2008), various conditions were compared, including conditions in which no new vertices were introduced. A change of sign was easy to detect, but there was no advantage for detecting a new concavity versus a new convexity, a change of angle in a concave vertex versus in a convex vertex, or a change within a concave region versus a change within a convex region. Detection was affected by changes of perceived part structure, but there was no evidence of a special role for concavity or minima of curvature in guiding attention (Bertamini, 2008).

Because of the key role of a change of sign in change detection, Bertamini and Farrant (2005) suggested the term *bracketing*. The bracketing hypothesis says that although the presence of convexities and concavities is important, it is not the convexities or the concavities per se that guide attention, but their sequence and any change in this sequence, because these change perceived part structure.

The type of concavity advantage reported using the change detection paradigm, or at least the *prima facie* advantage, appears consistent with the advantage in detection of concavities using the visual search paradigm (Hulleman et al., 2000; Humphreys & Müller, 2000; Wolfe & Bennett, 1997). Within this literature, some authors have speculated that, because concavities have a central role in describing




Fig. 8 Top: Illustration of the change detection task used by Barenholtz et al. (2003). The new concave vertex affects part structure in a way that is not true for the convex vertex. Bottom: Illustration of the conditions in the change detection task used by Bertamini and Farrant (2005). This diagram highlights the type of changes and does not reflect the actual stimuli. Performance was high in conditions B and C, where there is a likely change in perceived part structure

shape, “fast selection of concave edges, perhaps via the activation of specialized detectors, may thus serve a useful computational purpose” (Humphreys & Müller, 2000, p. 197).

Hulleman et al. (2000) conducted four experiments and found flat slopes (<10 ms/item) when observers searched for

a target with a concave cusp among totally convex distractors. And there was a search asymmetry, because reversing the roles of target and distractor yielded inefficient search. This was true for a concave cusp, but also for other types of concavities, including a concavity with constant curvature. They concluded that concavities act as basic features in visual search. Humphreys and Müller (2000) reported similar findings, using different shapes. In addition, they found that for the same contours, the effect could reverse with a figure–ground reversal, thus supporting the key role of convexity and concavity.

As can be seen in Fig. 9, the concave target is a complex object with both convexities and concavities. One of Hulleman et al.’s (2000) conclusions was that to explain the asymmetry, one has to recognize that the search is not performed simply on the local property of the contour, because the concave cusp was matched to the convex cusp. However, it would be problematic to take this as a basic difference in attention allocation to concavities, as opposed to convexities. The issues are similar to those raised with respect to the change detection paradigm—namely, that there are fundamental asymmetries in the conditions (for an interesting discussion of asymmetries of design in search asymmetry studies, see Rosenholtz, 2001). The shape that includes a concavity is more complex, with changes of sign along its contour, and it is also likely to differ in terms of its perceived part structure. Hulleman et al. argued that part structure did not play an important role, since all types of concavities, including those that have no minima, gave similar results. This argument, in turn, relied on the assumption that concavities without a minimum do not produce parsing. The evidence for this is not strong, as we discuss elsewhere in this review.

With a different approach, it is possible to compare strictly concave and strictly convex elements in a visual search task. Bertamini and Lawson (2006) used strictly convex and strictly concave stimuli: circular objects and circular holes. Holes are useful exactly because a circular hole is a strictly concave stimulus. They did not find any visual search difference and concluded that the best explanation for the existing data is that it is not concavity per se that guides attention in visual search.

Is convexity coded explicitly in the brain?

Some behavioral data have been interpreted in terms of how convexity and concavity are coded in the brain. In a series of elegant studies, Suzuki (2001, 2003; Suzuki & Cavanagh, 1998) has documented what is known as the convexity aftereffect (see Fig. 10). Suzuki concluded that this rapidly adapting contour aftereffect may be mediated by high-level neural units that encode global configurations of orientation




Fig. 9 Visual search for a target that includes a concavity is easier than the search for a target that has no concavities (Hulleman et al., 2000). Interestingly, presence or absence of a curvature minimum did not make a difference

(e.g., convexity and concavity), probably in the inferotemporal cortex. Importantly, this aftereffect is position invariant, is scale invariant, transfers across surface attributes, and is modulated by attention.

Müller, Wilke, and Leopold (2009) studied V4 neurons of rhesus monkeys, using stimuli similar to those of Suzuki. They found not just selectivity for these shapes, but also changes consistent with shape adaptation and concluded that these neurons are involved in the coding of convexity and contribute to shape-specific adaptation.

In terms of brain recordings, much work has been done to characterize how shape is encoded in visual areas. In particular, a series of papers from Connor and collaborators (Brincat & Connor, 2004, 2006; Carlson et al., 2011; Pasupathy & Connor, 2002) have focused on shape processing in macaque area V4 and the inferotemporal cortex. One of their conclusions is that sparse coding is achieved in the midlevel visual cortex by emphasizing representation of acute convex and concave curvature. It is possible that these boundary features play an important role in shape reconstruction. Brincat and Connor (2004) concluded that the evidence supports parts-based shape representation theories that posit explicit coding of structural relations between parts.

In the case of the human visual system, an fMRI study by Haushofer, Baker, Livingstone, and Kanwisher (2008) tested shape adaptation, using convex and concave stimuli. Two

simple shapes were presented sequentially. These were similar in design to those used by Bertamini and Lawson (2008). Haushofer et al. found that the fMRI response in the anterior lateral occipital complex (LOC) for convex was higher for different than that for identical pairs. When the same stimuli were seen as concave, the responses for different and identical pairs were the same. Unlike the work discussed previously in monkeys' V4, the results from Haushofer et al. suggest a priority for convexity in terms of how shape is encoded in the human LOC, a critical region for shape analysis and object recognition.

Conclusions

We started this review by pointing out that contours are informative about solid shape. Research has indeed confirmed that this is the case. For instance, object recognition is equivalent for silhouettes and for shaded images (Hayward, 1998; Hayward, Tarr, & Corderoy, 1999; Lloyd-Jones & Luckhurst, 2002). In an extensive study with silhouettes and contours derived from 260 line drawings of everyday objects (Wagemans et al., 2008), a large number remained almost perfectly identifiable (112 stimuli with >90 % correct identification). However, there were also a substantial number that were almost completely unidentifiable (56 stimuli with <10 % correct identification), probably because of a loss of 3-D structure due to particular perspectives or a loss of distinctive surface markings. We suggested that a critical aspect of how the system uses contours is the extraction of local convexity and concavity information.

On the basis of the available evidence, we can be confident in saying that convexity and concavity information along a contour is something that the visual system computes and represents as part of how shape is analyzed. This information contributes to how solid shapes are perceived. Although some early scholars have discussed convexity and concavity, a clearer picture emerged in the 1980s thanks to the work of Koenderink, Hoffman, and others. However, empirical research continues, and more work is necessary to test how exactly the visual system uses this information. To summarize the literature, we have structured this review in a series of questions. For some the answer is available, but for others the answer will only come in the future with more




Fig. 10 Example of stimuli used to study the convexity aftereffect (Suzuki, 2001). Rapidly adapting contour after effects may be mediated by high-level mechanisms that encode global shape information (e.g., convexity and concavity)

research. Although there may be disagreements on this, here are the answers that seem appropriate to us, given the evidence (see Table 1 again): (1) Does convexity affect figure-ground organization? Yes. (2) Does convexity affect perceived metric depth? Not proven yet. (3) Is convexity coded explicitly in the brain? Yes. (4) Are convexities perceived as “parts”? Yes, or at least they play an important role in part parsing. (5) Are convex and concave extrema (maxima and minima) critical for how shape is perceived? Not proven; at least, the presence of extrema is not as important as the sign of curvature. (6) Is attention directed to convexities, concavities, or neither? Neither—that is, not in any automatic sense that is independent of the context. (7) Are convexities and concavities of a visual hole coded as in its complement (figure)? Not for perceptual tasks, where contour ownership is important; perhaps in memory.

Author Note Address correspondence to Marco Bertamini (M.Bertamini@liv.ac.uk): University of Liverpool, Department of Psychology, Eleanor Rathbone Building, Bedford Street South, Liverpool L69 7ZA, UK. J.W. receives long-term structural funding from the Flemish Government (METH/08/02).

References

- Albrecht, A. R., List, A., & Robertson, L. C. (2008). Attentional selection and the representation of holes and objects. *Journal of Vision*, 8(13), 1–10.
- Asada, H., & Brady, M. (1984). *The curvature primal sketch*. Technical Report, MIT AI memo 758.
- Arnheim, R. (1954). *Art and visual perception: A psychology of the creative eye*. Berkeley: University of California Press.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61, 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.
- Barenholtz, E., & Feldman, J. (2003). Visual comparisons within and between object parts: Evidence for a single-part superiority effect. *Vision Research*, 43, 1655–1666.
- Barenholtz, E., & Feldman, J. (2006). Determination of visual figure and ground in dynamically deforming shapes. *Cognition*, 101(3), 530–544. doi:10.1016/j.cognition.2005.12.00
- Barenholtz, E., & Tarr, M. J. (2008). Visual judgment of similarity across shape transformations: Evidence for a compositional model of articulated objects. *Acta Psychologica*, 128(2), 331–338.
- Baylis, G. C., & Driver, J. (1993). Visual attention and objects: Evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception and Performance*, 19(3), 451–470.
- Baylis, G. C., & Driver, J. (1994). Parallel computation of symmetry but not repetition in single visual objects. *Visual Cognition*, 1, 377–400.
- 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.
- Baylis, G. C., & Driver, J. (2001). Perception of symmetry and repetition within and across visual shapes: Part-descriptions and object-based attention. *Visual Cognition*, 8(2), 163–196.
- Bertamini, M. (2001). The importance of being convex: An advantage for convexity when judging position. *Perception*, 30, 1295–1310.
- Bertamini, M. (2006). Who owns the contour of a hole? *Perception*, 35, 883–894.
- Bertamini, M. (2008). Detection of convexity and concavity in context. *Journal of Experimental Psychology: Human Perception and Performance*, 34(4), 775–789.
- Bertamini, M., & Croucher, C. J. (2003). The shape of holes. *Cognition*, 87(1), 33–54.
- Bertamini, M., & Farrant, T. (2005). Detection of change in shape: An advantage for changes in part structure. *Acta Psychologica*, 120, 35–54.
- Bertamini, M., & Farrant, T. (2006). The perceived structural shape of thin (wire-like) objects is different from that of silhouettes. *Perception*, 35, 1265–1288.
- Bertamini, M., Friedenberg, J., & Argyle, L. (2002). No within-object advantage for detection of rotation. *Acta Psychologica*, 111, 59–81.
- Bertamini, M., Friedenberg, J., & Kubovy, M. (1997). Detection of symmetry and perceptual organization: The way a lock-and-key process works. *Acta Psychologica*, 95, 119–140.
- Bertamini, M., & Hulleman, J. (2006). Amodal completion and visual holes (static and moving). *Acta Psychologica*, 123, 55–72.
- Bertamini, M., & Lawson, R. (2006). Visual search for a figure among holes and for a hole among figures. *Perception & Psychophysics*, 58, 776–791.
- Bertamini, M., & Lawson, R. (2008). Rapid figure-ground responses to stereograms reveal an advantage for a convex foreground. *Perception*, 37, 483–494.
- Bertamini, M., & Mosca, F. (2004). Early computation of contour curvature and part structure: Evidence from holes. *Perception*, 33, 35–48.
- Bertamini, M., Helmy, M. S., & Hulleman, J. (2012). The role of convexity in perception of symmetry and in visual short-term memory. *Quarterly Journal of Experimental Psychology*. doi:10.1080/17470218.2012.717953
- Bertamini, M., Martinovic, J., Wuerger, S. M. (2008). Integration of ordinal and metric cues in depth processing. *Journal of Vision*, 8(2)10, 1–12. <http://journalofvision.org/8/2/10/> doi:10.1167/8.2.10
- Bhatt, R. S., & Bertin, E. (2001). Pictorial cues and three-dimensional information processing in early infancy. *Journal of Experimental Child Psychology*, 80(4), 315–332.
- Bhatt, R. S., Hayden, A., Reed, A., Bertin, E., & Joseph, J. (2006). Infants' perception of information along object boundaries: Concavities versus convexities. *Journal of Experimental Child Psychology*, 94(2), 91–113. doi:10.1016/j.jecp.2006.01.002
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94(2), 115–147.
- Braunstein, M. L., Hoffman, D. D., & Saidpour, A. (1989). Parts of visual objects: An experimental test of the minima rule. *Perception*, 18(6), 817–826.
- Brincat, S. L., & Connor, C. E. (2004). Underlying principles of visual shape selectivity in posterior inferotemporal cortex. *Nature Neuroscience*, 7(8), 880–886.
- Brincat, S., & Connor, C. (2006). Dynamic shape synthesis in posterior inferotemporal cortex. *Neuron*, 49, 17–24.
- Burge, J., Fowlkes, C. C., & Banks, M. S. (2010). Natural-scene statistics predict how the figure-ground cue of convexity affects human depth perception. *The Journal of Neuroscience*, 30, 7269–7280.
- Burge, J., Peterson, M. A., Palmer, S. E. (2005). Ordinal configural cues combine with metric disparity in depth perception. *Journal of Vision*, 5(6)5, 534–542.
- Carlson, E. T., Rasquinha, R. J., Zhang, K., & Connor, C. E. (2011). A sparse object coding scheme in area V4. *Current Biology*, 21, 288–293. doi:10.1016/j.cub.2011.01.013
- Cohen, E. H., & Singh, M. (2007). Geometric determinants of shape segmentation: Tests using segment identification. *Vision Research*, 47, 2825–2840.

- Cohen, E. H., Barenholtz, E., Singh, M., & Feldman, J. (2005). What change detection tells us about the visual representation of shape. *Journal of Vision*, *5*, 313–321.
- Corrow, S., Granrud, C. E., Mathison, J., & Yonas, A. (2011). Six-month-old infants perceive the hollow face illusion. *Perception*, *40*, 1376–1383.
- De Winter, J., & Wagemans, J. (2006). Segmentation of object outlines into parts: A large-scale integrative study. *Cognition*, *99*, 275–325.
- De Winter, J., & Wagemans, J. (2008). Perceptual saliency of points along the contour of everyday objects: A large-scale study. *Perception & Psychophysics*, *70*, 50–64.
- Dhandapani, R., & Kimia, B. B. (2002). Role of scale in partitioning shape. In *International conference on image processing* (pp. II-565).
- Fantoni, C., Bertamini, M., & Gerbino, W. (2005). Contour curvature polarity and surface interpolation. *Vision Research*, *45*, 1047–1062.
- Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review*, *112*, 243–252.
- Fowlkes, C. C., Martin, D. R., & Malik, J. (2007). Local figure-ground cues are valid for natural images. *Journal of Vision*, *7*(8), 2. doi:10.1167/7.8
- Friedenberg, J., & Bertamini, M. (2000). Contour symmetry detection: The influence of axis orientation and number of objects. *Acta Psychologica*, *105*, 107–118.
- Fulvio, J. M., & Singh, M. (2006). Surface geometry influences the shape of illusory contours. *Acta Psychologica*, *123*(1–2), 20–40.
- Gibson, B. S. (1994). Visual attention and objects: One versus two or convex versus concave? *Journal of Experimental Psychology: Human Perception and Performance*, *20*(1), 203–207.
- Gillam, B. J., Anderson, B. L., Rizwi, F. (2009). Failure of facial configurational cues to alter metric stereoscopic depth. *Journal of Vision*, *9*(1)3, 1–5, doi:10.1167/9.1.3
- Haushofer, J., Baker, C. I., Livingstone, M. S., & Kanwisher, N. (2008). Privileged coding of convex shapes in human object-selective cortex. *Journal of Neurophysiology*, *100*, 753–762.
- Hayden, A., Bhatt, R. S., Kangas, A., & Zieber, N. (2011). Parts, cavities, and object representation in infancy. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(1), 314–317.
- Hayward, W. G. (1998). Effects of outline shape in object recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 427–440.
- Hayward, W. G., Tarr, M. J., & Corderoy, A. K. (1999). Recognizing silhouettes and shaded images across depth rotation. *Perception*, *28*, 1197–1215.
- Hebb, D. (1949). *The organization of behavior*. New York: Wiley.
- Hill, H., & Bruce, V. (1993). Independent effects of lighting, orientation, and stereopsis on the hollow-face illusion. *Perception*, *22*, 887–897.
- Hoffman, D. D., & Richards, W. (1984). Parts of recognition. *Cognition*, *18*, 65–96.
- Hoffman, D. D., & Singh, M. (1997). Saliency of visual parts. *Cognition*, *63*, 29–78.
- Horowitz, T. S., & Kuzmova, Y. (2011). Can we track holes? *Vision Research*, *51*(9), 1013–1021. doi:10.1016/j.visres.2011.02.009
- Howard, I. P. (1996). Alhazen's neglected discoveries of visual phenomena. *Perception*, *25*, 1203–1217.
- Hulleman, J., & Olivers, C. N. L. (2007). Concavities count for less in symmetry perception. *Psychonomic Bulletin & Review*, *14*, 1212–1217.
- Hulleman, J., te Winkel, W., & Boselie, F. (2000). Concavities as basic features in visual search: Evidence from search asymmetries. *Perception & Psychophysics*, *62*(1), 162–174.
- Humphreys, G., & Müller, H. (2000). A search asymmetry reversed by figure-ground assignment. *Psychological Science*, *11*(3), 196–201.
- Jacobs, D. W. (1996). Robust and efficient detection of salient convex groups. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *18*(1), 23–37.
- Johnston, A., Hill, H., & Carman, N. (1992). Recognising faces: Effects of lighting direction, inversion and brightness reversal. *Perception*, *21*, 365–375.
- Kanizsa, G., & Gerbino, W. (1976). Convexity and symmetry in figure-ground organization. In M. Henle (Ed.), *Vision and artifact* (pp. 25–32). New York: Springer Publishing Co.
- Kennedy, J. M., & Domander, R. (1985). Shape and contour: The points of maximum change are least useful for recognition. *Perception*, *14*(3), 367–370.
- Kim, S.-H., & Feldman, J. (2009). Globally inconsistent figure/ground relations induced by a negative part. *Journal of Vision*, *9*(10), 8.1–13. doi:10.1167/9.10
- Koenderink, J. (1984). What does the occluding contour tell us about solid shape? *Perception*, *13*, 321–330.
- Koenderink, J. (1990). *Solid shape*. Cambridge: MIT Press.
- Koenderink, J. J., & van Doorn, A. J. (1982). The shape of smooth objects and the way contours end. *Perception*, *11*, 129–137.
- Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2001). Ambiguity and the “mental eye” in pictorial relief. *Perception*, *30*, 431–448.
- Koffka, K. (1935). *Principles of Gestalt psychology*. London: Routledge and Kegan.
- Kogo, N., Galli, A., & Wagemans, J. (2011). Switching dynamics of border ownership: A stochastic model for bi-stable perception. *Vision Research*, *51*, 2085–2098.
- Kogo, N., & Wagemans, J. (2012). The ‘side’ matters: How configurability is reflected in completion. *Cognitive Neuroscience*, in press.
- Koning, A., & Wagemans, J. (2009). Detection of symmetry and repetition in one and two objects. Structures versus strategies. *Experimental Psychology*, *56*(1), 5–17. doi:10.1027/1618-3169.56.1
- Lamote, C., & Wagemans, J. (1999). Rapid integration of contour fragments: From simple filling-in to parts-based shape description. *Visual Cognition*, *6*(3), 345–361.
- Langer, M. S., & Bülthoff, H. H. (2001). A prior for global convexity in local shape-from-shading. *Perception*, *30*, 403–410.
- Lappin, J. S., & Craft, W. D. (2000). Foundations of spatial vision: From retinal images to perceived shapes. *Psychological Review*, *107*, 6–38.
- Lappin, J. S., Norman, J. F., & Phillips, F. (2011). Fechner, information, and shape perception. *Attention, Perception, & Psychophysics*, *73*, 2353–2378.
- Lim, I. S., & Leek, E. C. (2012). Curvature and the visual perception of shape: Theory on information along object boundaries and the minima rule revisited. *Psychological Review*, *119*(3), 668–677.
- Liu, Z., Jacobs, D. W., & Basri, R. (1999). The role of convexity in perceptual completion: Beyond good continuation. *Vision Research*, *39*, 4244–4257.
- Lloyd-Jones, T. J., & Luckhurst, L. (2002). Outline shape is a mediator of object recognition that is particularly important for living things. *Memory & Cognition*, *30*, 489–498.
- Mach, E. (1959). *The analysis of sensations, and the relation of the physical to the psychological*. New York: Dover. original work published in German in 1897.
- Marr, D., & Nishihara, H. K. (1978). Representation and recognition of spatial-organization of 3-dimensional shapes. *Proceedings of the Royal Society of London B*, *200*(1140), 269–294.
- Müller, K.-M., Wilke, M., & Leopold, D. A. (2009). Visual adaptation to convexity in macaque area V4. *Neuroscience*, *161*(2), 655–662. doi:10.1016/j.neuroscience.2009.03.070
- Nelson, R., Thierman, J., & Palmer, S. E. (2009). Shape memory for intrinsic versus accidental holes. *Attention, Perception, & Psychophysics*, *71*, 200–206.

- 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: MIT Press.
- Palmer, S. E., Davis, J., Nelson, R., & Rock, I. (2008). Figure - ground effects on shape memory for objects versus holes. *Perception, 37*, 1569–1586.
- Panis, S., De Winter, J., Vandekerckhove, J., & Wagemans, J. (2008). Identification of everyday objects on the basis of fragmented outline versions. *Perception, 37*, 271–289.
- Panis, S., & Wagemans, J. (2009). Time-course contingencies in perceptual organization and identification of fragmented object outlines. *Journal of Experimental Psychology: Human Perception and Performance, 35*(3), 661–687.
- Pao, H.-K., Geiger, D., Rubin, N. (1999). Measuring convexity for figure/ground separation. In 7th IEEE International Conference on Computer Vision, Kerkyra, Greece, 1999.
- Pao, H.-K., & Geiger, D. (2001). A continuous shape descriptor by orientation diffusion. *Lecture Notes in Computer Science, 2134*, 544–559.
- Pasupathy, A., & Connor, C. E. (2002). Population coding of shape in area v4. *Nature Neuroscience, 5*, 1332–1338.
- Peterson, M. A., & Gibson, B. S. (1993). Shape recognition contributions to figure-ground organization in three-dimensional displays. *Cognitive Psychology, 25*, 383–429.
- Peterson, M. A., & Salvagio, E. (2008). Inhibitory competition in figure-ground perception: Context and convexity. *Journal of Vision, 8*(16)4, 1–13, doi:10.1167/8.16.4
- Peterson, M. A., & Skow, E. (2008). Inhibitory competition between shape properties in figure-ground perception. *Journal of Experimental Psychology: Human Perception and Performance, 34*(2), 251–267.
- Preparata, F., & Shamos, I. (1985). *Computational geometry: An introduction* (2nd ed.). New York: Springer-Verlag.
- 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 & Psychophysics, 63*, 476–489.
- Rosin, P. L. (2000). Shape partitioning by convexity. *IEEE Transactions on Systems Man and Cybernetics Part A, 30*, 202–210.
- Sabra, A. I. (1989). *The optics of Ibn al-Haytham: Books I-III: On direct vision*. London: Warburg Institute. English translation and commentary by Sabra, A.I.
- Sakai, K., & Inui, T. (2002). A feature-segmentation model of short-term visual memory. *Perception, 31*(5), 579–590.
- Salapatek, P. (1968). Visual scanning of geometric figures by the human newborn. *Journal of Comparative and Physiological Psychology, 66*, 247–258.
- Siddiqi, K., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception, 25*, 399–424.
- Singh, M., & Feldman, J. (2012). Principles of contour information: Reply to Lim and Leek (2012). *Psychological Review, 119*(3), 678–683.
- Singh, M., & Hoffman, D. (1998). Part boundaries alter the perception of transparency. *Psychological Science, 9*, 370–378.
- 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 & Psychophysics, 61*, 636–660.
- Suzuki, S. (2001). Attention-dependent brief adaptation to contour orientation: A high-level after effect for convexity? *Vision Research, 41*, 3883–3902.
- Suzuki, S. (2003). Attentional selection of overlapped shapes: A study using brief shape after effects. *Vision Research, 43*(5), 549–561.
- Suzuki, S., & Cavanagh, P. (1998). A shape-contrast effect for briefly presented stimuli. *Journal of Experimental Psychology: Human Perception and Performance, 24*(5), 1–27.
- Treder, M. S., & van der Helm, P. A. (2007). Symmetry versus repetition in cyclopean vision: A microgenetic analysis. *Vision Research, 47*, 2956–2967.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review, 95*(1), 15.
- Vaina, L. M., & Zlateva, S. D. (1990). The largest convex patches: A boundary-based method for obtaining object parts. *Biological Cybernetics, 62*, 235–236.
- Vandekerckhove, J., Panis, S., & Wagemans, J. (2007). The concavity effect is a compound of local and global effects. *Perception & Psychophysics, 69*, 1253–1260.
- van Doorn, A. J., Koenderink, J. J., Wagemans, J. (2011). Light fields and shape from shading. *Journal of Vision, 11*(3)21, 1–21. doi:10.1167/11.3.21
- Vecera, S. P., Flevaris, A. V., & Filapek, J. C. (2004). Exogenous spatial attention influences figure-ground assignment. *Psychological Science, 15*, 20–26.
- Wagemans, J., De Winter, J., Op de Beeck, H., Ploeger, A., Beckers, T., & Vanroose, P. (2008). Identification of everyday objects on the basis of silhouette and outline versions. *Perception, 37*(2), 207–244.
- Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin, 138*(6), 1172–1217. doi:10.1037/a0029333
- Wagemans, J., van Doorn, A. J., & Koenderink, J. J. (2010). The shading cue in context. *Perception, 1*, 159–178. doi:10.1068/i0401
- Wang, X., & Burbeck, C. A. (1998). Scaled medial axis representation: Evidence from position discrimination task. *Vision Research, 38*, 1947–1959.
- 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.
- Wolfe, J., & Bennett, S. C. (1997). Preattentive object files: Shapeless bundles of basic features. *Vision Research, 37*, 25–43.
- Xu, Y., & Singh, M. (2002). Early computation of part structure: Evidence from visual search. *Perception & Psychophysics, 64*, 1039–1054.
- Zhou, K., Luo, H., Zhou, T., Zhuo, Y., & Chen, L. (2010). Topological change disturbs object continuity in attentive tracking. *Proceedings of the National Academy of Sciences, 107*(50), 21920–21924.
- Zunic, J. D., & Rosin, P. L. (2002). A convexity measure for polygons. In *British machine vision conference, Cardiff, 2002*. British Machine Vision Association.