



Segmentation of object outlines into parts: A large-scale integrative study

Joeri De Winter, Johan Wagemans*

Department of Psychology, University of Leuven, Tiensestraat 102, B-3000 Leuven, Belgium

Received 15 June 2002; revised 17 March 2005; accepted 18 March 2005

Abstract

In this study, a large number of observers ($N=201$) were asked to segment a collection of outlines derived from line drawings of everyday objects ($N=88$). This data set was then used as a benchmark to evaluate current models of object segmentation. All of the previously proposed rules of segmentation were found supported in our results. For example, minima of curvature (i.e. locations along the contour where negative curvature takes an extreme value) were often used as segmentation points. The second point of a pair connected by a segmentation line often depended on more global shape characteristics such as proximity, collinearity, symmetry, and elongation. Based on these results, a framework is presented in which all of the previously proposed (and now empirically validated) segmentation rules or rules for part formation are integrated.

© 2005 Published by Elsevier B.V.

Keywords: Segmentation rule; Perceptual unit; Object contour; Curvature; Collinearity; Symmetry; Spatial relation; Spatial scale; Shape

When people look at objects, they often spontaneously segment them into parts (Biederman, 1987; Cave & Kosslyn, 1993; Hoffman & Richards, 1984; Hoffman & Singh, 1997; Kurbat, 1994a; Lamberts & Freeman, 1999; Lamote & Wagemans, 1999; Scholl, 2001; Singh, Seyranian & Hoffman, 1996; Van Lier & Wagemans, 1998). Parts seem to have a special status as ‘perceptual units’ to determine visual attention (e.g. Vecera,

* Corresponding author. Tel.: +32 16 32 59 69; fax: +32 16 32 60 99.

E-mail address: johan.wagemans@psy.kuleuven.be (J. Wagemans).

Behrmann, & Filapek, 2001; Vecera, Behrmann, & McGoldrick, 2000; Xu & Singh, 2002) and object encoding (e.g. Barenholtz, Cohen, Feldman, & Singh, 2003; Barenholtz & Feldman, 2003; Bertamini & Croucher, 2003; Bertamini & Farrant, 2005; Bertamini & Mosca, 2004). Parts themselves are often cognitively accessible by their name (e.g. Landau, Smith, & Jones, 1998). For instance, a hand consists of a palm and fingers, and each finger can be named and still be divided further. Segmenting objects into parts has many advantages. Biederman (1987) and Hoffman and Richards (1984) have argued that parts supply a first index to a possible object category or identity and that objects can thus be recognized on the basis of their parts. In many cases, a partly occluded object can still be identified by the remaining visible parts. Partial occlusion of object occurs frequently in everyday life. For instance, one usually cannot see the back of an object, unless it is transparent (Van Lier & Wagemans, 1999). A structural object description consisting of parts and their spatial relations is also much more stable and flexible. It is better suited for handling deformations of semi-rigid or non-rigid shapes than template descriptions (Pinker, 1984). In sum, object segmentation into parts occurs frequently and it appears quite useful.

Several models of object segmentation exist, yet it is still not known in detail how parts are formed. The existing models can generally be divided into two groups. One group postulates the existence of a limited set of basic shapes like generalized cylinders (Marr, 1982) or geons (Biederman, 1987). This predefined set of shape primitives can be compared to the letters of the alphabet. Using the letters (parts), one can create an unlimited set of words (objects). In the second group of theories, the parts are not predefined but geometric rules are used to find the parts in the image. These rules are shape-based and often only the outline of the shape is used. For example, some models are based on convexity in general (Rosin, 2000) or look explicitly for the largest convex patches (Vaina & Zlateva, 1990), while other models are more axis-based (e.g. Burbeck & Pizer, 1995; Rom & Medioni, 1993). The geometric properties that are most typically used, however, are special points of curvature along the outline contour, like local minima and maxima of curvature. In some models, these local cues are used in combination with more global geometric properties, like the distance between possible segmentation points or the collinearity of contour fragments.

Despite the importance of object segmentation into parts and the large number of theoretical models, not so many experimental studies have been performed on how humans segment object shapes into parts. The few available studies have tested only a small number of subjects and they have all used a limited number of stimuli, often meaningless outline shapes. In this paper, we will present a large-scale study of how humans segment outline shapes derived from line drawings of existing objects (i.e. the standardized set by Snodgrass & Vanderwart, 1980) and we evaluate how the most frequently observed segmentations of everyday objects compare to the known boundary-based segmentation models. We will also explore possible top-down influences on segmentation. Before we present our own study, we will briefly review the most prominent models of boundary-based object segmentation and the experiments conducted to evaluate these models (see also Singh & Hoffman, 2001).

1. Boundary-based object segmentation

1.1. The minima rule (Hoffman & Richards, 1984)

Hoffman and Richards (1984) proposed that objects are segmented according to the *minima rule*. Because in their view object segmentation into parts precedes object recognition, it must be based on very generic regularities that hold for almost any object. These regularities are found in the transversality and singularity principle. The transversality principle entails that a concave crease on the surface of an object is a very likely candidate for the segmentation of two forms. The singularity principle states that a three-dimensional (3D) concave crease almost always results in a two-dimensional (2D) concave discontinuity. This concave discontinuity is then further generalized to concave regions (i.e. contour fragments with negative curvature). Thus, 2D outline shapes can be segmented at concave regions. These regions have a local minimum of negative curvature ($m-$ for short, see Fig. 1), hence the name ‘minima rule’. The bounding contours of the resulting parts have predominantly positive curvature, so these parts are called ‘positive parts’. Note that this implies a different role of convex (positively curved) and concave (negatively curved) contour fragments, namely as ‘thing-like’ versus ‘glue-like’, respectively (e.g. Feldman & Singh, 2005; Koenderink & van Doorn, 1982). Note also that this creates an asymmetry which reverses when figure and background are reversed:

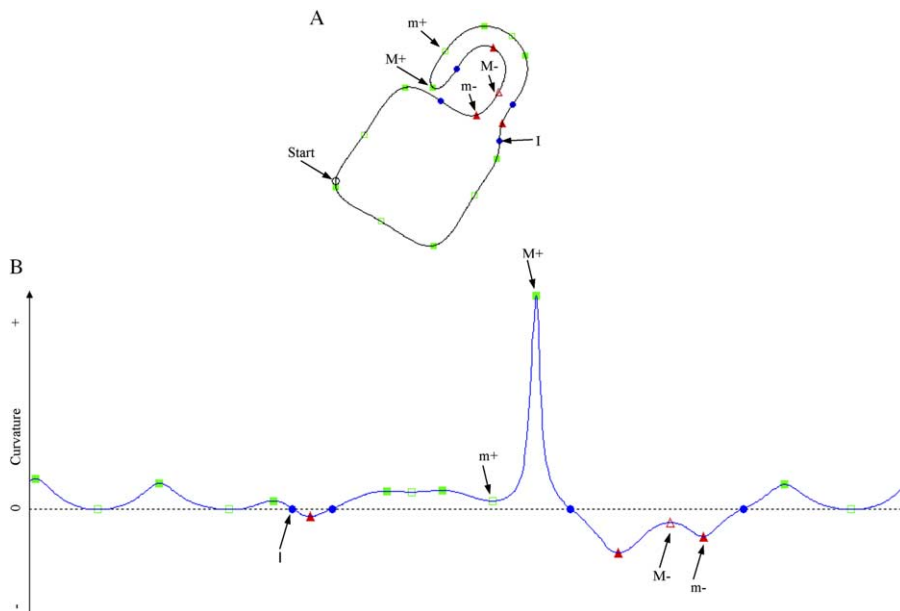


Fig. 1. An example of an outline shape (A) and its curvature graph (B) with curvature singularities marked on both: green filled squares represent the positive maxima ($M+$), green open squares the positive minima ($m+$), red filled triangles the negative minima ($m-$), red open triangles the negative maxima ($M-$) and blue circles are inflections (I). The curvature graph starts on the black dot and follows the outline in counter-clockwise direction. The minima rule segments the outline at $m-$ points (i.e. yielding predominantly convex parts).

Concave becomes convex and vice versa (e.g. Bertamini & Croucher, 2003; Driver & Baylis, 1995, 1996). Note, finally, that this rule does not state how a part itself is formed; it only defines possible borders of a part. Although often two consecutive *m*- segmentation points are connected to form a part, a strict application of this rule can lead to mistakes (e.g. Kurbat, 1994a; Singh, Seyranian, & Hoffman, 1999; Xu & Singh, 2002).

Hoffman and Richards (1984) also mentioned points of maximum positive curvature (*M+* for short, see Fig. 1) as a second type of segmentation point, bordering a so-called 'negative part', the hole that is left behind when a part is removed from an object. This type of segmentation point is not discussed in great detail (probably because they are fairly rare) but Hoffman and Richards (1984) do argue clearly that a part will never be bordered by an *m*- and an *M+* point. Other special points of curvature or curvature singularities (see Fig. 1) are positive minima (*m+* for short, where the positive curvature is locally minimal), negative maxima (*M-* for short, where the negative curvature is locally maximal), and inflections (*I* for short, where the curvature goes through zero when it changes from positive to negative or vice versa). In the study reported here, the focus of the analyses will be on these curvature singularities.

Braunstein, Hoffman, and Saidpour (1989) conducted three experiments to evaluate the minima rule. The stimuli ($N=8$) consisted of meaningless arbitrary surfaces of revolution filled with moving dots to give the illusion of a 3D shape so that figure and ground assignment was clear. The revolving surface was bordered by the generating curve either on the left or on the right so that negative minima (*m-*) become positive maxima (*M+*) and vice versa in different conditions. Subjects ($N=13$) had to choose the part that originated from the stimuli. The first experiment presented four response alternatives: Each time two parts (a positive and a negative one) were derived from the 3D shape and two were not. From the 75% correct answers, 65% consisted of parts flanked by minima (also called 'minima parts'). An analysis of variance (ANOVA) showed a significant effect of part type (positive versus negative) and not of object. These results were regarded as support for the minima rule. In the second experiment ($N=17$), only two response alternatives were presented, only one of which originated from the presented stimulus (either a positive or a negative part). Now, only 40% were correct answers and part type no longer had a significant effect. Although this result seemed to go against the minima rule, the authors gave the following alternative explanation. Because a correct minima part is only available in 50% of the cases, almost half of the subjects reported that they sometimes intentionally reversed figure and ground so that positive parts became negative parts. In the third and final experiment ($N=8$), subjects had to mark perceived part boundaries. There was an effect of boundary type (negative minima and positive maxima) and not of object. Inflections were not used as segmentation points. As a general conclusion, Braunstein et al. (1989) stated that the minima rule was supported by their experimental data.

Indirect empirical support for the minima rule of segmentation also comes from experiments demonstrating that concavities are basic features in visual search (e.g. Elder & Zucker, 1993, 1998; Hulleman, Te Winkel, & Boselie, 2000; Humphreys & Müller, 2000; Xu & Singh, 2002) and in detection of shape changes (e.g. Barenholtz et al., 2003; Bertamini & Farrant, 2005; Lamote & Wagemans, 1999).

1.2. Limbs and necks (Siddiqi, Tresness, & Kimia, 1996)

The model of Siddiqi et al. (1996) builds further on the minima rule and does specify how to determine the parts themselves. To make this possible, an interaction with the global shape is proposed. *Limbs* are formed when two *m*- points can be connected so that the part line, at least on one side, forms a good continuation with the outline (see Fig. 2A). Siddiqi et al. (1996) provided several possible definitions of good continuation (like co-circularity) and included a rather arbitrarily chosen maximum tolerance value (Siddiqi & Kimia, 1995), which was not based on empirical observations. *Necks* are produced when a maximally inscribed circle is also a local minimum of the diameter (see Fig. 2B). The segmentation points are not necessarily positioned on *m*- points.

Empirical support for the limbs and neck model comes from experiments conducted by the authors themselves (Siddiqi & Kimia, 1995). Here, subjects had to draw segmentation lines on 15 meaningless silhouettes and on silhouettes of six animals and of one tree leaf, all presented on a computer screen. In the first experiment ($N=5$), they measured intra-subject consistency, which was high: 87% for meaningless shapes and 72% for biological forms. (How broad the tolerance range was for a segmentation point to be considered the same has not been specified in the article.) In the second experiment ($N=14$), inter-subject consistency was also measured. For meaningless stimuli, 81% of the part lines were in agreement for 11 out of the 14 subjects, while this decreased to 60% for the biological forms. From these results, they concluded that the notion of a part is intuitive and perceptually clear.

The third and final experiment ($N=14$) addressed the similarity between the perceived parts and the parts proposed in the model. Three results stand out. (1) A large number of perceived parts are also parts proposed in the model: 90% for the meaningless shapes and 67% for the biological forms. Two explanations were given for the lower similarity for biological forms: (a) Top-down knowledge can augment the absence of perceptual cues, leading to over-segmentation. Conversely, it is possible that perceptual cues are overridden by top-down knowledge, leading to under-segmentation. (b) Parts that are not rigidly connected to the object can move in such a way that perceptual evidence for a part border is destroyed (e.g. the trunk of an elephant) and thus becomes less clear. (2)




Fig. 2. Examples of a typical limb (A) and neck part (B) (adapted from Siddiqi et al., 1996). The dashed line represents the part cut. The circle in the neck represents the maximally inscribed circle.

About half of the part lines proposed by the model were also perceived by the subjects: 58% for meaningless stimuli and 49% for biological forms. Siddiqi et al. (1996) explained these lower frequencies by the fact that subjects were free in the number of part lines they drew. For a number of reasons, fewer lines were drawn than necessary according to the model (e.g. curved lines were not allowed, some supposedly meaningless shapes were perceived as existing 3D objects, etc.). (3) The deviation between proposed and perceived part lines was generally less than 5 pixels and less than 10° in orientation. Overall, it was concluded that there was a high consistency in segmentation and a high overlap between perceived and proposed neck and limb segmentation lines.

In a more recent paper, Dhandapani & Kimia (2002) improved this model further in several ways. They added the notion of salience to resolve numerous conflicting limb and neck hypotheses. Noisy shapes were analyzed at different spatial scales and a multiple-scale partitioning scheme was implemented in which the strongest negative minima at the coarsest scale are used to resolve conflicts at finer spatial scales. To allow for non-straight part lines, they employed a so-called Euler Spiral as the part curve (interpolating two tangent point pairs by minimizing the total curvature variation). These improvements generally led to fewer parts and parts being more consistent with the previously reported human partitioning data (Siddiqi et al., 1996).

1.3. The short-cut rule (Singh, Seyranian, et al., 1999)

Because the limbs and neck model of Siddiqi et al. (1996) does not always detect intuitive parts, Singh, Seyranian, et al. (1999) proposed an augmentation to the minima rule. This *short-cut rule* (see Fig. 3) specifies that the human visual system prefers to connect segmentation points that are in close proximity to form a part. Proximity is an important factor determining perceptual organization in other contexts too and it is easy to quantify (e.g. Kubovy, Holcombe, & Wagemans, 1998; Kubovy & Wagemans, 1995). The line connecting two segmentation points (i.e. the part line) has to adhere to several restrictions. (1) It must be a straight line. (2) It must cross a local axis of symmetry. (3) It must connect two points on the outline, so that (4) at least one point has minimal negative curvature. The length of the part cut influences the salience of the part. Part salience can be understood as the degree in which the part ‘parts’ itself from the body of the object. The more salient the resulting part is, the better it functions as a good first index into memory




Fig. 3. Examples of the preferred parts according to the short-cut rule (adapted from Singh, Seyranian, et al., 1999). The dashed lines represent the part lines. (A) is an example of a typical cross-shape or + form and (B) is an example of a typical L shape.

and the easier object recognition becomes. In a related paper by two of the same authors (Hoffman & Singh, 1997), three additional features were considered important for part salience: the area of the part relative to the object, the part-boundary strength, and the degree of part protrusion (for more details, see Hoffman & Singh, 1997). However, these other factors determining part salience are not incorporated in the short-cut rule.

It is useful to point out that the short-cut rule is not all that different from the rule leading to a neck. In both rules, part lines connect two points on the outline and they are more attractive when they bridge a shorter distance. In both cases, they also cross a local symmetry axis. Although this is not stated explicitly by Siddiqi et al. (1996), it follows from their definition of a neck and the definition of a local symmetry axis (e.g. Blum, 1973). The major theoretical distinction is that the definition of a neck does not necessarily involve negative minima but even this does not create a big difference in practice: Necks are often formed at negative minima and even according to the short-cut rule, a narrow passage in itself is enough to form a part line in cases when salient negative minima are absent.

Singh, Seyranian, et al. (1999) presented five similar experiments with crosses (+ forms, as in Fig. 3A) and elbows (L shapes, as in Fig. 3B). By changing a number of proportions in these shapes (e.g. widening one dimension), they concluded that in most cases the preferred part lines are indeed the shortest lines. In the + forms, two minima are connected, whereas in the L shapes an *m*- is connected to a point which is not an *m*-. Furthermore, they noticed that it is even possible that none of the segmentation points are negative minima. In cases where the negative minima are not salient, it is possible that the short-cut rule pulls the segmentation points away from nearby negative minima. With highly salient negative minima, this is not possible. However, for reasons not known to us, this interaction with salience has not been included in the definition of the short-cut rule.

Additional criticisms have been formulated by Rosin (2000). He argued that many different methods for finding local symmetry axes have been proposed in the past (for a review, see Kurbat, 1994b) but Singh, Seyranian, et al. (1999) have not specified which one to use. Rosin also gave examples in which the shortest cut does not cross the local symmetry axis but does form an intuitive part, and other examples where the length of the part line is the same but the salience of the parts is clearly different. To handle these problems, Rosin (2000) proposed an alternative way of calculating the salience of a part, incorporating all known influences into one single measure. Finally, in line with previous observations, he also stressed that a straight line is not always the most intuitive part line.

1.4. Conclusion

Although we have not tried to give an exhaustive overview, we think it is clear that each segmentation model has its limitations and no unified theory of segmentation based on the boundary of an object currently exists. On the one hand, there are always examples of shape segmentation following some rule that do not correspond to intuitive parts, and certain intuitive parts cannot be derived from any existing segmentation model. In many cases, the original formulations of the boundary segmentation rules (e.g. minima rule) or rules to obtain part cuts (e.g. necks or limbs) appear to be too strict (Singh & Hoffman, 2001). On the other hand, several ideas are common in most models in the literature. First,

most models include m - points as possible segmentation points to border the part. This does not mean that other special points (like $M+$ and I), or any other point on the outline for that matter, are not possible as segmentation points. It only means that the latter points are less likely as candidate segmentation points. Second, a short part line is preferred to a longer one. Third, recent work explicitly acknowledges the interaction between local and global factors. For example, Siddiqi et al. (1996) and Singh, Seyranian, et al. (1999) have stressed that part cuts cross local symmetry axes and Singh & Hoffman (2001) have explicitly argued that boundary strength can interact with the short-cut rule (e.g. one factor compensating for the other).

Another common characteristic of the current literature is that the empirical basis to validate the different models is rather scarce. While the minima-rule has received the most experimental support, it is not a segmentation model for parts (but only for segmentation points). Also, most of the research has used only a small number of subjects and stimuli, and most of these were meaningless outline shapes. As an exception, Siddiqi et al. (1996) have used a mixture of fifteen meaningless and seven biological shapes but no manmade objects, for instance.

2. Aims and scope of the present study

2.1. Aims

In the light of the above state-of-the-art, we thus believe that the field would benefit from a study in which a large set of outlines is segmented by a large number of subjects. In such a benchmark data set, the relative strengths and weaknesses of the different models would be easier to evaluate, and even relatively rare types of parts could reveal themselves. To provide such a benchmark data set is the most important goal of this paper.

We investigate how people segment outlines of existing objects into parts. Our outline shapes are all based on line-drawings of everyday objects from the stimulus set developed by Snodgrass & Vanderwart (1980). This gives us the opportunity to explore the influence of knowledge on visual segmentation. This is done by selecting outline shapes that are either easy or difficult to recognize (based on previously obtained results). Although some authors mention the possible influence of object knowledge on segmentation (e.g. Schyns, Goldstone, & Thibaut, 1998; Schyns & Murphy, 1994; Siddiqi et al., 1996; Singh, Seyranian, et al., 1999), it has not been examined systematically in previous research. Hence, our study must be considered exploratory in this respect. In line with the growing attention for global factors complementing the early focus on special points along the outline contour (curvature singularities), our study will examine more global factors too. For example, the distance between segmentation points, the collinearity of contour fragments, and the symmetry of the whole object, or of the resulting parts, will all turn out to affect segmentation. Hence, a second aim of our study was to analyze the influence of all known or previously proposed segmentation factors including curvature singularities as well as more global shape factors and cognitive influences.

Third, based on the results from this large-scale study and on our own synthesis of the relevant principles from the literature, we will develop an integrative framework of shape

segmentation in which all of the empirically validated factors (local curvature singularities, global shape properties, and top-down influences) are incorporated. The ambitions of the current version of the framework are modest: It is a purely qualitative framework in which the factors stressed by previous segmentation models are allowed to interact with each other, but the relative weight of each factor and the parameters needed to determine their interactions quantitatively have not been specified.

2.2. Scope

Our study is also limited by the specific choices that we have made regarding its scope. The present study joins the dominant research tradition reviewed above in starting from image-based geometric rules rather than from predefined volumetric object parts. We believe this is justified by what we currently know about object recognition and the way this might be implemented in the brain. Based on his review of the available literature, [Bar \(2003\)](#) postulated that object recognition takes place in two phases. A first rapid bottom-up analysis of a coarse representation of the input (not unlike a blurred silhouette) in the prefrontal cortex serves as a first initial guess, which is then back-projected to the temporal cortex to be integrated with further bottom-up analyses at finer spatial scales (see also [Hochstein & Ahissar, 2002](#)). To the extent that parts are relevant to object recognition, it seems very unlikely that parts would be detected by predefined volumetric shapes. Indeed, this would require separate feedforward and feedback loops for the generation and testing of hypotheses regarding all possible object parts. In contrast, it seems much more plausible if the first bottom-up analysis of the coarse input representation included a primitive, geometry-based part decomposition, which could yield a match with a single structural part-based object description (possibly, at a categorical level) to be tested in the subsequent feedback loop. Given the short time needed to recognize objects and the cortical processing data referenced by [Bar \(2003\)](#), not more than one or two cycles seem to be possible and thus bottom-up geometric properties must play a major role (see also [Brincat & Connor, 2004](#); [Pasupathy & Connor, 2002](#)). Additional support for the early computation of part structure has been provided recently (e.g. [Barenholtz et al., 2003](#); [Bertamini & Farrant, 2005](#); [Bertamini & Mosca, 2004](#); [Xu & Singh, 2002](#)).

In line with this tradition relying on geometric rules, we will use 2D object outlines rather than real 3D objects or 2D line-drawings. Although our 3D world is rich in visual cues (like color, motion, shape, depth, textures, etc.) that help us to recognize objects, we can easily recognize most objects solely by their shape, even in simple 2D representations (e.g. [Hoffman & Richards, 1984](#); [Kovács et al., 2003](#); [Wagemans et al., 2005](#)). It has also been shown empirically that humans do see parts in outline shapes (e.g. [Singh, Seyranian, et al., 1999](#)). Furthermore, most segmentation models can handle only outline shapes, not fully rendered 3D objects. In sum, the reduction from ecological 3D objects to 2D outline stimuli has many advantages and few disadvantages: In addition to providing us with a tool to compare recognizable and unrecognizable shapes with similar geometric properties, working with 2D outlines derived from natural objects keeps the complexity of the geometric analyses within reasonable limits, while at the same time being much more relevant to everyday object perception than most of the limited stimulus sets used so far in the literature. Of course, as always with experimental work, we will have to be

careful in avoiding unjustified generalizations from our results, obtained within the confines of the above scope, to everyday object segmentation and recognition.

3. Methods

3.1. Subjects

Two hundred and one subjects, all first-year psychology students at the University of Leuven, participated in this study as a mandatory component of their curriculum. It took place in a large movie theatre used as a lecture hall, following two hours of lectures on the history of philosophy and metaphysics. Our study was only part of a longer session of several (unrelated) paper-and-pencil tests and experiments. It is important to note that these conditions are not ideal to motivate subjects for an experiment. On the other hand, it is an efficient method to quickly get a large set of data (which can then be filtered to remove the clearly unreliable data).

3.2. Materials

The stimulus set consisted of 88 outline shapes derived from the 260 line drawings of everyday objects by Snodgrass and Vanderwart (1980). To convert the line drawings into outline shapes we proceeded as follows: First, silhouettes were made (by filling-in the interior surfaces in black), and their outlines were then extracted automatically and spline-fitted to obtain smooth curvature values at all points along the contour (for more details, see Wagemans, Notebaert, & Boucart, 1998, Appendix, and De Winter and Wagemans, 2004). In a previous large-scale study with these silhouette and outline versions of the complete Snodgrass and Vanderwart set the degree of recognizability was examined on a large group of subjects (see De Winter & Wagemans, 2004; Wagemans et al., 2005). On the basis of these normative data, we could now create two conditions with two different types of outline stimuli: ‘easy to recognize’ and ‘difficult to recognize’. We defined easy outlines as being recognized by at least 70% of the previous subjects, whereas difficult outlines were identified correctly by less than 30%. In the analyses below, we will include ‘recognizability’ as a factor and we will examine the effects of other variables separately in each of the two subsets. Using that criterion, we created 44 balanced pairs of stimuli, each consisting of an easy and a difficult stimulus, which were matched by keeping the number of inflections as equal as possible.¹ This automatically guarantees an equal number of protrusions and intrusions for each stimulus of a balanced pair. We also wanted stimuli with enough curvature variation: The number of inflections per balanced pair ranged from 4 to 84.

These 44 balanced pairs were divided into four equivalent sets of stimuli, each consisting of 11 easy and 11 difficult stimuli, with a matched distribution of numbers of

¹ The mean difference in number of singularities between the two stimuli within the balanced pairs (easy-difficult) was -1.84 ($SD=6.93$) for I , 0.16 ($SD=4.62$) for $M+$, and -0.05 ($SD=3.73$) for $m-$. These are small numbers compared to the average number of singularities in each outline.

inflections. Each subject received a booklet of seven pages (A4 format in landscape orientation): one page with instructions, five pages with four outlines each and one page with two outlines, for a total of 22 outlines per subject. The order of the stimulus pages in the booklet was randomized for each subject.

3.3. Procedure

Each subject received written instructions and a set of 22 outline figures based on real objects as described above. Subjects were instructed to look carefully at each stimulus to try to identify it before segmenting it into parts. Under each outline, a line was drawn on which subjects could write the name of the object or a predetermined expression to indicate that they could not recognize it. After that, subjects had to segment the outline shape. Segmentation was defined as the partitioning in salient or important parts. They were instructed to draw segmentation lines, either straight or curved, cutting the outlines in exactly two places. It was also emphasized that each figure had to be segmented, even when it was not recognized.

To permit a quantitative analysis of the segmentation data, a computer program was developed. This program made it possible to transfer the paper-and-pencil data on the computer by drawing (with the mouse) the segmentation lines superimposed over the outline displayed on screen. The program calculated the segmentation points and saved the (x, y) coordinates to data files. To keep the program simple, the segmentation lines on the screen (and in the data reported below) were straight lines. Most of the drawn lines were straight anyway and we were mainly interested in the placement of the segmentation points resulting from the segmentation lines. Ten volunteers performed the task of transferring the data ($201 \times 22 = 4422$ segmented outlines) as accurately as possible. To maximize the consistency of this process, detailed instructions and a demonstration of the program were provided.

4. Results and discussion

A fairly large number of data was excluded from analysis. There were three categories of faults that led to exclusion. (1) The segmentation line did not cut two points on the outline. Sometimes this was true for the majority of the 22 outlines, sometimes for only a number of outlines per subject. (2) Some stimuli were not segmented at all. (3) There was also a rest category of other ‘errors’, like filling-in the inner details of an outline (e.g. drawing an eye on the horse). In addition to exclusion of single, improperly segmented figures, subjects who made mistakes like this on seven or more (out of 22) outlines were also excluded from analysis. In total, 21.4% of the data had to be excluded. This large number is a consequence of the circumstances of data acquisition as mentioned earlier. We considered it better to work with fewer but more reliable data.

To give the reader an initial impression of our data, we start by reporting some general figures about all the stimuli and segmentations. All outlines together had 1338 inflections ($I = 35.4\%$), 1247 positive maxima ($M+ = 32.9\%$), 879 negative minima ($m- = 23.2\%$), 237 positive minima ($m+ = 6.3\%$), and 87 negative maxima ($M- = 2.2\%$) (see Appendix

A for how we determined the singularities). Per stimulus shape, the average number (and SD) of each type of these singularities was 15.2 (13.0), 14.1 (8.5), 10.0 (8.6), 2.7 (2.3), and 1.0 (1.4), respectively. Over all outlines and subjects, 10,773 good part cuts (see above) were made, 6085 (56.5%) in the easy condition and 4688 (43.5%) in the difficult condition. Each outline was segmented 122.4 times on average (SD=80.13) and each subject drew 68.6 part cuts on average (SD=19.47). Of all segmentation points, 74.3% had negative curvature and 25.7% had positive or zero curvature.

In the remainder of this section, we will first present the results concerning individual segmentation points, followed by the results concerning the part cuts (i.e. the pairs of segmentation points connected by the segmentation lines drawn by the subjects) and the resulting parts. The third and final part of this section contains a visual analysis of our rich data set, exploring the principles that determine segmentation of object outlines into parts by means of some self-made visual tools.

4.1. Analysis of segmentation points

The major goal of this type of analysis was to determine whether outlines were segmented at (or near) negative minima of curvature, as predicted by the minima rule by Hoffman and Richards (1984). We tested this idea in many different ways.

4.1.1. Singularities closest to a segmentation point

To get an indication of the relative importance of each singularity type as a possible segmentation point, we calculated the outline-distance for each segmentation point to its nearest singularity point. In Appendix A, we explain how we determined the singularities (the local curvature points of interest) on the outlines and how we determined the distance between a segmentation point and a singularity. For all the segmentation points, we then looked which type of singularity was the nearest neighbor and we tabulated the number of segmentation points closest to each singularity type for all the stimuli (see Table 1). In both conditions (easy and difficult to recognize), negative minima ($m-$) have the highest attraction as possible segmentation points, as predicted by the minima rule of Hoffman and Richards (1984). Of gradually less importance as possible segmentation points are, respectively, inflections (I), positive maxima ($M+$), negative maxima ($M-$), and positive

Table 1
Number of segmentation points with a particular type of curvature singularity as their nearest neighbor

Singularity type	Condition			
	Easy		Difficult	
	Count	Adjusted (%)	Count	Adjusted (%)
$m-$	7081	68.4	4842	63.3
I	2758	19.0	2339	18.7
$M+$	1855	12.6	1947	18.0
$M-$	509		310	
$m+$	132		50	
Total	12,335		9488	

‘Adjusted’ corrects for the unequal number of singularities of each type (see Appendix B).

minima ($m+$). Although these last two types of singularities occur less frequently, the negative ones ($M-$) are more important than the positive ones ($m+$) in this case too. In order to simplify the following analyses, we further disregard these latter two types of curvature singularities.

To get a more balanced indication of the importance of each singularity type as a possible segmentation point, we expressed the frequencies of each type also relative to the total number available of each type (column ‘Adjusted’-see Appendix B for the details of the procedure). Since there are fewer $m-$ points, the adjusted percentage of $m-$ increases, suggesting an even higher importance of these singularities as segmentation points. Inflections are most frequent, so their importance is reduced somewhat, to become almost equal to that of the $M+$ points, especially in the difficult condition.

4.1.2. Singularities and segmentation points within an outline-distance of maximum 10 pixels

In another analysis, we only regarded a singularity as a possible segmentation point if it was close enough to the segmentation point itself. The maximum allowed outline-distance was set to 10 pixels, a threshold deemed high enough to allow for some noise in the segmentation data and small enough to reject unlikely singularity points. This threshold represents on average about 1% of the contour. If several singularities were closer than 10 pixels to a segmentation point, only the closest singularity was considered. This selection accounts for 69.7% of all segmentation points and thus still reflects the data reasonably well. Of the segmentation points further than 10 pixels from a singularity, 72.5% had negative curvature (adjusted).

The data summarized in Table 2 show a clear trend, which is equal in both conditions: Negative minima ($m-$) are clearly the most favorite segmentation points, followed by inflections (I) and then positive maxima ($M+$). This result is consistent with many of the models discussed earlier in this article: Segmentation occurs mainly at $m-$ points. Moreover, when looking at the absolute counts, there seems to be support for some top-down influence: For each type of curvature singularities, more segmentation points are indicated on easy-to-recognize outlines than on difficult ones. This trend could be due to over-segmentation: Segmentation may be applied on the basis of knowledge about the object, possibly without any apparent visual cues.

Table 2
Number of segmentation points with a particular type of curvature singularity at an outline-distance of maximum 10 pixels

Singularity type	Condition		Condition	
	Easy		Difficult	
	Count	Adjusted (%)	Count	Adjusted (%)
$m-$	5554	72.7	4008	69.9
I	2167	20.2	1989	21.2
$M+$	769	7.1	723	8.9
Total	8490		6720	

‘Adjusted’ corrects for the unequal number of singularities of each type (see Appendix B).

In Fig. 4, the distribution of all segmentation points within 10 pixels from a curvature singularity are shown, according to the type of nearest singularity ($m-$, I , and $M+$), in the two conditions (easy vs. difficult), and in the whole range of distances to the nearest singularity, from 0 (i.e. segmentation point falls exactly on the singularity) to 10 pixels. This figure shows clearly that $m-$ are always more frequent than the other singularities, at all the tolerance ranges between 0 and 10 pixels, and that the cognitive influence is strongest and most systematic for $m-$.

4.1.3. Segmentation popularity of singularities

In the above analyses, we collapsed segmentation points into different categories but we did not look at individual segmentation points. However, it is also interesting to examine whether subjects select the same segmentation point more consistently in one condition than another. We therefore measured the frequency of each particular singularity as a segmentation point, a concept we denote ‘popularity’ (we prefer to use a distinct term instead of frequency because later on we will also use the frequency of popularity). The same threshold as before was applied: Each segmentation point within a maximum outline-distance of 10 pixels from a particular singularity was considered as lying on that singularity. The data are summarized in Table 3 and plotted in Fig. 5.

We performed an ANOVA with recognizability and singularity type as independent variables and popularity as the dependent variable. There was a main effect of recognizability [$F(1, 2222)=43.19, P<0.001$] and of singularity type [$F(2, 2222)=490.05, P<0.001$]. There was also an interaction effect [$F(2, 2222)=12.05, P<0.001$]. So, not only more points are used when subjects recognize the object, but the same




Fig. 4. Number of observations (No of obs) of each level of the outline-distance between a segmentation point and the nearest $m-$, I or $M+$ singularity, with the outline-distance rounded to an integer value (range between 0 and 10).

Table 3
Mean popularity of singularity points as segmentation points (SD in parentheses)

Singularity type	Condition		
	Easy	Difficult	
<i>m</i> –	14.3 (9.6)	>>	10.7 (8.9)
<i>I</i>	5.2 (4.9)	>>	3.9 (3.2)
<i>M</i> +	3.0 (3.3)	=	2.7 (2.5)

'>>' stands for 'larger than', with statistical significance at $P < 0.001$.

segmentation points are also used more consistently. The effect of recognizability was significant for *m*– points [$F(1, 763) = 29.93, P < 0.001$] and for *I* points [$F(1, 993) = 25.66, P < 0.001$], but not for *M*– points [$F(1, 524) = 2.27, P = 0.133$], although the trend was in the same direction.

Looking at the frequency plots in Fig. 5, there is a clear difference in the shape of the distributions for *m*– points compared to *I* and *M*– points. The popularity of *I* and *M*– points have a much higher frequency for low popularities than for high popularities. This means that many subjects chose different *I* and *M*– points as segmentation points. This lower popularity reflects a larger variability and thus a lower salience of these points as segmentation points. For *m*– points the distribution is more uniform. Points with a high popularity are relatively frequent and the maximum popularity is clearly higher too. Subjects apparently chose the same *m*– points more frequently. This implies that these points have a higher salience. This became even clearer when we performed a visual analysis of the segmentation data (see below).

4.1.4. Correlation between the segmentation popularity of singularities and scale

It is a widely accepted idea that visual shapes are analyzed at different spatial scales. Using the scale-space notion of Witkin (1986), it is possible to trace the existence of a specific singularity over the different spatial scales. The number of subsequent scales a singularity 'survives' is expressed by its scale: The longer a singularity survives the increasing smoothing, the higher its scale is. It is quite natural to assume that curvature singularities are visually more salient when they have a higher scale. Here, we test whether the segmentation popularity of singularities (how frequent a singularity is chosen as a segmentation point) correlates with its scale. The calculation of the segmentation popularity is described in Appendix C (for an example, see Fig. 6B and E). The procedures for calculating the smoothed outlines, the scale-space representation, and the scale are described in Appendix D (for an example, see Fig. 6D and F).

Table 4 summarizes the Spearman correlations between the scale of a singularity and its segmentation popularity over all outlines. They are strongest and positive for *m*– points, and almost equal for the two conditions (easy and difficult to recognize). Correlations for *M*– points are also almost equal in both conditions, but now they are smaller and negative. For inflections, there is no correlation in the easy condition, and only a marginal one in the difficult condition.

These data suggests that a large-scale *m*– singularity (observable over small and large spatial scales) has a higher attraction as a segmentation point than a small-scale *m*– singularity. A weaker and opposite effect is observable for *M*– singularities: They

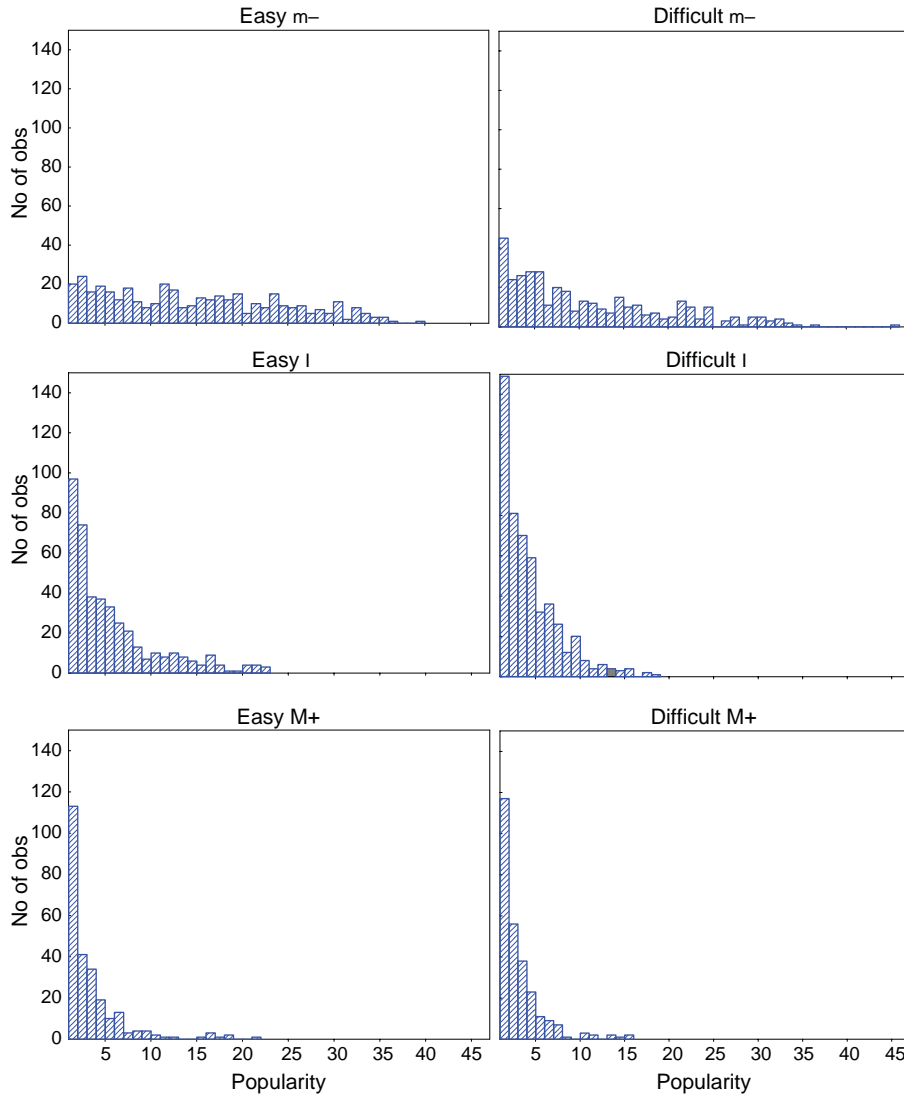


Fig. 5. Histograms (number of observations or No of obs) of the popularity values of specific curvature singularities as segmentation points. From top to bottom: negative minima (m^-), inflections (I) and positive maxima (M^+); in the left column are the data for the easy-to-recognize condition, in the right column for the difficult condition.

become less attractive when they have a higher scale. The fact that the popularity of inflections does not correlate with their scale adds support to the previous finding that these curvature singularities have little attraction in general as a segmentation point. The fact that there is virtually no effect of recognizability suggests that the scale is predominantly a low-level factor. However, the importance of this factor must not be

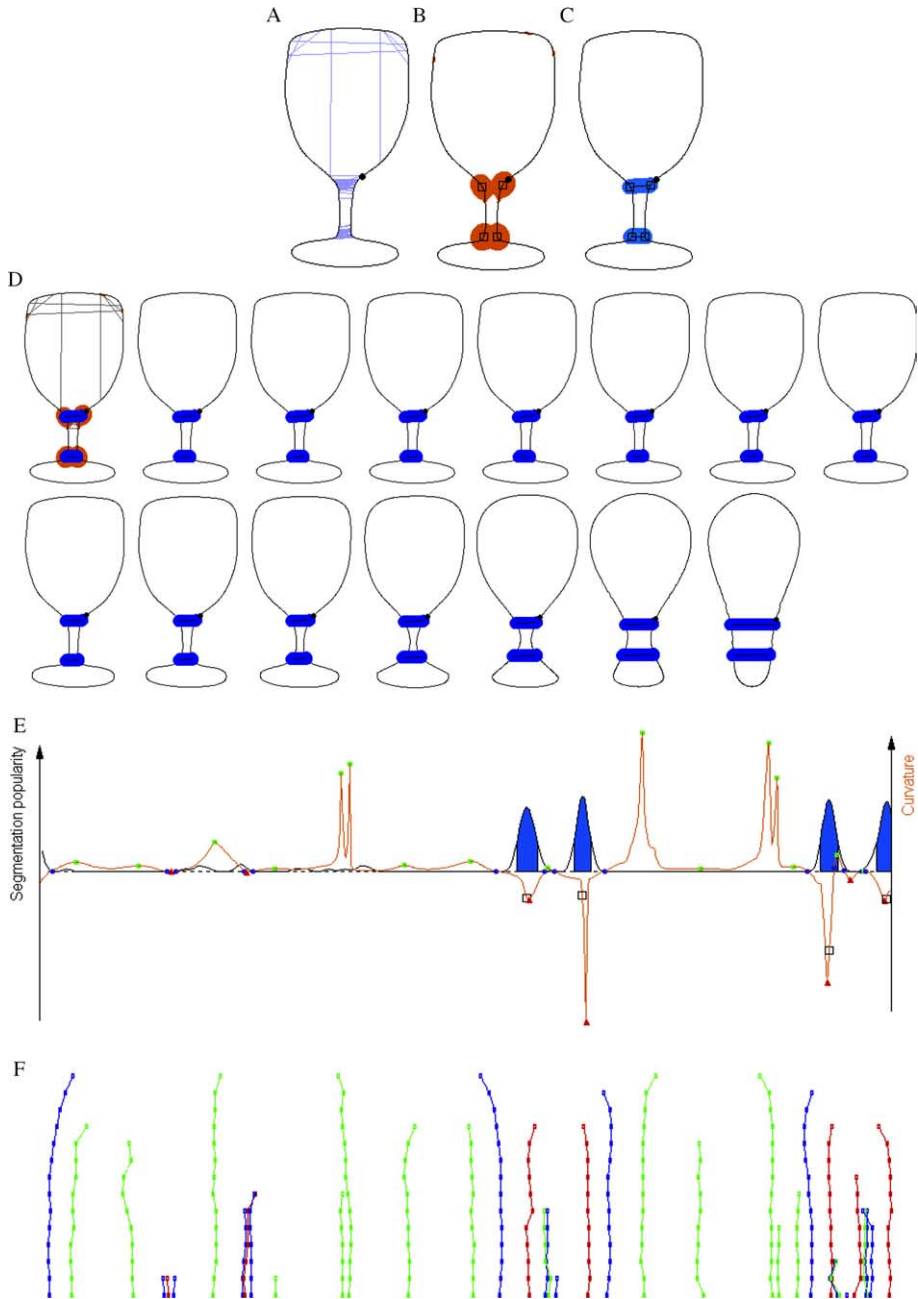


Fig. 6. (A) The outline of a wineglass as segmented by our subjects (light blue lines). (B) The orange blobs represent the popularity (bigger radius is higher popularity). The black squares represent the (anchor) segmentation points (also in C and E; see also Appendix E). (C) The thick blue lines represent the summarized part cuts (with a thin black line at the center). The thickness of the blue lines represents the between-subjects

Table 4
Correlations between the segmentation popularity of all singularities and their scale

Singularity	Condition					
	Easy			Difficult		
	<i>N</i>	<i>r</i>	<i>P</i>	<i>N</i>	<i>r</i>	<i>P</i>
<i>m</i> −	679	0.35	0.00	631	0.36	0.00
<i>M</i> +	943	−0.19	0.00	894	−0.20	0.00
<i>I</i>	1082	0.01	0.62	1068	0.08	0.01

overstated: For *m*− points, it explains 12.6% of the variance and for *M*+ points only 3.8%. Nevertheless, the correlations are statistically significant and opposite in effect, and hence they suggest that the popularity of a singularity as a segmentation point is at least partially determined by its spatial scale.

4.1.5. Correlation between the segmentation popularity of extrema and absolute curvature

An obvious question is whether the attraction of curvature extrema (positive maxima, *M*+, or negative minima, *m*−) as segmentation points is determined by the amount of curvature. To address this question, we use all segmentation points within a distance of 10 pixels away from an extremum. We then determine the smoothed popularity using the procedure explained above (using a smoothing value of 5 pixels, and no saliency threshold to have more points in the sample) and correlate it with the absolute curvature in its nearest-neighbor extremum. Our sample includes 11,054 segmentation points: 9562 closest to *m*− and 1492 closest to *M*++; 6323 from the set of easy-to-recognize objects and 4731 from the difficult ones. The correlations were positive for *m*− (.286, $P < 0.0001$) and negative for *M*+ (−0.134, $P < 0.0001$). The correlations were generally weaker for the easy-to-recognize objects than for the difficult ones: 0.261 ($P < 0.0001$) versus 0.319 ($P < 0.0001$), for the *m*− points; and −0.116 ($P < 0.005$) versus −0.148 ($P < 0.0001$), for the *M*+ points, respectively. This pattern of correlations confirms that deep negative minima are strong cues to segmentation, whereas segmentations around positive maxima are most likely influenced by other factors than local curvature as such. We will return to this observation below.

consistency of the part cut (thicker is higher). In each of these outlines, the black dot represents the starting point (of counter-clockwise tracing along the contour) for the segmentation popularity and curvature graph (E), as well as the scale-space representation (F). (D) The original wineglass outline with segmentation lines and its highly salient part cuts in thick blue lines superimposed. The following versions of the wineglass (left to right and downwards) represent Gaussian smoothed versions with increasing sigma with the same highly salient part cuts. (E) The popularity graph (black line, left axis) represents the popularity of segmentation points along the outline. The blue sections have a popularity value above threshold (=10) and correspond to the selected part cuts. The curvature graph is superimposed (in orange). Note that the anchor points (open squares) are all located near an *m*− point. The X-axis follows the outline counter-clockwise, the Y-axis represents the segmentation popularity of each point or curvature value. Green squares represent *M*+ singularities, red triangles *m*− points and blue dots inflections. (F) The representation of the scale space. The begin position of a scale matches the x-position on the curvature graph (E) above it (singularities not connected to other singularities are not displayed for added clarity).

4.2. Analysis of part cuts and part outlines

In this section, we examine which segmentation points were paired by being connected by a part line or part cut. In addition, we explore the nature of the resulting part on the basis of its curvature description at different spatial scales. In the final part of this section, we report some statistics on some characteristics of the resulting parts and we provide some examples to illustrate them.

4.2.1. Part cuts

As before, a segmentation point was considered to be lying on a singularity when the outline-distance between segmentation point and singularity was 10 pixels maximum. The data are summarized in Table 5. The trend is the same as with the segmentation points. An *m*− point was most often used as a segmentation point and now a part line was also more often going through two *m*− points than through any other pair of singularities. Note that an *m*− point was also frequently connected to a non-singularity (indicated by ‘?’ in Table 5). The choice of the second point in a pair follows the previously established trend for separate segmentation points: *m*− points are preferred, then *I* points, and finally *M*+ points. The trends are quite similar for the two recognizability conditions. The main difference is again that there are fewer segmentation lines in the difficult condition (a trend already observed before), which could indicate a possible influence of stored knowledge on object segmentation. Note also that there is a trend to use more non-singularities in the easy condition, which could be interpreted as evidence for segmentation influenced by object knowledge in the absence of visual cues in terms of curvature singularities. In fact, non-singularities have generally been used rather frequently. To help us understand the other factors that have determined the choice of the segmentation points, we decided to analyze the segmentation data across several spatial scales (see next paragraph) and to analyze the data visually (see next section).

Table 5
Number of segmentation lines connecting two types of segmentation points

Segmentation points		Condition			
		Easy		Difficult	
<i>m</i> −	?	1447	24%	921	20%
<i>m</i> −	<i>m</i> −	1370	22%	1031	22%
<i>m</i> −	<i>I</i>	1018	17%	743	16%
<i>m</i> −	<i>M</i> +	349	6%	282	6%
<i>I</i>	?	436	7%	424	9%
<i>I</i>	<i>I</i>	284	5%	336	7%
<i>I</i>	<i>M</i> +	145	3%	150	4%
<i>M</i> +	?	181	2%	193	3%
<i>M</i> +	<i>M</i> +	47	1%	49	1%
?	?	871	14%	580	12%
	Total	6148		4709	

‘?’ denotes that the segmentation point was not within an outline-distance of 10 pixels from a curvature singularity.

4.2.2. Part outlines and spatial scale

In this part of the analyses, we explore the nature of the resulting parts, namely the fraction of the outline located between the two segmentation points of a part cut, which we call the ‘part outline’. Traveling along the outline, we always take the shortest route of the two possible ones, and we describe the shape changes that the part outlines make by reporting the sequence of curvature singularities on them (ignoring the inflections that necessarily occur between $m-$ and $M+$). In doing so, we follow Richards and Hoffman (1985), who have proposed to describe shape as sequences of so-called ‘codons’, typical strings of curvature singularities that act as primitive building blocks for complete contours of shapes (see also Richards, Koenderink, & Hoffman, 1987). When shapes are segmented at pairs of negative minima, the resulting part outline should be an $m- M+ m-$ tuple in most cases. In line with our previous analysis, we will test this idea at multiple spatial scales.

To keep the reported data within acceptable limits, we only selected the part cuts with a high between-subjects consistency. These represent highly salient part cuts and are therefore the most important ones to report.² The selection procedure is described in Appendix E and visualized in Fig. 6. With this procedure, 3927 highly salient part cuts were selected out of a total of 10,773 part cuts (i.e. 36% of all of the part cuts in the total set; 2630 ‘easy’ and 1237 ‘difficult’). Aggregating across individuals then yields 188 different salient parts selected (124 from the ‘easy’ condition and 64 from the ‘difficult’ condition).

After selecting high salience part cuts, we assigned the closest neighboring singularities (again within an outline-distance of 10 pixels) to the two segmentation points of a highly salient part cut. We then followed these two reported singularity points from spatial scale to spatial scale and also report the part outline bordered by these two points as tuples of curvature singularities. We could always follow the pairs of segmentation points over different spatial scales because the smoothing procedure only displaced their location, without deleting points. If there was no singularity within the predefined range, the actual segmentation point itself was chosen and traced over the consecutive spatial scales; these will be reported by way of a question mark (?).

Again, to restrict the reported data to a manageable amount, we only report tuples at two special spatial scales. (1) Because the spatial scale of a segmentation point (as defined above) reflects a single physical event, we considered it informative to determine the scale of a part cut too (i.e. generally the highest scale of the two segmentation points, see Appendix F). (2) Equally interesting is the largest possible spatial scale before the smoothed outline becomes totally convex, because this automatically yields the simplest possible tuples. We do not report the tuples at the finest (original) scale, because they vary greatly between part cuts, even when the two segmentation points are the same. This would make it nearly impossible to present a summary of the data. So, for each highly salient part cut, we report the nearby singularities at its non-smoothed scale, follow the

² It is for the same reason that we report only the codon description of the part with the shortest part outline. Adding all the complementary tuples would yield more data from ‘parts’ that are probably not intended as such but rather as the remainder of the object when the most intuitive parts are taken out.

resulting part outline over all spatial scales, and report the contained tuple at the two specified scales.

Tables 6 and 7 are organized as follows. We grouped the data by type of part cut and sorted these by popularity. As before, we only considered $M+$, $m-$ and I points as singularities (ignoring the very infrequently used $m+$ and $M-$ singularities). The column *part cut* reports the type of part cut by its two neighboring singularities, located within an outline-distance of 10 pixels of its two segmentation points. A ‘?’ denotes that no singularity was within an outline-distance of 10 pixels. In the next two columns, we report the tuples or sequences of curvature singularities in terms of $M+$ and $m-$ only, ignoring the inflections that necessarily occur between them (this entails no loss of information but yields a more compact representation). We report tuples of maximum length 3; longer tuples are summarized and denoted by ‘...’. To avoid duplication of symmetric part cuts and symmetric tuples, the singularities are sorted so that $m-$ always comes first. For instance, following the outline counter-clockwise, both $M+M+m-$ and $m-M+M+$ would be reported as $m-M+M+$.

As described before, we report the tuples at two spatial scales (columns *scale* and *end*). The *scale* column reports the tuples at the level of the part cut scale. The subsequent *end* column reports the tuple that is formed at the largest scale and is evolved from the tuple in the previous *scale* column. The *n* column represents the number of part cuts that have exactly the same reported configuration (same type of part cut and same tuples). These can be on the same outline, but are mostly on different outlines. For instance, if $n=5$ for part cut $m-\dots m-$, then there are five part cuts that have two $m-$ singularities located closely to their two segmentation points and the reported tuples are exactly the same. It states nothing about the number of subjects that have marked those part cuts (see later). The *cumm* column reports the *n* column expressed as a cumulative percentage. The *average frequency* column represents the average number of subjects that marked a part cut configuration of that type (standard deviation between brackets). The *weighted frequency* column reports the total number of subjects that marked a part cut of a given configuration (and is equal to *n* multiplied by *average frequency*). The data for each part cut are summarized (under the line) by reporting the sum of *n* and *weighted frequency* and by reporting the average of the *average frequency*.

We observed several trends that are similar in both the easy and difficult conditions (in Tables 6 and 7, respectively). First, the most popular part cuts have always two $m-$ singularities and many part cuts have at least one $m-$ singularity. The most salient ones typically yield a tuple consisting of only one $M+$ singularity at the part cut scale (‘scale’) or none at all at the largest smoothing scale (‘end’, where the part outline has predominately positive curvature because the whole outline is then almost totally convex). Thus, subjects effectively select part cuts resulting in part outlines with $m-M+m-$ configurations at that scale and/or at the largest possible scale in about a quarter to a third of the cases. The second most popular part cut has one $m-$ singularity and one I singularity. The less popular part tuples also have a least one $m-$ segmentation point in many cases and they proportionally more often occur in longer sequences (especially at the scale of the part cut). Note that these are much more isolated cases, each accounting for only a very small portion of part cuts (they occur in only 1 outline each, on average). We give some examples of such isolated cases in Fig. 7.

Table 6
Highly salient part cuts and their tuple configuration at the scale of the part cut and the largest scale for the easy condition (see text for details)

Part cut	Tuplet length	Scale	End	<i>n</i>	Cumm	Average frequency	Weighted frequency
<i>m</i> –... <i>m</i> –	0	?	?	1		19.0	(0.0) 19.0
	1	<i>M</i> +	<i>M</i> +	15		24.8	(7.5) 372.0
			?	10		19.1	(6.3) 191.0
	2	<i>M</i> + <i>M</i> +	<i>m</i> –	2		24.5	(10.6) 49.0
			<i>M</i> +	4		27.3	(4.3) 109.0
			<i>M</i> + <i>M</i> +	1		14.0	(0.0) 14.0
			<i>m</i> – <i>M</i> +	2		22.7	(3.5) 45.3
			<i>m</i> – <i>M</i> +	2		14.5	(2.1) 29.0
			<i>M</i> +	1		21.0	(0.0) 21.0
	3	<i>m</i> – <i>M</i> + <i>M</i> +	<i>M</i> +	2		31.5	(4.9) 63.0
			?	1		31.0	(0.0) 31.0
			<i>M</i> + <i>m</i> – <i>M</i> +	1		29.0	(0.0) 29.0
			<i>M</i> + <i>M</i> + <i>M</i> +	1		29.0	(0.0) 29.0
			<i>m</i> – <i>M</i> + <i>m</i> –	1		20.0	(0.0) 20.0
	>3	...	<i>M</i> +	1		10.0	(0.0) 10.0
			<i>M</i> +	5		22.6	(8.0) 113.0
			<i>M</i> + <i>M</i> +	1		28.0	(0.0) 28.0
			50	40.3%	22.4	1143.3	
<i>m</i> –... <i>I</i>	1	<i>M</i> +	?	8		21.3	(7.1) 170.0
			<i>M</i> +	7		20.7	(7.9) 145.0
	2	<i>m</i> – <i>M</i> +	?	2		24.0	(2.8) 48.0
			<i>M</i> +	2		23.0	(8.5) 46.0
			<i>m</i> – <i>M</i> +	1		23.0	(0.0) 23.0
			<i>m</i> –	1		20.0	(0.0) 20.0
			<i>M</i> + <i>M</i> +	1		16.0	(0.0) 16.0
	3	<i>m</i> – <i>M</i> + <i>M</i> +	<i>M</i> +	1		13.0	(0.0) 13.0
			?	1		25.0	(0.0) 25.0
			<i>M</i> +	3		22.3	(13.1) 67.0
	>3	...	<i>M</i> + <i>M</i> +	1		10.0	(0.0) 10.0
			28	62.9%	19.8	583.0	
<i>m</i> –...?	1	<i>M</i> +	?	5		23.2	(6.3) 116.0
			<i>M</i> +	5		19.4	(4.0) 97.0
	2	<i>m</i> – <i>M</i> +	<i>m</i> – <i>M</i> + <i>m</i> –	1		17.0	(0.0) 17.0
			<i>M</i> +	4		20.8	(7.5) 83.0
	>3	...	?	1		15.0	(0.0) 15.0
			<i>m</i> – <i>M</i> + <i>M</i> +	1		15.0	(0.0) 15.0
			17	76.6%	18.4	343.0	
<i>m</i> –... <i>M</i> +	1	<i>M</i> +	<i>M</i> +	4		19.3	(3.6) 77.0
			?	1		15.0	(0.0) 15.0
	2	<i>M</i> + <i>M</i> +	?	1		25.0	(0.0) 25.0
			<i>M</i> +	1		15.0	(0.0) 15.0
	3	<i>m</i> – <i>M</i> + <i>M</i> +	<i>M</i> +	1		33.0	(0.0) 33.0
			?	1		24.0	(0.0) 24.0
>3	...		9	83.9%	21.9	189.0	

(continued on next page)

Table 6 (continued)

Part cut	Tuplet length	Scale	End	<i>n</i>	Cumm	Average frequency	Weighted frequency	
<i>I...?</i>	1	<i>M+</i>	?	2		16.5 (0.7)	33.0	
	3	<i>M+m-M+</i>	<i>M+</i>	1		13.0 (0.0)	13.0	
	>3		<i>M+</i>	?	2		16.5 (0.7)	33.0
				<i>m-</i>	1		27.0 (0.0)	27.0
					1		12.0 (0.0)	12.0
			7	89.5%	17.0	118.0		
<i>I...I</i>	0	?	?	1		12.0 (0.0)	12.0	
	1	<i>M+</i>	?	2		19.5 (7.8)	39.0	
	>3		<i>M+</i>	1		16.0 (0.0)	16.0	
				?	1		26.0 (0.0)	26.0
			5	93.5%	18.4	93.0		
<i>?...?</i>	1	<i>M+</i>	<i>M+</i>	5		16.0 (2.3)	80.0	
				5	97.6%	16.0	80.0	
<i>M+...I</i>	1	<i>M+</i>	?	2		21.5 (3.5)	43.0	
	>3		<i>M+</i>	1		15.0 (0.0)	15.0	
					3	100.0%	18.3	58.0

In Fig. 7A to D, part cuts (marked in red) have one *m-* and a second variable singularity (mostly *M+*). Most of these cases occur for something bent around a joint. For instance, where the tail of the rhinoceros connects with the body, it forms only one salient *m-* point, but in another pose (not hanging down) it could have two salient *m-* points. Note that most people prefer the shortest part cut instead of segmenting the tail vertically so that two *m-* points would be connected but the part line would be longer. Sometimes the bending is more permanent, as in the earpieces of the glasses, but that does not prevent them from still being perceived as separate parts.³ Note also that the *M+* points, selected to form a pair with the salient *m-* points, often have low salience (see negative correlation with scale reported above). It is interesting to point out that this type of part was also found by Siddiqi et al. 1996, (pp. 416–418, Figs. 25 and 26), who explicitly described ‘trunks’ and ‘tails’ as ‘limbs for which perceptual evidence has been partially removed by the process of bending’. They considered this as one source of oversegmentation in the set of animals compared to the nonsense shapes.

In Fig. 7E and F, we show examples of the rather infrequent *I...I* part cuts. It might be that these part cuts are also predominantly determined by cognitive influences since there appears very little visual information to mark the parts. On the other hand, these inflections are rather close to some slight indentations (*m-*). Indeed, we observed that many inflections selected as segmentation points were often located close to an *m-* singularity and we confirmed this statistically (see Table 8 bottom row). If an inflection is the closest

³ Although both earpieces have been segmented, only one is selected as highly salient. The segmentation lines on the other earpiece are spread out in too many directions, resulting in decreased segmentation popularity (below the cut-off value) of the points in that vicinity.

Table 7

Highly salient part cuts and their tuple configuration at the scale of the part cut and the largest scale for the difficult condition (see text for details)

Part cut	Tuplet length	Scale	End	<i>n</i>	Cumm	Average frequency	Weighted frequency	
<i>m</i> –... <i>m</i> –	1	<i>M</i> +	<i>M</i> +	16		21.2 (7.3)	339.0	
			?	4		20.0 (5.4)	80.0	
	2	<i>M</i> + <i>M</i> +	<i>M</i> + <i>M</i> +	3		19.0 (6.2)	57.0	
			<i>m</i> – <i>M</i> +	1		32.0 (0.0)	32.0	
			<i>m</i> – <i>M</i> +	?	1		16.0 (0.0)	16.0
			<i>m</i> – <i>M</i> +	1		14.0 (0.0)	14.0	
			<i>M</i> +	1		13.0 (0.0)	13.0	
	3	<i>m</i> – <i>M</i> + <i>M</i> +	<i>m</i> – <i>M</i> + <i>M</i> +	1		32.0 (0.0)	32.0	
			<i>M</i> +	1		16.0 (0.0)	16.0	
			<i>M</i> + <i>m</i> – <i>M</i> +	1		32.0 (0.0)	32.0	
	>3	1		22.0 (0.0)	22.0	
			<i>M</i> + <i>M</i> +	1		15.0 (0.0)	15.0	
					32	50.0%	21.0	668.0
	<i>m</i> –... <i>I</i>	1	<i>M</i> +	<i>M</i> +	5		20.2 (6.8)	101.0
?				5		18.8 (4.9)	94.0	
2		<i>M</i> + <i>M</i> +	<i>M</i> +	1		18.0 (0.0)	18.0	
			<i>m</i> – <i>M</i> +	1		12.0 (0.0)	12.0	
>3		...	<i>M</i> + <i>m</i> –	1		11.0 (0.0)	11.0	
			<i>M</i> + <i>M</i> + <i>M</i> +	1		17.0 (0.0)	17.0	
					14	71.9%	16.2	253.0
<i>m</i> –...?	1	<i>M</i> +	<i>M</i> +	4		19.0 (4.1)	76.0	
	2	<i>M</i> + <i>M</i> +	<i>M</i> +	1		12.0 (0.0)	12.0	
	3	<i>M</i> + <i>M</i> + <i>M</i> +	<i>M</i> +	1		29.0 (0.0)	29.0	
	>3	...	<i>m</i> – <i>M</i> + <i>m</i> –	1		29.0 (0.0)	29.0	
			<i>M</i> + <i>M</i> +	1		17.0 (0.0)	17.0	
					8	84.4%	21.2	163.0
<i>m</i> –... <i>M</i> +	1	<i>M</i> +	<i>M</i> +	4		19.3 (8.1)	77.0	
	>3	...	?	1		19.0 (0.0)	19.0	
			<i>M</i> + <i>M</i> +	1		24.0 (0.0)	24.0	
					6	93.8%	20.8	120.0
<i>M</i> +... <i>I</i>	>3	...	<i>M</i> +	2		24.0 (4.2)	48.0	
					2	96.9%	24.0	48.0
<i>I</i> ...?	>3	...	<i>M</i> +	1		26.0 (0.0)	26.0	
					1	98.4%	26.0	26.0
? ...?	>3	...	<i>m</i> – <i>M</i> + <i>m</i> –	1		19.0 (0.0)	26.0	
					1	100.0%	19.0	26.0

singularity to a segmentation point, then there is also an *m*– singularity within an outline-distance of 10 pixels in 62.5 and in 70.6% of the cases, for easy and difficult to recognize objects, respectively. There are also regularly *M*+ singularities within 10 pixels, but less often (37.5 and 47.1%, respectively). In other words, it could very well be that subjects

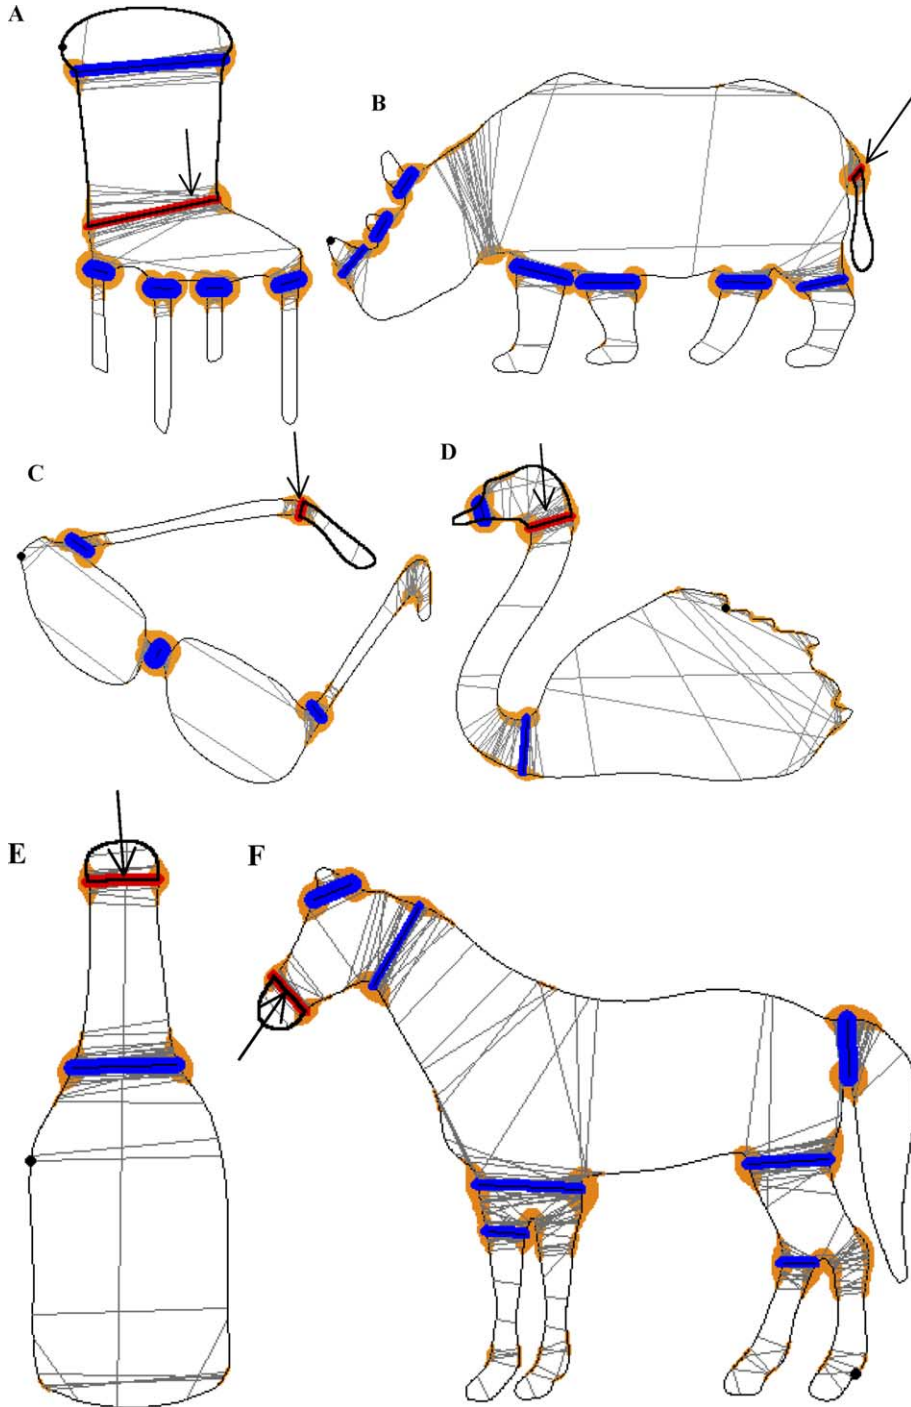


Table 8

For each singularity type closest to a segmentation point (row), the percentage of cases in which a singularity of another type is within an outline-distance of 10 pixels is given (column)

Type	Condition					
	Easy			Difficult		
	<i>m</i> −	<i>M</i> +	<i>I</i>	<i>m</i> −	<i>M</i> +	<i>I</i>
<i>m</i> −	–	7.1%	37.2%	–	3.3%	37.0%
<i>M</i> +	33.3%	–	75.0%	50.0%	–	87.5%
<i>I</i>	62.5%	37.5%	–	70.6%	47.1%	–

For instance, if an inflection is the closest point to a segmentation point (bottom row) then there is also at least one *m*− singularity within an outline-distance of 10 pixels in 62.5% (easy) and in 70.6% (difficult) of the cases.

were in fact indicating *m*− singularities (or to a lesser degree *M*− singularities) instead of inflections and then the overwhelming majority of highly salient part cuts would consist of two *m*− singularities, as predicted by the minima rule by Hoffman and Richards (1984). The choice of inflections may also result from the fact that for smooth part boundaries, the inflections provide a more stable estimate of the tangent direction of the contour on each side of a negative minimum. To the extent that good continuation of the contours of parts (parameterized as their co-circularity) determines segmentation, this advantage of *I* over *m*− as more stable anchors may have affected the results (see also Singh & Hoffman, 2001, Fig. 30b and c).

Despite the overall similarities, we also observed several systematic differences between the easy and difficult condition. First, not only are there more part cuts as reported before, but there are also more than double the number of highly salient part cuts in the easy condition compared to the difficult condition. Second, there are more part-tuplet configurations where the tuplet string is (very) long at the level of the part-cut scale in the easy conditions. Third, there is less variation in chosen part cuts (only 4 types) in the difficult condition and all have at least one *m*− segmentation point. All these differences could be due to cognitive influences. As reported earlier, part cuts are marked more consistently when the shape is easy to recognize. While segmentations in the difficult condition were led mostly (or exclusively) by lower-level influences, knowledge about the object (e.g. its function or the functions of its parts) could have augmented or overridden low-level factors such as the scale of the part cut in the easy condition. As a consequence, the resulting parts in the recognized objects often contained more complex curvature variations than the very primitive type (*m*− *M*+*m*−), hypothesized by Richards and Hoffman (1985) as the prototypical codon for a standard object part.

A good example to illustrate this interplay between low-level visual factors and more high-level cognitive influences is shown in Fig. 8. In Fig. 8A, we have indicated all the marked part cuts, along with the most popular individual segmentation points (marked in



Fig. 7. (A)–(D) Examples of part cuts (marked in red and indicated by an arrow) with one *m*− and a second variable singularity (mostly *M*−). (E)–(F) Examples of part cuts (marked in the same way) with two inflections as segmentation points (see text for more details).

orange) and the two most popular part cuts (marked in blue), showing that the Yale key is clearly segmented into three main parts. First, there is the part that slides into the lock and effectively unlocks the lock. This part has many singularities at the original spatial scale and many subjects have indeed segmented this large part even further into smaller parts. However, at its own part-cut scale (i.e. strongly smoothed, as shown in Fig. 8B), this part cut of the $m-\dots m-$ type has yielded a part outline forming an $m-M+M+m-$ configuration (the $m-$ points in this tuplet are the original segmentation points displaced). The second part is the handle of the key (see Fig. 8B). This part allows the turning of the key and thus has a different function than the first part. Note that the distances between the four segmentation points constituting these two major part cuts are roughly equal, and that these segmentation points could easily be connected vertically too (i.e. in the direction of the elongation axis), as some subjects indeed did. Because these segmentations parallel to the elongation axis are also orthogonal to a local symmetry axis, the short-cut rule would have predicted more equal proportions of the two possible segmentations but the majority of subjects clearly chose to connect the $m-$ pairs orthogonally to the main axis. We suggest that this is due to the knowledge that people have about keys and their function, and therefore create parts that are more consistent with its function. The third main part (see

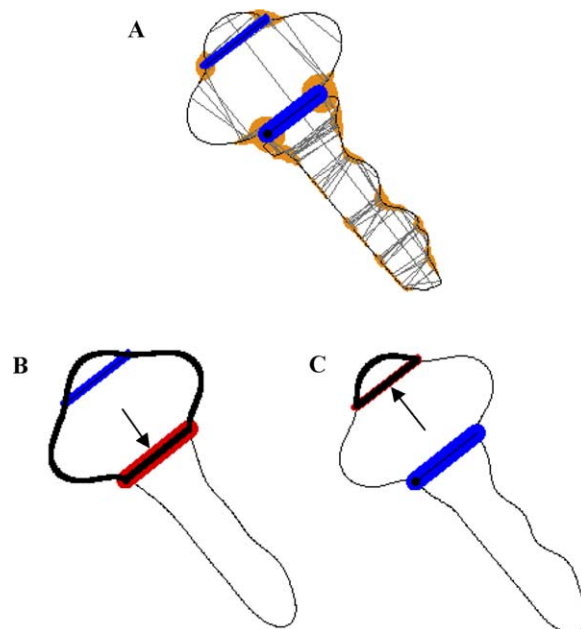


Fig. 8. (A) Original outline of the Yale key with segmentations superimposed and the two most popular part cuts in blue. (B) and (C) The outline smoothed at the scale of the two part cuts (marked in red and indicated by an arrow), with the shortest outline path marked in thicker black. (B) By choosing the shortest path the complementary path is ignored in the reported data (while also being a meaningful part in this example). (C) Another example of a preferred part cut for this outline: at the part-cut scale the formed tuplet has configuration $m-M+m-$ (while there are several $M+$ points at the original scale).

Fig. 8C) is the upper part of the key where typically a hole is located to attach the key to a key ring. Although many parts can be established by using only visual cues, we suggest that segmentation is more consistent when the resulting parts have a distinct function.

4.2.3. Shape characteristics of the resulting parts

In addition to analyzing the part cuts themselves and the contour properties of the resulting parts (e.g. their 1D codon description), it is also interesting to examine some of the more spatial (2D) characteristics of the resulting parts (see Fig. 9 for illustrations). Hoffman and Singh (1997) have proposed two factors as possible determinants of part salience. One is the *relative size* of the part (i.e. the area of the part divided by the area of the whole outline). The second is what they called '*stick out*': the degree to which a part protrudes from the object (i.e. the length of the contour segment of the part divided by the length of the part cut). The larger a part relative to the whole object, and the more it sticks out, the more salient it is supposed to be. A third measure that may be interesting to report is the *aspect ratio* of the part. To compute this, we took the part cut as the base and then drew a rectangle orthogonal to it, extending it until it reached the part's most extreme point (see Fig. 9). The aspect ratio is then the height of this rectangle divided by its original base⁴ (i.e. another measure of how much the part sticks out from the rest of the shape). Here too, one may assume that more elongated parts (i.e. parts that are much 'higher' than they are 'wide') are more salient. A fourth factor that we considered relevant too is the so-called '*compactness*' of the part. Zusne (1970) has defined compactness of overall shapes as the area of the shape divided by the squared contour length. Intuitively, shapes are most compact if their area is as large as possible for the shortest possible contour length. A round shape is very compact, whereas a shape with a lot of parts sticking out has a low compactness. We believe that the salience of a part may be related to the *inverse of compactness* (i.e. less compact parts are more salient in the sense that they stick out more). We therefore define the inverse of compactness of our parts as the squared length of the contour of the part (i.e. contour segment between the two segmentation points defining a part cut, plus the length of the part cut itself), divided by the area of the part itself.

To select a representative sample of salient parts for these computations, we decided to work with the aggregated popularity (as explained in Appendix C and E), using a smoothing value of 5 pixels and a threshold of 5 (i.e. at least 5 subjects should have indicated the same part cut or one very close to it). In the scale-space analysis reported above, the threshold value was set to 10 but then only 188 different salient parts were selected. Using these new settings, we now have a set of 327 different parts (188 'easy' and 139 'difficult'), obtained from aggregating 5302 individual part cuts (i.e. 49% or about half of the total; 3302 'easy' and 2000 'difficult').

⁴ Alternative parameterizations are possible, of course, but each has its own advantages and disadvantages. For example, taking the bounding rectangle and computing the aspect ratio with the length of the side that extends the base rather than the base itself would underestimate the stick out of very skewed or curved parts. On the other hand, trying to determine the length of the elongation axis from the midpoint of the base to the top of the part is not trivial, especially not for skewed parts with rounded tops or for strongly curved parts with strongly tapered tops. Our parameterization seemed to establish a good compromise between being rather simple and being still valid for the majority of cases.




Fig. 9. Illustration of a salient part (as determined by our selection criteria) and some of its shape characteristics. On top is the outline of the whole object (hammer), with its most salient part cuts indicated by thick grey lines and one typical part encircled. Below are then the critical parts and contour fragments of this example part to show how we computed its relative size, stick out, aspect ratio and inverse of compactness of one selected part. The distributions of these parameter values in our set of selected salient parts are shown in Fig. 10 and some examples corresponding to the median and third quartile values are shown in Fig. 11.

Fig. 10 contains the frequency distributions (histograms) of all possible values of these parameters (in several bins) within our subset of 327 most salient parts, separately for the parts in easy-to-recognize and difficult-to-recognize objects. All of these distributions are strongly left-skewed (i.e. a large proportion of cases with low parameter values), with

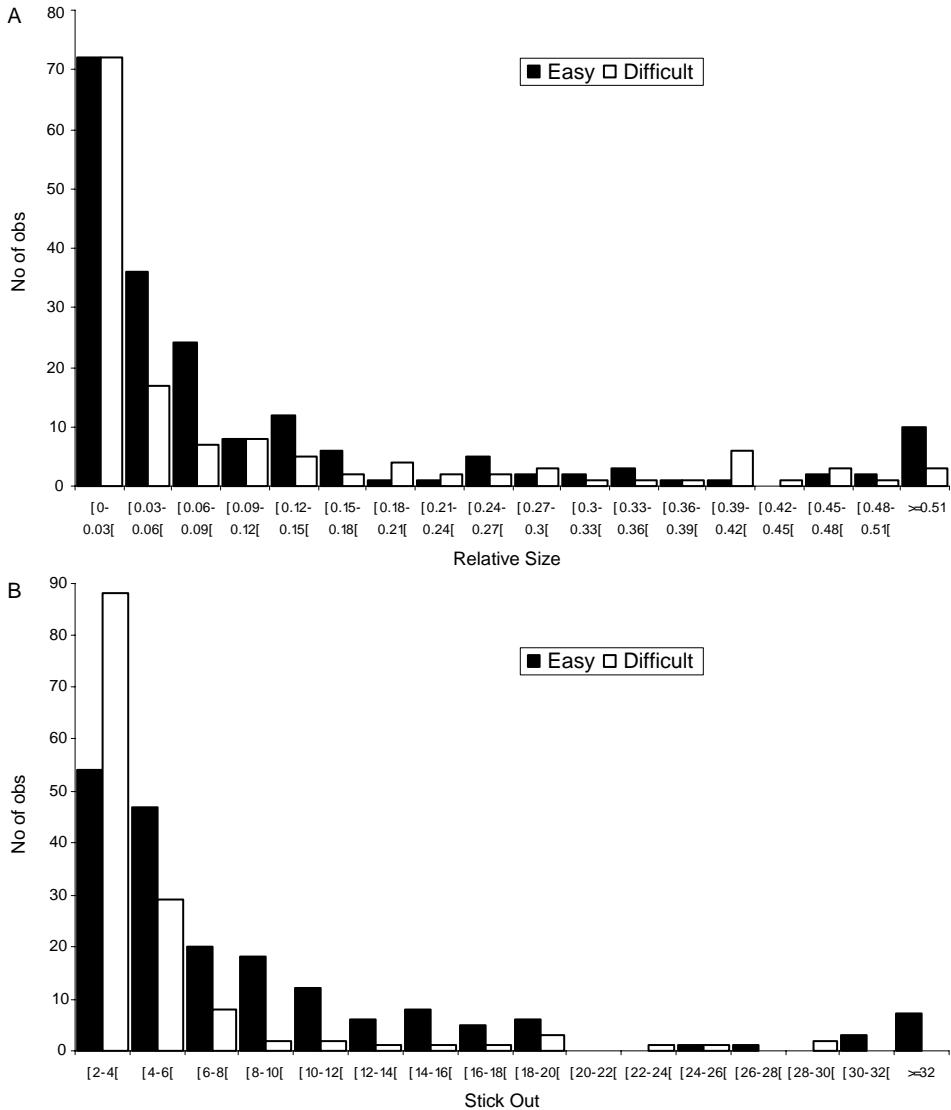


Fig. 10. Frequency distributions (number of observations or No of obs) of the parameter values (in different bins) in our set of selected salient parts (using the criteria explained in the text), separately for parts in easy-to-recognize objects (black bars) and difficult-to-recognize objects (white bars). (A) Relative size, (B) stick out, (C) aspect ratio, and (D) inverse of compactness.

a long tail to the right (i.e. relatively few cases of each bin with high parameter values). With the exception of the lowest parameter values, the cases from easy-to-recognize objects usually outnumber the cases from the difficult-to-recognize objects. Table 9 includes the statistics of these distributions (e.g. mean, SD, first quartile, second quartile or median, third quartile, etc.), both for the whole set and for the two subsets of different

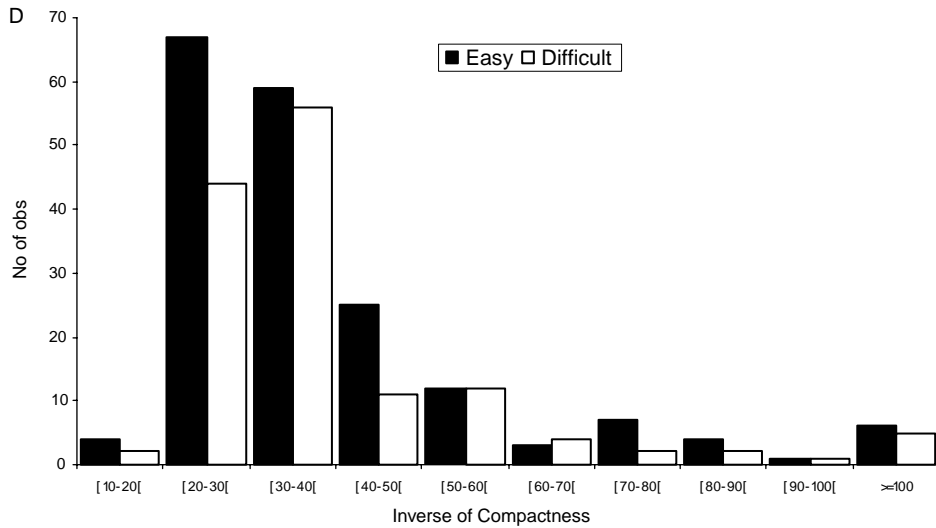
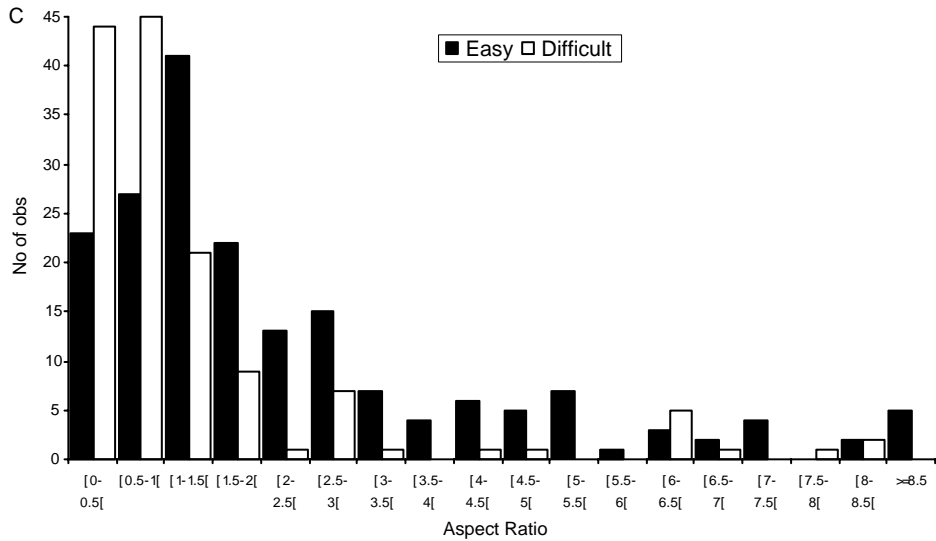


Fig. 10 (continued)

recognizability, as well as the primitive measurements from which they are computed (e.g. contour segment length) and some other relevant parameters (e.g. number of $m-$ and $M+$ included in the contour segment). Fig. 11 illustrates some typical parts, representative for the median value ($Q2$) and third quartile ($Q3$) of each of the four parameters (A and B for relative size, C and D for stick out, E and F for aspect ratio, and G and H for inverse of compactness, respectively).

Table 9

Average (Av), standard deviation (SD), minimum (Min), maximum (Max), first quartile ($Q1$), second quartile ($Q2$ or median), third quartile ($Q3$) values of some relevant shape characteristics of the subset of most salient parts (relative size, stick out, aspect ratio, and inverse of compactness), along with some of the parameter values needed to compute them (part area, part length, part base, part height) and some other relevant characteristics (part frequency and the number or count of $m-$ and $M+$ in the contour fragments of those parts)

Condition		Part frequency	Part area	Part length	Part base	Part height	Count $m-$	Count $M+$	Relative size	Stick out	Aspect ratio	Inverse of compactness
Easy	Av	17.6	3131.0	248.7	41.2	63.2	2.0	3.9	0.11	9.10	2.59	40.26
	SD	8.9	3850.7	175.0	40.9	46.6	2.5	2.9	0.15	10.50	3.18	24.99
	Min	5.0	91.0	38.1	6.0	6.4	0.0	1.0	0.00	2.2	0.12	18.0
	Max	37.0	18,606.0	884.6	270.8	235.0	14	16	0.68	98.5	22.17	254.0
	$Q1$	10.0	475.8	95.2	17.4	25.0	0.0	2.0	0.01	3.7	0.84	27.5
	$Q2$	16.0	1476.5	212.3	27.0	51.0	1.0	3.0	0.05	5.6	1.57	32.7
	$Q3$	25.0	4452.8	361.3	50.8	92.6	3.0	5.0	0.12	10.2	3.01	43.1
Difficult	Av	14.4	3971.6	243.8	64.0	49.4	2.1	4.2	0.10	5.03	1.31	40.98
	SD	8.2	6232.2	215.0	69.7	45.0	3.3	4.1	0.16	4.87	1.67	29.97
	Min	5.0	57.0	35.0	7.3	3.0	0.0	1.0	0.00	2.1	0.04	17.3
	Max	33.0	37,953.0	1005.2	398.1	208.1	30	33.0	0.75	28.8	8.42	319.3
	$Q1$	7.0	421.5	88.1	22.6	16.1	0.0	2.0	0.01	2.6	0.43	28.2
	$Q2$	12.0	1145.0	158.2	36.1	35.0	1.0	3.0	0.03	3.4	0.78	33.2
	$Q3$	21.0	4559.0	354.4	71.9	60.7	3.0	5.0	0.13	4.8	1.25	40.5
All	Av	16.2	3488.3	246.6	50.9	57.3	2.1	4.1	0.10	7.37	2.04	40.57
	SD	8.7	5012.0	192.7	56.1	46.4	2.9	3.5	0.15	8.79	2.72	27.18
	Min	5.0	57.0	35.0	6.0	3.0	0.0	1.0	0.00	2.1	0.04	17.3
	Max	37.0	37,953	1005.2	398.1	235.0	30	33.0	0.75	98.5	22.17	319.3
	$Q1$	8.0	458.0	93.3	19.2	22.1	0.0	2.0	0.01	3.0	0.60	27.6
	$Q2$	15.0	1375.0	186.6	30.1	43.9	1.0	3.0	0.04	4.4	1.16	32.8
	$Q3$	23.0	4462.5	356.2	58.5	78.7	3.0	5.0	0.12	7.9	2.41	41.8

The statistics are tabulated separately for the parts in easy- and difficult-to-recognize objects, as well as for the whole set.




Fig. 11. Examples of parts with some shape characteristics corresponding to the median (Q_2) and third quartile (Q_3) values from the whole set of selected salient parts. Each panel shows the whole outline on top (with the salient part cuts indicated by thick grey lines) and an enlarged version of the part below. The parameter value and the frequency of selection are mentioned too. (A) Glasses (easy) with relative size representative for Q_2 . (B) Motorcycle (easy) with relative size representative for Q_3 . (C) Windmill (easy) with stick out representative for Q_2 . (D) Door knob (difficult) with stick out representative for Q_3 . (E) Toll (difficult) with aspect ratio representative for Q_2 . (F) Axe (easy) with aspect ratio representative for Q_3 . (G) Apple (easy) with inverse of compactness representative for Q_2 . (H) Traffic light (difficult) with inverse of compactness representative for Q_3 .

Some of these factors do indeed correlate well with part salience (or, at least with our measure of aggregated popularity of the part cut). In decreasing order, the correlations were 0.423 for aspect ratio ($P < 0.0001$), 0.401 for stick out ($P < 0.0001$), -0.122 for the inverse of compactness ($P < 0.05$), and 0.088 for relative size ($P = 0.113$). The first two parameters correlate positively with part salience in both subsets, albeit slightly stronger in the ‘difficult’ condition than in the ‘easy’ condition: 0.546 versus 0.363, respectively, for aspect ratio, and 0.461 versus 0.370, respectively, for stick out (all $P < 0.0001$). The last two parameters do not correlate with part salience in the subset of parts from the ‘easy’ condition: 0.011 and 0.022 for the inverse of compactness and relative size, respectively. Within the subset of parts from the ‘difficult’ condition, the correlation with part salience is slightly positive for relative size (0.180, $P < 0.05$), as expected, but, contrary to what was expected, it was negative for the inverse of compactness (-0.287 , $P < 0.001$). This seems to suggest that parts from the difficult-to-recognize objects are more salient when they are more compact (i.e. larger area for the same contour length), rather than the other way around (i.e. longer perimeter for the same area). However, the scatter plot of the data points entering this correlation shows that this value may be strongly biased by a few outliers. Moreover, the visual examples shown below will make it clear that parts in the difficult-to-recognize objects are sometimes determined by the demand characteristics of the task (e.g. shapes cut in four pieces by two diagonals, rather than protruding parts being chopped off).

4.3. Visual analysis

To help us understand better all the factors that have determined the choice of the segmentation points, we designed and implemented a computer program that could represent the segmentation data in a comprehensive visual manner. This enabled us to look at each stimulus in turn with the segmentations aggregated across subjects summarized and superimposed on the stimulus itself. An example is given in Fig. 12 but all the stimuli are made available in this manner on our website to encourage other researchers to use them for their own purposes too (<http://www.psy.kuleuven.ac.be/~winterj/Segmentation/index.html>). Light grey lines represent segmentation lines. The circles represent the popularity of a segmentation point: The bigger the circle, the higher the popularity of that segmentation point. On the right, the raw (original) data are collected. On the left, all segmentation points within a distance of maximally 10 pixels from a singularity are clustered (each point is allocated only to its nearest singularity), thereby shifting segmentation lines towards their closest singularity. Red circles represent the




Fig. 12. An example with a clearly visible neck (separating the tail from the main body of the fish) and limbs (the fins of the fish).

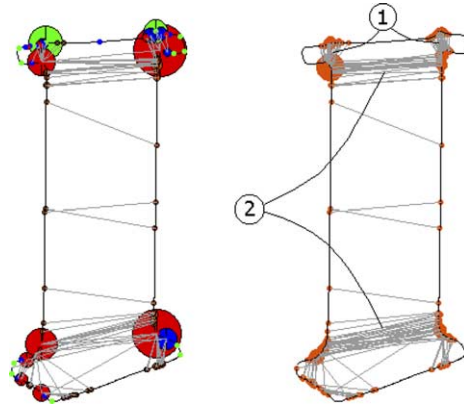


Fig. 13. An example where the short-cut rule is applied to form two parts at the top (1), while other instances (2) provide counter-examples to the short-cut rule.

aggregated segmentations near $m -$ points, while green circles do the same for $M +$ points and blue ones for I points. Using these representations, we will give examples of the typical factors influencing segmentation along with examples of specific types of parts as proposed in previous models (reviewed above).

Fig. 12 gives a good example of the parts described by Siddiqi et al., 1996; (see also their Fig. 11 on p. 406). At the back of the fish, we can see a clear 'neck'. Note that there is a fairly widespread range of segmentation lines, possibly due to the low salience of the individual points along the contour of the neck (Hoffman & Singh, 1997). In contrast, the part lines of the 'limbs' are drawn more consistently (e.g. the fins, especially those along the bottom side of the fish). This could be due to the very good continuation of the contour fragments at both sides of each of the fin parts, as well as to a higher salience of the two $m -$ points (i.e. deeper concavities).

In this example, (from the easy condition) the head of the fish is segmented off by a number of subjects although visual cues in terms of curvature singularities are almost completely absent. This does not happen very frequently in our database but it seems to be due to a cognitive influence of stored knowledge about the object (see also Fig. 14, where the mouth of the horse is segmented from the rest of the head). This could be a first source for the use of other points than $m -$ points as segmentation points.

In Fig. 13, we have two clear examples (indicated by number 1) of the *short-cut rule* (Singh, Seyranian, et al., 1999). An $m -$ point forms the first point of the segmentation line, while the other point is the closest point on the other side of the surface and not necessarily an $m -$ point (or any other singularity). Because this object from the difficult condition is not recognized (as a swing) by most subjects, and cognitive factors are not at work here, this is thus a second (independent) source for the use of non-singularities as segmentation points. A fairly large number of subjects connect two $m -$ points while crossing a relatively big part of the object (indicated by number 2 in Fig. 13). We observed that single salient $m -$ points easily lead to decomposition when accompanied by object or part symmetry (or main orientation axes, as observed in other examples not shown here).

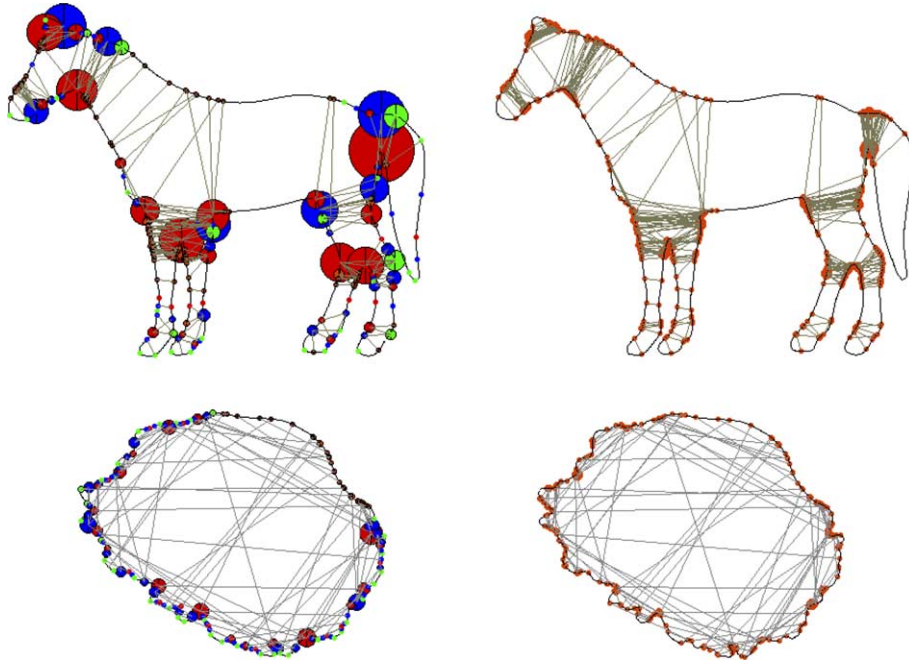


Fig. 14. The outline in the upper half is from the easy condition. The outline in the lower half is its matched counter-part from the difficult condition. Note that the segmentation lines are much more variable in the latter condition.

Symmetry is a strong visual regularity which is easily detected (e.g. Olivers & Van der Helm, 1998; Van der Helm & Leeuwenberg, 1996; Wagemans, 1995, 1997, 1999), so it is not surprising that object symmetry as well as more local symmetry can aid in the formation of parts. The need to cross a symmetry line was also pointed out by Singh, Seyranian, et al. (1999) in their short-cut rule, but it appears strong enough to work more generally than in the case where the shortest distance is chosen. In fact, in this particular case, the overall symmetry/elongation turned out to be more important than the shortest cut because clearly fewer subjects have drawn the shorter segmentation lines here. This is another example of natural segmentations being less strict than the geometric rules proposed in the literature (similar to the counterexamples against ‘necks’ and ‘limbs’ given by Singh & Hoffman, 2001; Singh, Seyranian, et al., 1999).

In the example given in Fig. 13 some corners are ‘cut off’ by segmentation lines crossing a local symmetry axis. In some other examples (mostly in the difficult condition), we also observed that the outline was simply divided in two or in four, with the segmentation lines following the global shape diagonals or the main orientation axes (which were sometimes also symmetry axes). This was mainly the case when almost no m – points or only low-salient m – points were present and it could also be the result of the task demands (i.e. the subjects had to segment each outline at least once). In those cases, the outlines typically looked more like a shapeless convex ‘object’ or blob (e.g. lower half of Fig. 14). The resulting ‘parts’ then are also more convex blobs rather than




Fig. 15. The short-cut rule is sometimes overruled by a diagonal cut, connecting an $m-$ point to an $M+$ point, in the case of a joint (real elbow from outline drawing of an arm).

typical elongated parts protruding from the main body of the object (see also above: the positive correlation between compactness and salience in the ‘difficult’ condition).

In Fig. 14, we give an example of a matched stimulus pair (both outlines have an equal number of inflections). As was mentioned previously, the outline of the horse (from the easy condition) is segmented more often and also more consistently (bigger circles) than the outline of the lettuce (from the difficult condition). It is not clear whether this difference is due to the generally more salient $m-$ points in the easy condition or to the influence of knowledge about the object. The fact that most of the resulting parts of the horse correspond to known (existing) parts could suggest the influence of knowledge. It is also possible (in fact, quite likely; see Biederman, 1987; Hoffman & Richards, 1984; Hoffman & Singh, 1997) that salient segmentation points, and thus easy segmentation into parts, make the recognition of the object itself easier. So, the two factors (recognizability and salience of segmentation points) may very well go together. In retrospect, we noticed that most of the easy-to-recognize outlines had generally more salient points, whereas the difficult-to-recognize outlines had typically a more convex, blob-like shape with only a few salient points. The same is true with respect to more salient parts sticking out in the recognizable objects. Future research may attempt to dissociate the two factors, recognizability and more low-level visual factors such as point or part salience, if at all possible.

Fig. 15 shows an outline drawing of an arm, containing a real elbow. Geometric ‘elbows’ and L shapes had been tested experimentally before by Singh, Seyranian, et al. (1999). In contrast to these controlled, parametric experiments, in which the curvature did not play a role and the shortest cuts were preferred even for the smoothest elbows (Singh et al.’s Experiment 5), the part cuts now clearly follow the diagonal between the $m-$ point (on the interior of the arm’s bent) and the $M+$ point (on the exterior). This type of segmentation is in contradiction with the Hoffman and Richards’s (1984) claim that an $m-$ point will never be connected to an $M+$ point, but it was anticipated by a more recent theoretical analysis linking the diagonal to the local symmetry axis (e.g. Singh & Hoffman, 2001). The part line connecting an $m-$ point with an $M+$ point was also observed in cases where a long but narrow part of an object with relative little variation in width was bent rather strongly (e.g. the bent in the ear piece of the spectacles; Fig. 7C). Perhaps some people regard the parts at both sides of the bent as belonging to one bigger, formerly




Fig. 16. Sometimes one m -point and collinearity with the part line is enough to form a part.

straight part which has been bent to form a new functional part. This observation provides some empirical support to Leyton's (1989) process-based view of parts, in which parts are regarded as the outcome of causal processes like protrusion acting on the boundary of the shape, as well as two related ideas: first, the mathematically sound symmetry-curvature duality principle, stating that curvature singularities are the end-points of local symmetry axes (Leyton, 1987) and, second, the nice computational and psychophysical work on the so-called shape triangle with parts, protrusions, and bends (e.g. Kimia, Tannenbaum, & Zucker, 1995; Siddiqi, Kimia, Tannenbaum, & Zucker, 2001). However, it should not be taken as empirical support for Leyton's (1992) later theoretical work, in which these original ideas were stretched much beyond their original limits (see Hendrickx & Wagemans, 1999).

In Fig. 16, we show an example of a difficult-to-recognize object where one salient m -point is used to form a part (on the top right side of the object). The segmentation line is then collinear with one of the two line segments at either side of the m -point. This cannot be the result of cognitive influences since most subjects did not recognize the object (as a sled). So, collinearity also seems a general factor influencing the creation of a part line (it was also observed frequently in other examples). Again, this comes as no surprise because collinearity or 'good continuation' is a well-known Gestalt principle of perceptual grouping (e.g. Claessens & Wagemans, 2005a; Feldman, 1999; Kovács, 1996), which has been re-introduced in the literature several times (e.g. as co-circularity by Parent & Zucker, 1999; as relatability by Kellman & Shipley, 1991). Like symmetry, it is detected quickly, even when it is less than perfect (e.g. Boucart & Humphreys, 1992; Claessens & Wagemans, 2005b; Feldman, 1997; Field, Hayes, & Hess, 1993; Wagemans, Van Gool, Lamote, & Foster, 2000) and it is also known to affect shape perception and segmentation (e.g. Feldman, 1999; Singh & Hoffman, 1998, 1999; Singh, Hoffman, & Albert, 1999a). Note also that a smaller part (upper left) is cut off from the main body (see also Singh et al., 1996). It seems that smaller parts are cut off from its main body as if they are seen as added to it or having grown out of it.

5. General discussion

This study started out by formulating three aims. The first goal of this study was to supplement the large literature on object segmentation into parts with good (i.e. reliable and valid) benchmark data. The current literature consists of a wealth of different segmentation principles and rules for part formation but relatively little strong empirical evidence. Most of the experiments so far had used only a small number of subjects and stimuli (mostly meaningless outline shapes). Using outlines derived from Snodgrass and

Vanderwart's (1980) set of line drawings of everyday objects (De Winter & Wagemans, 2004), we could investigate segmentation of natural object outlines that were either easily recognizable or not, while still being able to test the relevance of geometric principles related to contour curvature properties (e.g. curvature singularities). Our data set consists of 88 outline stimuli, segmented more than 122 times on average, yielding a total set of more than 10,000 part cuts. We have analyzed this huge data set in several different ways and we have presented plenty of visual examples illustrating our findings. We make all of this material available on our website and we hope that shape theorists will use it extensively to test their ideas on shape and object segmentation. It is also possible to rely on our data to derive empirically validated object parts for other purposes. For example, using the techniques presented here (and data obtained in a similar fashion), we have been able to demonstrate that object naming is facilitated when primed with valid parts, compared to invalid parts or other contour fragments (De Winter & Wagemans, 2005a).

A second goal of our study was to analyze the role of previously proposed factors determining segmentation, including curvature singularities (especially, negative minima or $m-$) as well as more global factors (such as symmetry, elongation, proximity, good continuation) and cognitive influences. We will first summarize some of our main findings and we will then present a framework which we find useful to get a handle on all of the factors that are shown to determine segmentation of object outlines into parts. Although it should not be considered as a new theory of segmentation, providing this synthesis was the third goal of our study.

5.1. Summary of the main findings

The first factor that clearly influences segmentation has a rather local quality. The data clearly demonstrate the importance of $m-$ points and (more generally) points with negative curvature. The more salient the $m-$ point is, the more consistently it was used as segmentation point. Its salience is determined by the amount of normalized curvature, the form of the two flanking sides, and—as our data showed—by its scale (see Hoffman & Singh, 1997, for a detailed review of how its salience can be established and see Rosin, 2000, for an alternative way of calculating salience). Using another data set, acquired in a study asking subjects explicitly to mark salient points, we will examine and report these influences in more detail elsewhere (De Winter & Wagemans, 2005b). Our present data indicated that negative curvature in itself is also a strong attractor, especially when other influences are also present (as in the case of 'necks').

The other visual influences have a more global quality. The second visual influence that frequently contributes to the part lines is collinearity and curvilinearity. In this case, the part line is collinear or curvilinear with the object outline on at least one side of the part line. The stronger the degree of linearity, preferably with both sides, the higher its salience. Siddiqi et al. (1996) also pointed to the importance of a similar notion (i.e. 'good continuation') when they proposed their limb part. But in addition, we observed that two $m-$ points on two sides are not always necessary (as with a limb). One $m-$ point and good collinearity or curvilinearity on one side can also generate a part. This is an obvious relaxation of a rather strict rule, which is a regular finding in our data set.

Third, the shorter the distance between segmentation points, the more likely it is that a segmentation line will be formed. This is in agreement with the short-cut rule of Singh, Seyranian, et al. (1999). The distance itself specifies its salience (see also Kubovy et al., 1998). Note that this does not imply that longer part lines are automatically suppressed. Especially in the presence of other influences, parts with longer part lines can be formed in the same object as parts with shorter part lines. In this sense, we have found clear exceptions to the short-cut rule (e.g. Fig. 13). However, we do want to stress that shorter part lines lead more easily to more consistent parts. Siddiqi et al. (1996) had a similar measure incorporated in their definition of a neck. By relaxing their requirement for a maximally inscribed circle with minimal radius to a short distance between segmentation points in general, we can account better for the observed data. Sometimes the circle is not fully inscribed within the object shape (see Fig. 11H, a real-object version of a theoretical case made in earlier papers; e.g. Singh, Seyranian, et al., 1999, Fig. 16d). By having this more general definition, we can still form the observed intuitive necks while at the same time accounting for parts that look like necks but which have a circle that is not fully inscribed.

Fourth, symmetry and main orientation axes are frequently crossed when forming part lines. The stronger the symmetry or the more elongated the object, the higher its salience. The need to cross a symmetry axis was previously also stated for the short-cut rule (Singh, Seyranian, et al., 1999) and plays a central role in axis-based models (e.g. Burbeck & Pizer, 1995) and models stressing the importance of smoothed local symmetry (e.g. Kimia et al., 1995; Siddiqi & Kimia, 1995; Siddiqi et al., 1996, 2001). Local symmetry axes are also sometimes crossed to cut-off parts or even followed to divide the object in equal pieces (as in cutting a pie). This occurred more often for difficult-to-recognize blob-like outlines. It seems that pieces and visual parts are sometimes conceptually mixed-up (perhaps due to task demands). Also the main orientation seems to play an important role. But instead of crossing it, it is also sometimes followed so that a part sticking out from the main body gets cut-off (a good example is a tooth brush where the brush is separated from its main body, the handle). The role of symmetry and elongation axes in shape perception has been established previously (e.g. Davi & Proffitt, 1993; Quinlan, 1995; Quinlan & Humphreys, 1993) and the connection between symmetry detection and part formation has also been examined explicitly before (e.g. Baylis & Driver, 1995; Bertamini, Friedenber, & Kubovy, 1997; Driver, Baylis, & Rafal, 1992). Another reason to expect local symmetry axes to play a role is that they always terminate at positive maxima (Leyton, 1987). If the part line then crosses the symmetry axis orthogonally, the resulting part will be a good, symmetrical part (see also Kurbat, 1994a). Our research confirms the relevance of these principles.

Fifth, cognitive knowledge about the object and its parts can form new parts, strengthen parts or even sometimes override visual parts. Thus, it can create extra parts with little or no visual support, or it can strengthen other factors or even override visual influences. This comes as no surprise if one considers how important parts are in cognitive tasks like object categorization and naming (e.g. Schyns and Murphy, 1994; Tversky & Hemenway, 1984), even in children (e.g. Giralt & Bloom, 2000; Landau et al., 1998). Cognitive parts are learned from previous experiences with objects. These parts frequently have names, like the hand of a human with all its fingers, and can be based on their function. In the case of outlines, 'cognitive' influences may be needed to

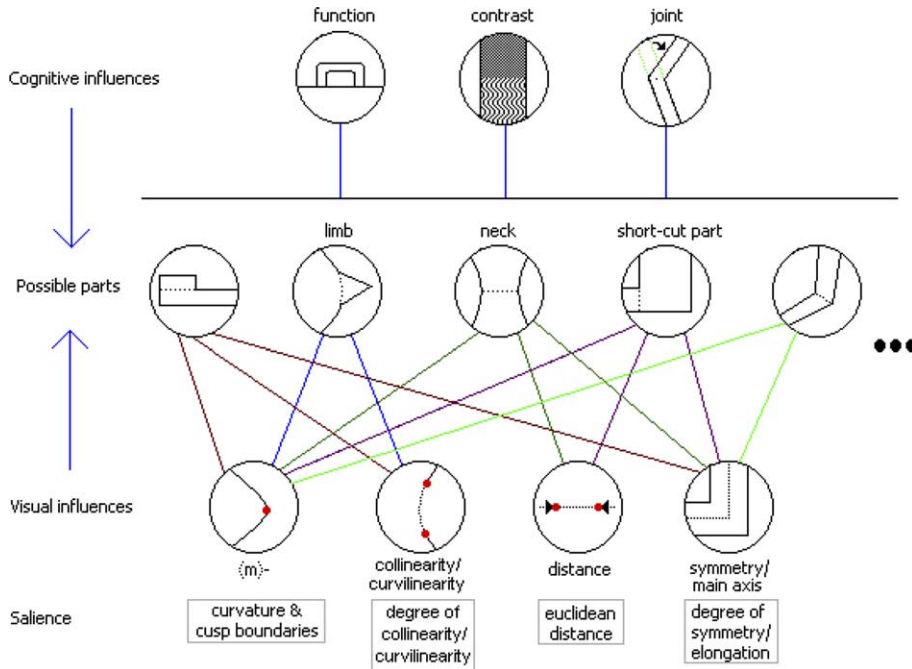


Fig. 17. Framework to understand the interactions between the different factors influencing the segmentation of outlines into parts (see text for further explanation).

replace visual cues which happen to be invisible in the outline itself (e.g. high contrast in texture, shade of grey or color). Although we did not directly observe instances of this, Nilson, Joneleit, Smith, and Hoffman (2001) demonstrated the effect that color can have on the learning of perceived parts. Finally, a joint can also create parts, like in the elbow of the arm. In static depictions, the movement of a joint, which might be visible in more natural conditions, can only work cognitively, triggered by the remaining visual traces (e.g. smoothed local symmetry and curvature singularities). So, most of what we call ‘cognitive influences’ would work as visual influences in the case of natural objects (as opposed to impoverished outlines). This suggests that the importance of cognitive influences should normally not be overestimated.

5.2. Integrative framework

We are now ready to try to synthesize our findings and to embed them into what we know from the literature so far. We propose the framework in Fig. 17 to do so. As reviewed in the first part of this paper, two major traditions concerning object segmentation are (1) models postulating a limited set of basic shape primitives (e.g. geons; Biederman, 1987) and (2) models in which parts are not predefined but are created by applying geometric rules. Some of these rules concern only individual segmentation points (e.g. minima rule; Hoffman & Richards, 1984), while others stipulate how parts come about by combining

local cues to segmentation with more global factors (e.g. limbs and necks; Siddiqi et al., 1996; short-cut rule; Singh, Seyranian, et al., 1999).

We think it is useful to distinguish these two levels explicitly and to show how they interact. The level marked by ‘visual influences’ represents the different local, quasi-local and more global influences known to influence object segmentation and part formation. The level marked by ‘possible parts’ lists the different types of parts previously identified in the literature, as well as some others that result from a further relaxation of the principles. For example, the part at the left does not follow the shortest-cut rule but it makes use of the minima of curvature (or the cusps) and it is probably influenced by good continuation and the object’s main elongation. The lines connecting the two levels clarify how the different parts arise from the combination of several low-level influences. For example, a limb results from a combination of collinearity and local curvature singularities (or cusps), while a neck makes use of points with negative curvature (not necessarily minima) where the maximally inscribed circle has the smallest possible diameter (which we relax further to the shortest possible distance). The short-cut rule connects local minima (or cusps) with other contour points using the shortest possible cut that crosses a local symmetry axis. The elbow segmentation does not use the shortest cut but it connects two curvature singularities and makes use of principles related to smoothed local symmetry. To indicate that this combination of geometric rules does not yield an exhaustive list of part types, we have added three points on the right of this diagram.

Another feature of our framework is that it acknowledges explicitly the active interplay between these factors. The strength of each factor varies (in a way that is sometimes easy to quantify, e.g. proximity and collinearity), and this variation leads to intricate interactions. The joint presence of these visual influences leads to more consistent segmentations and a relatively weak factor can be compensated by another, stronger one. This idea of interactions between different geometric principles has been discussed before (e.g. Siddiqi et al., 1996; Singh & Hoffman, 2001; Singh, Seyranian, et al., 1999) but we have now established clear empirical evidence for it. For example, a part cut using a very strong (deep) negative minimum can overcome a longer distance just as a very narrow neck does not need a local minimum. Likewise, segmentation can be pulled away from a local minimum when the neighboring contour segments are very strongly co-circular or the resulting parts themselves are very regular (symmetric). We believe it is very important to address these interactions in future experimental and computational work because it is a very characteristic feature of human segmentation of natural shapes.

Finally, our framework adds cognitive influences because their role was also clearly visible in our results. Usually, these top-down factors seem to moderate the interplay between the low-level visual influences and it is not clear how important they would be in real-world objects rather than impoverished outline stimuli.

5.3. Future work

Several questions still remain. First, what is the relative contribution of each of the influences (apart from its salience)? It seems unlikely that each influence has an equal contribution. It appears very probable, for instance, that m - points have a higher

influence than the other influences (e.g. they might be necessary although not sufficient). The contribution of each influence can be examined in future research by varying the salience of each influence while keeping the others as constant as possible. This necessarily requires parametric variation of much more simple, less naturalistic shapes, an approach which we consider complementary to the present one.

Second, it is not clear whether recognition leads to more consistent segmentation due to cognitive influences or due to the higher salience of the relevant factors, which was generally observed in the case of known objects. As suggested above, it is our impression that the influence of top-down factors was relatively moderate and so the higher consistency would then be caused more by the higher salience of the other factors. However, we do not want to exclude the possibility that more salient parts make objects easier to recognize (Hoffman & Singh, 1997) and so the two types of factors would always go together. This could be investigated by more carefully balancing the easy-and-difficult-to-recognize outline pairs (if at all possible). Alternatively, one could use nonsense shapes, instead of difficult-to-recognize outlines, which are easier to match to the easy-to-recognize outlines in terms of visual factors. However, to be able to create two groups of visually similar outlines, one would also have to know how to measure the salience of all the factors separately. Although our data contain suggestions on what makes each factor more salient, a more accurate way of establishing or measuring the salience of these factors lies outside the scope of this study.

Third, it is possible that the salience of the part itself contributes to the salience of the part cut. By this, we mean that the salience of the resulting part could be an extra influence on the salience of the part cut: Part cuts then become even more salient when more salient parts are formed. This is impossible to determine from the present study, but it could be examined by varying the salience of the part (e.g. varying the ratio of the part area relative to the length of the cut, following the suggestion by Hoffman & Singh, 1997) while holding the other visual influences constant as much as possible. Again, such parametric experiments would constitute a valuable addition to our present line of work.

In sum, the present framework integrates the previously proposed and currently well-established influences on segmentation and part formation. Our integrative framework offers a refreshing way of looking at shape segmentation as a whole by giving the different influences a central role (instead of possible predefined parts) and by allowing them to interact with each other to form parts and thus allowing to form parts that have not been described in previous models. The interaction of all influences is determined by their respective salience, resulting in a strengthening of weakening of each other's force. This integrative framework thus enables future research to study segmentation of shape in a more diverse and coherent manner than the past stand-alone models. Future, more parametric research with simpler, less naturalistic shapes will be needed to determine the weight of each of the factors separately and study how they interact precisely in more complex shapes. Only then, the framework might be developed further into a computational model, leading to new, quantitative predictions. Our present work has only provided some of the foundations for this future work by thoroughly examining a large and reliable benchmark data set.

Acknowledgements

This research was supported by a research grant from the University Research Council (OT/00/007) and from the Fund for Scientific Research (FWO-Vlaanderen G.0189.02) to JW. This study is part of a larger research program with financial support from the University Research Council (GOA/2005/03-TBA) to the Laboratory of Experimental Psychology. We would also like to thank Felix Wichmann, Sven Panis, Donald Hoffman and an anonymous reviewer for their helpful comments on a previous draft.

We would like to encourage other researchers to use our benchmark data set, made available on the following website: <http://www.psy.kuleuven.be/~winterj/Segmentation/index.html>. We would be happy to assist.

Appendix A. The determination of singularities and outline-distance

We applied the following method to determine the singularities. First, the mathematical extremes were selected. Three procedures were applied to avoid that noise consisting of very small changes in curvature or consisting of irregular changes on very short parts of the outline would influence our later analyses too much. All the threshold values that we used were chosen after checking the effect of each on a representative sample of the outline figures. Moreover, in the end, all outlines were rechecked (with the chosen thresholds) so that only a minimum of unwanted side effects occurred. (1) All singularities with an absolute curvature value smaller than 0.01 were considered as zero curvature. (2) For consecutive extremes that differ in curvature value only by a small amount, we selected only those that were sufficiently different from the preceding or following extreme. In an effort to incorporate some global shape information, this threshold was adjusted for each outline and was set to half of the average curvature value. (3) Linear smoothing of the curvature value was applied such that the smoothed value for a point on the outline was the average of the point itself and the three points before and three points after that point.

Once all extremes were selected in this manner, the inflections were determined. Between two extremes with opposite sign, the curvature should be zero at least once. Since we are working with discrete points, however, a point with zero curvature does not always exist. The point with the value closest to zero was then selected. When several curvature points were inflection candidates (due to applying the above noise reduction), the middle value was chosen.

To summarize, the following calculations were made for each outline in the following order: Linear smoothing, absolute thresholding, average curvature, serial thresholding, and inflections.

The outline-distance from a segmentation point to a singularity is the distance between them, travelled along the outline. Because we are working with discrete points connected by very short straight lines (typically less than 1 pixel long), this is easy to calculate. We calculated the Euclidean distance in pixels from point to point and accumulated this along the outline. So, point 1 had the value of the Euclidean distance from point 0 to point 1. Point 2 had the value of point 1 added with the distance between point 1 and 2, etc.

Table A1

Adjusting the number of segmentation points for the unequal distribution of curvature sign

Curvature sign	Segmentation points	Count	Ratio	Adjusted count	Adjusted percent (%)
Negative	16,144	35,977	2.65	42,686	82.6
Positive	5600	59,325	1.61	8996	17.4
Total	21,714	95,302		51,682	100.0

Appendix B. Adjusting the frequency data for the unequal number of singularities or points of a particular curvature sign

To adjust for the imbalance in the distribution of the different types of singularities or points of different curvature sign, we calculated the adjusted percentage as follows (see Table A1). In the example, we calculate the adjusted percentage for the number of segmentation points that have positive or negative curvature. First, we counted the number of segmentation points (column ‘Segmentation points’) of each particular type (here, negative and positive curvature). Next, the raw numbers of available points of each particular type (column ‘Count’) were determined (here, the number of pixels with either negative or positive curvature). Then, for each particular type, the total was divided by the count for that type (column ‘Ratio’). The ‘Adjusted count’ column represents the number of segmentation points multiplied by that ratio. Finally, we calculated the ‘Adjusted percent’ column by calculating the ratio for each type in ‘Adjusted count’ as a percentage of the total. So, the total for adjusted percent is always 100%.

Appendix C. Calculating the segmentation popularity of points

First, we calculated the segmentation frequency of each point on the outline. This is the number of times that a point is marked as a segmentation point; it ranges from zero to the number of subjects that segmented that outline. We then applied a Gaussian smoothing technique on the segmentation frequency to avoid noise. The smoothing technique calculated for point A the weighted average of the frequency of point A and of points in the neighborhood of A. The weighting was done by a Gaussian function; the further the points are away from point A the less influence they have on the frequency value of point A. We used a Gaussian function because we considered it more and more unlikely that a subject meant to mark point A the further it was away from point A, and we assumed that this likelihood was distributed according to a Gaussian function. The width of the Gaussian curve (i.e. the smoothing parameter) is expressed in image pixels and represents the standard deviation of the curve. With the threshold set too high, there is an increasing chance that two or more separate segmentation points are smoothed together to become one. With the threshold set too low, too much noise persists in the data and several maxima of popularity will be located in close proximity. Both are undesired and, after extensive visual inspection, a value of 5 was chosen as a good compromise avoiding both of these

drawbacks. Within a certain range, this chosen value is not critical to the outcome of our analyses (e.g. 4 or 6 would give similar results).

Thus, after smoothing, we have a smoothed segmentation popularity value for every point on the outline (called ‘segmentation popularity’). The higher this value, the more subjects have chosen that point (or points nearby) as a segmentation point of a part cut.

Appendix D. Calculating the scale-space representation and the scale

Before being able to construct the scale-space representation, we first had to calculate Gaussian smoothed outlines of all outlines with correction for shrinkage.⁵ We used the technique as described by Lowe (1988). In short, each pixel is convoluted with a Gaussian function over a range of standard deviations (σ). σ had a starting value of 1 and it increased by a factor of $\sqrt{2}$ for each subsequent larger smoothing. The curvature was also calculated as described by Lowe (1988): $k = X'Y'' - Y'X'' / (X'^2 + Y'^2)^{3/2}$ with X' , Y' and X'' , Y'' being the convolutions with the first and second derivatives of the Gaussian function, respectively, (for a detailed description, see Lowe, 1988). We stopped smoothing when the *next* outline no longer had inflections. Hence, we selected the level right before the shape became totally convex. This seems to be a good stop criterion since the next scale has only positive curvature and thus has little information of interest. Then, each smoothed outline was corrected for shrinkage (Lowe, 1988). This has no effect on the configuration of the shape, but the overall size of the shape is much more in keeping with the original shape.

We could then construct a scale-space representation as described by Witkin (1986). In short, each singularity on the outline is connected to a corresponding singularity at the next larger scale and creates the so-called ‘contour’ (Witkin, 1986). We will call this the ‘Witkin-contour’ to avoid confusion with the outline shape or shape contour as such. The only restriction that we imposed was that the corresponding singularities had to be within a 10-pixel range from scale to scale (the limit used by Witkin is not reported). Otherwise, the singularity was considered as a point of a different Witkin-contour. This scale-space allowed us to determine the length of each Witkin-contour, called ‘scale’. It is defined as the spatial scale above which the Witkin-contour disappears and it is expressed as the log function of σ . This scale is considered to reflect a single physical event (Witkin, 1986), for instance, the growth of a convex part (Leyton, 1989). Finally, the true location of a Witkin-contour is given by its location at the finest (non-smoothed) scale.

Appendix E. Selecting highly salient part cuts

First, we calculated the segmentation popularity as described in Appendix C. Then, we calculated the local maxima of the segmentation popularity (popularity maxima) and

⁵ We would like to thank Paul Rosin for supplying his program code for smoothing outlines on his website. This code has been adapted for closed outlines.

considered these (anchor) points as segmentation points if their popularity was greater than or equal to 10 (i.e. a first popularity threshold, chosen after extensive testing, see below). Subsequently, we shifted each of the two end points of a part cut towards its closest anchor point if it was located within an outline-distance of 10 pixels. This procedure results in a number of part cuts being located in the exact same position, allowing us to calculate the frequency of each part cut (i.e. the number of subjects that marked that same part cut). Part cuts with a frequency lower than 10 (i.e. a second threshold) were discarded. Note that we have two thresholds (both with a value of 10): one for selecting anchor points based on the segmentation popularity of points and one for selecting part cuts based on the number of subjects that had marked the same part cut.

However, sometimes two anchor points were located in (very) close proximity. This has the effect that neighboring segmentation points are divided among the surrounding adjacent anchor points and thus lowering the number of segmentation points shifted to each anchor point. This results in two closely located part cuts both with a lowered frequency and it could imply that none of the two part cuts were selected while intuitively one expects a selection of at least one part cut. To avoid this, only the anchor point with the highest popularity was retained when two anchor points were located within a distance of 20 pixels. We chose 20 pixels because this is the minimum distance between two anchor points such that they would not compete with each other (since the distance allowed to shift a segmentation point to an anchor point was 10 pixels). We applied this correction before shifting end points towards anchor points.

When choosing the first threshold value (for the anchor points) the following aspects were considered. Too high a threshold results in the selection of only very highly salient segmentation points and automatically results in only a small subset of the data, which would not reflect the total data well enough. Too low a threshold results in an indiscriminating selection (i.e. too many low-salient segmentation points or noise). Looking at each outline in turn, we observed that many segmentation points had either a value lower than 5 or a value higher than 10. So, a threshold of 10 was chosen. The selection of the second threshold value was based on similar observations, resulting in a chosen threshold of 10.

Appendix F. Calculating the scale of a part cut

While it is clear how to determine the scale of a single event, it is not clear how one determines the scale of multiple points. Since we prefer to report the tuples for a higher rather than a lower scale (which would yield longer, less informative tuples), we decided to take the highest scale of the two segmentation points defining a part cut as the scale of the part cut. If one of the segmentation points has no singularity associated with it, this has the additional advantage that the scale of the other associated singularity automatically becomes the scale of the part cut. If both segmentation points had no singularity associated with them, then the assigned scale is zero (since it is not possible to calculate a scale for an arbitrary point). This happened only for six part cuts, so it is hardly affecting the analyses.

References

- Bar, M. (2003). A cortical mechanism for triggering top–down facilitation in visual object recognition. *Journal of Cognitive Neuroscience*, *15*, 600–609.
- Barenholtz, E., Cohen, E. H., Feldman, J., & Singh, M. (2003). Detection of change in shape: An advantage for concavities. *Cognition*, *89*, 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.
- 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*, 1323–1342.
- Bertamini, M., & Croucher, C. J. (2003). The shape of holes. *Cognition*, *87*, 33–54.
- Bertamini, M., & Farrant, T. (2005). Detection of change in shape and its relation to part structure: An advantage for changes in part structure. *Acta Psychologica*, in press.
- Bertamini, M., Friedenberg, J. D., & Kubovy, M. (1997). Detection of symmetry and perceptual organization: The way a lock-and-key process work. *Acta Psychologica*, *95*, 119–140.
- Bertamini, M., & Mosca, F. (2004). Early computation of contour curvature and part structure: Evidence from holes. *Perception*, *33*, 35–48.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Blum, H. (1973). Biological shape and visual science (Part I). *Journal of Theoretical Biology*, *38*, 205–287.
- Boucart, M., & Humphreys, G. W. (1992). The computation of perceptual structure from collinearity and closure: Normality and pathology. *Neuropsychologia*, *30*, 527–546.
- Braunstein, M. L., Hoffman, D. D., & Saidpour, A. (1989). Parts of visual objects: An experimental test of the minima rule. *Perception*, *18*, 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.
- Burbeck, C. A., & Pizer, S. M. (1995). Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research*, *35*, 1917–1930.
- Cave, C. B., & Kosslyn, S. M. (1993). The role of parts and spatial relations in object identification. *Perception*, *22*, 229–248.
- Claessens, P. M., & Wagemans, J. (2005a). Perceptual grouping in Gabor lattices: Proximity and alignment. *Perception and Psychophysics*, in press.
- Claessens, P. M., & Wagemans, J. (2005b). Perceptual grouping in discrete periodic patterns and perturbed zigzag lattices: Proximity and collinearity, submitted for publication.
- Davi, M., & Proffitt, D. R. (1993). Frames of reference and distinctive figural characteristics affect shape perception. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 867–877.
- De Winter, J., & Wagemans, J. (2004). Contour-based object identification and segmentation: Stimuli, norms and data, and software tools. *Behavior Research Methods, Instruments, and Computers*, *36*(4), 604–624.
- De Winter, J., & Wagemans, J. (2005a). Part priming of object naming, submitted for publication.
- De Winter, J., & Wagemans, J. (2005b). Perceptual salience of points along the contour of everyday objects: A large-scale study, submitted for publication.
- Dhandapani, R., & Kimia, B. B. (2002). Role of scale in partitioning shape. *Proceedings of the International Conference on Image Processing II*, 565–568.
- Driver, J., & Baylis, G. C. (1995). One-sided edge-assignment in vision: 2. Part decomposition, shape description, and attention to objects. *Current Directions in Psychological Science*, *4*, 201–206.
- Driver, J., & Baylis, G. C. (1996). Edge-assignment and figure-ground segmentation in short-term visual masking. *Cognitive Psychology*, *31*, 248–306.
- Driver, J., Baylis, G. C., & Rafal, R. D. (1992). Preserved figure-ground segregation and symmetry perception in visual neglect. *Nature*, *360*, 73–75.
- Elder, J. H., & Zucker, S. W. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, *33*, 981–991.
- Elder, J. H., & Zucker, S. W. (1998). Evidence for boundary-specific grouping. *Vision Research*, *38*, 143–152.

- Feldman, J. (1997). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research*, 37, 2835–2848.
- Feldman, J. (1999). The role of objects in perceptual grouping. *Acta Psychologica*, 102, 137–163.
- Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review*, 112, 243–252.
- Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local ‘association field’. *Vision Research*, 33, 173–193.
- Giralt, N., & Bloom, P. (2000). How special are objects? Children’s reasoning about objects, parts, and holes. *Psychological Science*, 11, 497–501.
- Hendrickx, M., & Wagemans, J. (1999). A critique of Leyton’s theory of perception and cognition. *Journal of Mathematical Psychology*, 43, 314–345.
- Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, 36, 791–804.
- Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition*, 18, 65–96.
- Hoffman, D. D., & Singh, M. (1997). Saliency 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, 162–174.
- Humphreys, G. W., & Müller, H. (2000). A search asymmetry reversed by figure-ground assignment. *Psychological Science*, 11, 196–201.
- Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, 23, 141–221.
- Kimia, B. B., Tannenbaum, A. R., & Zucker, S. W. (1995). Shapes, shocks, and deformations: I. The components of shape and the reaction-diffusion space. *International Journal of Computer Vision*, 15, 189–224.
- Koenderink, J. J., & van Doorn, A. J. (1982). The shape of smooth objects and the way contours end. *Perception*, 11, 129–137.
- Kovács, I. (1996). Gestalten of today: Early processing of visual contours and surfaces. *Behavioural Brain Research*, 82, 1–11.
- Kovács, G., Sáry, G., Köteles, K., Chadaide, Z., Tompa, T., Vogels, R., & Benedek, G. (2003). Effects of surface cues on macaque inferior temporal cortical responses. *Cerebral Cortex*, 13, 178–188.
- Kubovy, M., Holcombe, A. O., & Wagemans, J. (1998). On the lawfulness of grouping by proximity. *Cognitive Psychology*, 35, 71–98.
- Kubovy, M., & Wagemans, J. (1995). Grouping by proximity and multistability in dot lattices: A quantitative Gestalt theory. *Psychological Science*, 6, 225–234.
- Kurbat, M. A. (1994a). Structural description theories: Is RBC/JIM a general-purpose theory of human entry-level object recognition? *Perception*, 23, 1339–1369.
- Kurbat, M. A. (1994b). A network model for generating differential symmetry axes of shapes via receptive fields. *Spatial Vision*, 8, 433–442.
- Lamberts, K., & Freeman, R. P. J. (1999). Building object representations from parts: Tests of a stochastic sampling model. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 904–926.
- Lamote, C., & Wagemans, J. (1999). Rapid integration of contour fragments: From simple filling-in to parts-based shape description. *Visual Cognition*, 6, 345–361.
- Landau, B., Smith, L., & Jones, S. (1998). Object perception and object naming in early development. *Trends in Cognitive Sciences*, 2, 19–24.
- Leyton, M. (1987). Symmetry-curvature duality. *Computer Vision, Graphics, and Image Processing*, 38, 327–341.
- Leyton, M. (1989). Inferring causal history from shape. *Cognitive Science*, 13, 357–387.
- Leyton, M. (1992). *Symmetry, causality, mind*. Cambridge, MA: MIT Press.
- Lowe, D.G. (1988). Organization of smooth image curves at multiple scales. *Proceedings of the Second International Conference on Computer Vision* (pp. 558–567). Tampa, FL.
- Marr, D. (1982). *Vision: A computational investigation into the human presentation and processing of visual information*. San Francisco, CA: W.H. Freeman.
- Nilson, C., Joneleit, K., Smith, R., & Hoffman, D. D. (2001). Color and part [Abstract]. *Journal of Vision*, 1(3), 91a.

- Olivers, C. N. L., & Van der Helm, P. A. (1998). Symmetry and selective attention: A dissociation between effortless perception and serial search. *Perception and Psychophysics*, *60*, 1101–1116.
- Parent, P., & Zucker, S. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *PAMI-11*, 823–839.
- Pasupathy, A., & Connor, C. E. (2002). Population coding of shape in area V4. *Nature Neuroscience*, *5*(12), 1332–1338.
- Pinker, S. (1984). Visual cognition: An introduction. *Cognition*, *18*, 1–63.
- Quinlan, P. T. (1995). Evidence for the use of scene-based frames of reference in two-dimensional shape recognition. *Spatial Vision*, *9*, 101–125.
- Quinlan, P. T., & Humphreys, G. W. (1993). Perceptual frames of reference and two-dimensional shape recognition: Further examination of internal axes. *Perception*, *22*, 1343–1364.
- Richards, W., & Hoffman, D. D. (1985). Codon constraints on closed 2D shapes. *Computer Vision, Graphics, and Image Processing*, *31*, 265–281.
- Richards, W. A., Koenderink, J. J., & Hoffman, D. D. (1987). Inferring three-dimensional shapes from two-dimensional silhouettes. *Journal of the Optical Society of America*, *A4*, 1168–1175.
- Rom, H., & Medioni, G. (1993). Hierarchical decomposition and axial shape description. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *PAMI-15*, 973–981.
- Rosin, P. L. (2000). Shape partitioning by convexity. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, *30*, 202–210.
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, *80*, 1–46.
- Schyns, P. G., Goldstone, R. L., & Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, *21*, 1–54.
- Schyns, P. G., & Murphy, G. L. (1994). The ontogeny of part representation in object concepts. In D. L. Medin, *The psychology of learning and motivation: Advances in research and theory* (Vol. 31) (pp. 305–349). San Diego, CA: Academic Press, 305–349.
- Siddiqi, K., & Kimia, B. B. (1995). Parts of visual form: Computational aspects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *17*, 239–251.
- Siddiqi, K., Kimia, B. B., Tannenbaum, A. R., & Zucker, S. W. (2001). On the psychophysics of the shape triangle. *Vision Research*, *41*, 1153–1208.
- Siddiqi, K., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception*, *25*, 399–424.
- Singh, M., & Hoffman, D. D. (1998). Part boundaries alter the perception of transparency. *Psychological Science*, *9*, 370–378.
- Singh, M., & Hoffman, D. D. (1999). Completing visual contours: The relationship between relatability and minimizing inflections. *Perception and Psychophysics*, *61*, 943–951.
- Singh, M., & Hoffman, D. D. (2001). Part-based representations of visual shape and its implications for visual cognition. In: T.F. Shipley, & P. J. Kellman, *Advances in psychology. From fragments to objects: Segmentation and grouping in vision* (vol. 130) (pp. 401–459). Amsterdam: Elsevier.
- Singh, M., Hoffman, D. D., & Albert, M. K. (1999). Contour completion and relative depth: Petter's rule and support ratio. *Psychological Science*, *10*, 423–428.
- Singh, M., Seyranian, G. D., & Hoffman, D. D. (1996). *Cuts for parsing visual shapes* (Memo 96-33). University of California, Irvine, Institute for Mathematical Behavioral Sciences.
- Singh, M., Seyranian, G. D., & Hoffman, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception and Psychophysics*, *61*, 636–660.
- Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory*, *6*, 174–215.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, *113*, 169–193.
- Vaina, L. M., & Zlateva, S. D. (1990). The largest convex patches: A boundary-based method for obtaining object parts. *Biological Cybernetics*, *62*, 235–236.
- Van der Helm, P. A., & Leeuwenberg, E. L. J. (1996). Goodness of visual regularities: A nontransformational approach. *Psychological Review*, *103*, 429–456.

- Van Lier, R., & Wagemans, J. (1998). Effects of physical connectivity on the representational unity of multi-part configurations. *Cognition*, *69*, B1–B9.
- Van Lier, R., & Wagemans, J. (1999). From images to objects: Global and local completions of self-occluded parts. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 1721–1741.
- Vecera, S. P., Behrmann, M., & Filapek, J. C. (2001). Attending to the parts of a single object: Part-based selection limitations. *Perception and Psychophysics*, *63*, 308–321.
- Vecera, S. P., Behrmann, M., & McGoldrick, J. (2000). Selective attention to the parts of an object. *Psychonomic Bulletin and Review*, *7*, 301–308.
- Wagemans, J. (1995). Detection of visual symmetries. *Spatial Vision*, *9*, 9–32.
- Wagemans, J. (1997). Characteristics and models of human symmetry detection. *Trends in Cognitive Sciences*, *1*, 346–352.
- Wagemans, J. (1999). Toward a better approach to goodness: Comments on Van der Helm and Leeuwenberg (1996). *Psychological Review*, *106*, 610–621.
- Wagemans, J., De Winter, J., Op de Beeck, H., Ploeger, A., Beckers, T., & Van Horebeek, J. (2005). Identification of silhouette and contour versions of line drawings of everyday objects, in preparation.
- Wagemans, J., Notebaert, W., & Boucart, M. (1998). Lorazepam but not diazepam impairs identification of pictures on the basis of specific contour fragments. *Psychopharmacology*, *138*, 326–333.
- Wagemans, J., Van Gool, L., Lamote, C., & Foster, D. H. (2000). Minimal information to determine affine shape equivalence. *Journal of Experimental Psychology: Human Perception and Performance*, *26*, 443–468.
- Witkin, A.P. (1986). Scale space filtering. In: A.P. Pentland (Ed.), *From pixels to predicates: Recent advances in computational and robotic vision* (pp. 5–19). Norwood, NJ: Ablex.
- Xu, Y., & Singh, M. (2002). Early computation of part structure: Evidence from visual search. *Perception and Psychophysics*, *64*, 1039–1054.
- Zusne, L. (1970). *Visual perception of form*. London: Academic Press.