ON THE ANATOMY OF VISUAL FORM

Kaleem Siddiqi
Kathryn J. Tresness
Benjamin B. Kimia

Brown University
Laboratory for Engineering Man-Machine Systems
Providence, RI, USA 02912

ABSTRACT

Part based representations allow for recognition that is robust in the presence of occlusion, movement, growth, and deletion of portions of an object. We propose a general "form from function" principle arising from the interactions of objects in their environment, which, together with properties of visual projection, gives rise to two kinds of parts: LIMB-BASED parts arise from a pair of negative curvature minima with evidence for "good continuation" of boundaries on one side; NECK-BASED parts arise from narrowings in shape. We then test this hypothesis by requiring subjects to partition a variety of biological and nonsense 2D shapes into perceived components. We examine: 1) whether a subject determines components consistently across different trials of the same partitioning task, 2) whether there is evidence for consistency between subjects for the same partitioning task, and 3) how the perceived parts compare with the parts proposed by the "form from function" principle. The results are interpreted as suggesting that there are high levels of both intra-subject as well as inter-subject consistency and that a large majority of the perceived parts do in fact correspond to the proposed limbs and necks. Our proposal for parts envisions a role for them as an intermediate representation between local image features, *e.g.*, edges, and global object models; the implications are discussed in relation to figure/ground segmentation, and certain visual illusions.

1 Introduction

Part-based representations allow for recognition that is robust in the presence of occlusion, movement, deletion, or growth of portions of an object. In the task of forming high-level object-centered models from low-level image-based features, parts can serve as an intermediate representation. Evidence for part-based representations in human vision is strong [1, 11]. The computer vision literature is also rich with approaches to partitioning shape, see [14, 13] for a review. For example, Koenderink and van Doorn observe that as evidenced in works of art, elliptic (concave or convex) regions in three-dimensional (3D) objects have a 'thing-like' character whereas hyperbolic (saddle-shaped) regions act only as the glue that holds them together [9]. The dividing lines between these two types of regions are called parabolic lines, the projection

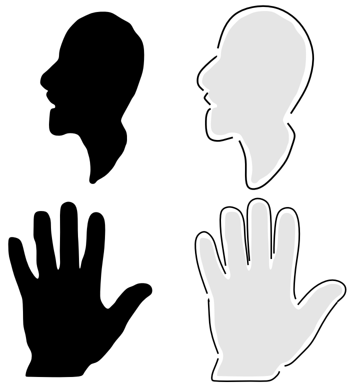

Figure 1: Hoffman and Richards' theory of curve partitioning segments a contour at negative curvature minima. The figure on the right of each pair depicts the parts as solid lines, with the shape shown in gray.

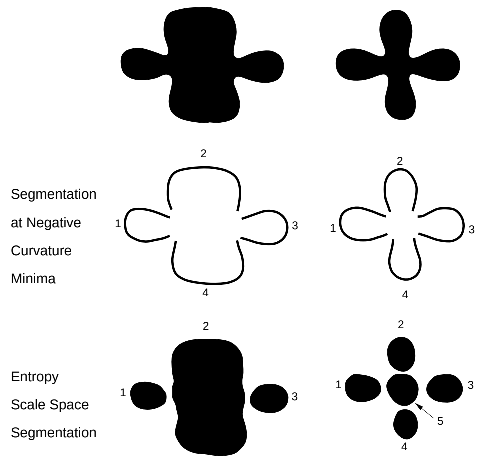

Figure 2: The shape on the top right is stretched through its region to create the shape on the top left. The minima rule does not reflect the change in percept to a shape with 3 parts (middle), whereas Kimia *et al.*'s proposal for parts at "necks" [7] does (bottom).

of which leads to inflection points on the shape's contour [8]. Hence, they suggest that 3D objects should be segmented along parabolic lines, and two-dimensional (2D) shapes at inflection points of the contour. Hoffman and Richards have pointed out that this proposal does not explain the reversal of perceived parts in a reversed 2D shape, or the perception of parts for everywhere parabolic surfaces [3]. Alternatively, they propose a theory of parts that relies on two "regularities of nature", in contrast to relying on the shapes of parts as described by primitives, leading to the following partitioning rule for plane curves: "divide a plane curve into parts at negative minima of curvature," *i.e.*, at points where the curve is concave and the amount of bending is a local maximum. This rule successfully explains several figure-ground reversal illusions [3] and can lead to intuitive part decompositions, Figure 1. Braunstein *et al.* have carried out a psychophysical test of this rule [2].

Whereas the minima rule provides intuitive parts for the above cases, for others it leads to unintuitive parts: we present three classes of examples. First, observe that the perceived parts may change when a portion of the contour of a shape is stretched, while the remaining contours are left intact [7], Figure 2. Second, observe that when a portion of the boundary of a shape is shifted with respect to another portion, Figure 3 (bottom), narrow regions that are no longer bound by negative curvature minima may be perceived as part boundaries [6]. This partitioning at "necks" was suggested by the computational framework of Kimia *et al.* [7] where necks correspond to "second order shocks" that arise in the course of deforming shape. Note that this argues for a role for interaction of a shape's boundary points through its region. Third, observe that when the boundaries associated with a negative curvature minima pair do not show evidence for "continuation", Figure 3 (top), the decomposition suggested by the minima rule is not intuitive [13]. Again, note the necessary role for regional

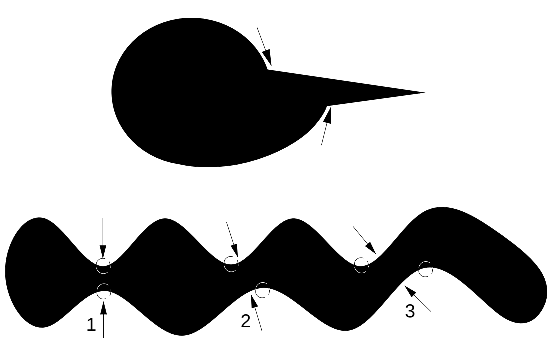

Figure 3: TOP: In the absence of "good continuation", a decomposition based on a pair of negative curvature minima may not be intuitive. BOTTOM: When a portion of the boundary of a shape is shifted with respect to another portion, narrow regions that are no longer bound by negative curvature minima may be perceived as part boundaries [6]. The negative curvature minima are marked by circles.

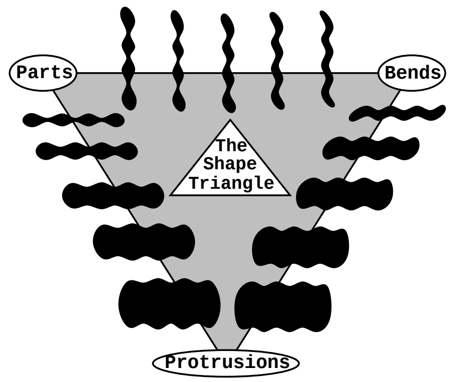

Figure 4: The nodes of the *shape triangle* represent three cooperative/competitive processes acting on shape, namely, *parts*, *protrusions*, and *bends* [6]. The sides of the triangle represent continua of shapes whose extremes correspond to the node and whose perception varies from one description to another.

interaction of boundary features.

Leyton [10, 11] contrasts approaches such as Hoffman and Richards' minima rule, where parts are viewed as rigid segments to be broken off, with a *process-based* view of parts. This view provides a causal explanation of a part as the outcome of a historical "protrusion" process acting on the shape's boundary, based on a biologically relevant Symmetry-Curvature duality principle [11]. Kimia *et al.* [6] propose a continuum connecting these two extreme views, and capturing the distinction between the *parts* extreme and the *protrusions* extreme. This continuum is one of three separate continua that reflect the cooperation and competition between three processes, *parts*, *protrusions*, and *bends*, to describe shape, Figure 4.

The focus of this paper is on the parts node of the shape triangle. Specifically, we address the question "How can parts of 2D shapes be reliably obtained?" Three approaches are possible. First, the question may be addressed in light of constraints imposed by object recognition, such as occlusion, movement of parts, and stability under changes in viewing conditions and object deformation [13]. Second, the question may be addressed in the context of how the human visual system may have acquired a notion of parts. What properties of objects and their projections as manifested in the retinal image, if any, may be used to robustly determine parts, even in the absence of such visual cues as color, texture, shading *etc.*? This may be viewed as addressing an "ecological aspect" of the question. Third, the question may be posed in the context of whether there is any evidence for a perceptual notion of parts in the human visual system. In particular, does a human subject determine components of known and unknown shapes in a consistent manner, over different trials? If so, is there consistency in the manner in which different subjects determine components of the same shapes? Finally, if there is evidence for an intra-subject as well as inter-subject

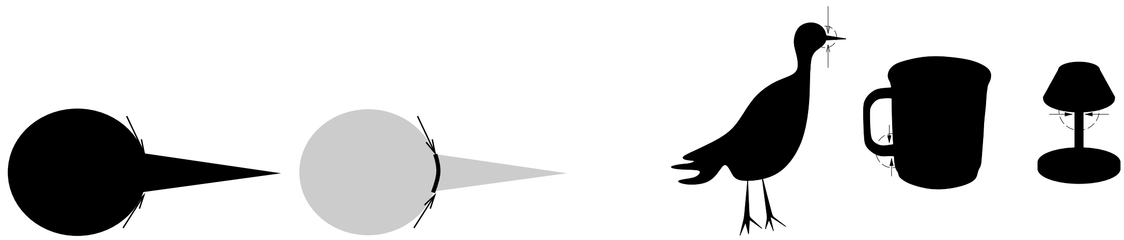

Figure 5: LEFT: Limb-based parts result from part-lines whose end-points are negative curvature minima, *and* whose boundaries boundaries continue smoothly from one end-point to the other. RIGHT: Examples of limb-based parts.

consistent partitioning mechanism, what relation does it bear to our computational proposal and algorithm [13]?

2 Form From Function

How might the human visual system have acquired a notion of parts? Hoffman and Richards have sought to explain a notion of parts based on regularities of nature. In their proposal, the defining regularity is found not in the shape of a part, *e.g.*, as defined by primitives, but in properties of part intersection: *transversality regularity* states that "when two arbitrarily shaped surfaces are made to interpenetrate they always meet in a contour of concave discontinuity of their tangent planes", *e.g.*, a straw in a glass of water. This provides a definition of parts in the 3D world. On the other hand, a different kind of regularity, *singularity regularity*, or "lawful properties of the singularities of the retinal projection", allows for the inference of parts from a 2D retinal projection, at negative curvature minima of the shape's contour.

Whereas we are in general agreement with Hoffman and Richards that a notion of parts should be motivated by a uniformity of nature, and not on the shapes of parts as defined by primitives, we suggest that a different type of regularity plays an important role. Observe that objects do not exist in isolation, rather, they must interact with their environment and with other objects in it. An object's two-dimensional form is influenced by both the nature of its *interaction* with other objects in a three-dimensional world, and the nature of its *projection* onto the retinal image. First, the interaction of objects has an influence on their form:

Principle 1 *Objects in their interaction with other objects in the environment specialize function differently across their volumes, either by evolution or by design.*

The *independent* specialization of function in two different, but connected regions is accompanied by sharp changes in a number of material and shape properties, *e.g.*, the beak of a bird is often hard, pointed, and distinct in color and texture from its head which is round and soft. Focusing on the shape properties alone, we can expect a sharp change in the surface of the object at the interface between the two connected regions, *e.g.*, the join between a bird's beak and its head, the shape of each region being *independently* suited to its particular function. The projection of these sharp

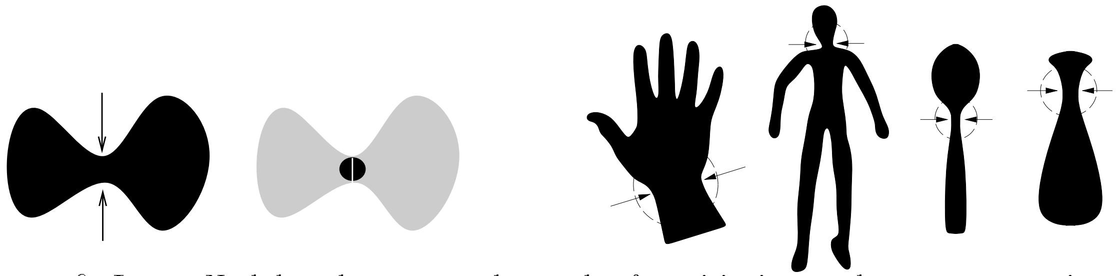

Figure 6: LEFT: Neck-based parts are the result of partitioning at the narrowest regions, namely, at part-lines whose lengths are a local minimum, *and* which are the diameter of an inscribed circle. These necks correspond to the second-order shocks of [7]. RIGHT: Examples of neck-based parts.

changes yields a pair of high curvature points. Moreover, when the constraint of the join is removed, in the vicinity of the interface the surface should have smooth continuations on (at least) one side. In the retinal image, the boundary at one high curvature point should smoothly continue to the boundary at the other one. This is a restatement of the Gestalt principle of "good continuation" for parts, and provides the first type of parts, the *limb-based parts*, Figure 5 (left). Examples of limb-based parts are shown in Figure 5 (right).

Second, Principle 1 is complemented by an assumption about the nature of *projection* of objects, which also leads to limb-based parts. Consider that when one portion of an object occludes another, *e.g.*, the cylindrical container of a mug occluding its handle, Figure 8, the projection of the occluding contour leads to a pair of concave discontinuities with evidence for good continuation:

Principle 2 *The projection of a portion of an object, partially occluding another, gives rise to a pair of concave discontinuities where the boundary at one discontinuity "smoothly" continues to the boundary at the other discontinuity.*

Thus, projection may also lead to limb-based parts, but where the pairs of negative curvature minima are in fact concave discontinuities. To formalize these ideas: *a* PART-LINE *is a curve whose end-points rest on the boundary of the shape, which is entirely embedded in it, and which divides it into two connected components; a* PARTITIONING SCHEME *is a mapping of a connected region in the image to a finite set of connected regions separated by part-lines.* In the case of limb-based parts, the strongest form of smooth continuation is co-circularity [12]. Formally:

Definition 1 (Limb) *A* LIMB *is a part-line going through a pair of negative curvature minima with co-circular boundary tangents on (at least) one side of the part-line, Figure 5 (left).*

While for limb-based parts a distinction in function between two connected regions leads to a distinction in their shapes, when such a distinction is also accompanied by a need for *articulation* between the two regions, the space necessary for movement leads to a narrowing of the connecting region. This resulting "neck" can be a determinant of parts, *e.g.*, a human neck or wrist, or a tail of a fish, Figure 6 (right). This leads to our second type of parts, the *neck-based parts*. Formally:

Definition 2 (Neck) *A* NECK *is a part-line which is also a local minimum of the diameter of an inscribed circle, Figure 6 (left).*

Articulation, however, is not the only cause of such narrowing. An emphasis on economy of mass in conjunction with the specialization of each region's shape to its particular function can also lead to a neck, *e.g.*, the neck of a spoon, Figure 6 (right). It is only natural that such narrowings be interpreted as the join between two parts, after all, a very narrow neck is the first place to break when an object falls or undergoes a collision. In addition, necks are often stable grasping sites. Of course, there are cases where a narrowing does not suggest functionally separate regions, *e.g.*, the neck of a vase, Figure 6 (right). However, it appears that the connection between distinction in function and necks is strong enough, and occurs frequently enough, that "partition at a neck" becomes the dominant rule; the neck of a vase is perceived as an equally good partitioning place as the neck of a person or the tail of a fish.

The above arguments in support of limb-based and neck-based parts together constitute the "**form-from-function**" principle which maintains that:

Principle 3 *Form is influenced by an* interaction *among objects and between objects and their environment.*

It is our contention that such interaction leads to the specialization of function across regions. This specialization of function results in distinctions in visual properties such as color, texture, *etc.*, and shape. Conversely, distinctions in these visual properties can lead to the recovery of the specialization in function, governed by general principles. Our visual system uses such principles to recover parts, as suggested by our limb-based/neck-based proposal. This proposal, motivated by a "form-from-function" regularity, differs in a number of properties from parts motivated by transversality, including the provision of roles for the regional interaction of boundary features, for "good continuation" and for occlusion junctions, Figures 7 and 8.

3 Experiments

We carried out a set of experiments to examine "psychophysical" aspects of how parts of 2D shapes may be obtained. Specifically, we addressed three questions. First, do human subjects determine components of 2D shapes consistently across different trials of the same partitioning task? Second, is there evidence for consistency between subjects for the same partitioning task? Third, if there is evidence for an intra-subject and inter-subject consistent partitioning mechanism, how do the perceived parts compare to our proposed limbs and necks? In these experiments, subjects were required to partition a variety of biological and nonsense 2D shapes into perceived components by drawing part-lines across them.

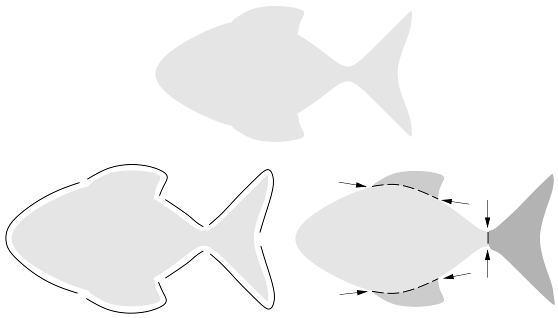

Figure 7: The minima rule decomposes the boundary of the fish (top) into seven pieces (bottom left), some of which appear to bound the same part. A decomposition of the shape into four regions, based on limbs and necks (bottom right), appears to be more intuitive.

Figure 8: Whereas transversality is a familiar part of everyday experience [3], in a 2D image the occlusion of one object by another object, or by another part of itself, is more likely to occur. Only the right view is one of a handle "interpenetrating" a cup.

Subjects. The 14 subjects (6 male, 8 female) were students and staff members at Brown University, or friends of the authors. Each subject participated voluntarily and had no prior knowledge of the purpose of the experiment, or of theories of visual perception of parts.

Shapes. 15 nonsense shapes, Figure 9, and 7 biological shapes, Figure 10, were used. The nonsense shapes and the *fish* and *leaf* shapes were hand drawn using the *idraw* program on a SUN workstation. The remaining shapes were obtained from a public domain database for Macintosh computers. Each image was black on white, with a resolution of 512×512. While the biological shapes represented familiar objects whose components had names and associated functions, the nonsense shapes represented the opposite extreme.

Procedure. The subject was seated in a dimly lit room in front of a SUN workstation with an 8 bit color monitor (16" or 19") with a resolution of 1152×900. The subject faced an interactive program that displayed the shapes in a random sequence in a display window located at the center of the screen. The subject was required to break off perceived components for each shape by drawing part-lines across them with the aid of a mouse. For every part-line drawn, the program stored its intersection points with the shape.

Experiment One: Intra-Subject Consistency. The goal of the first experiment was to measure a subject's consistency across different trials of the same partitioning task. Five subjects performed the same experiment twice, with an interval of roughly six months between the two trials, as described by the procedure above. The results are presented in Figure 11. It is evident that a subject's partitioning strategy does not vary much across time; 87% of the part-lines for the nonsense shapes and 72% of the part-lines for the biological shapes were drawn for both trials. Apparently a consistent mechanism might underlie the partitioning of biological as well as nonsense shapes. The next experiment examines the degree of similarity of this mechanism

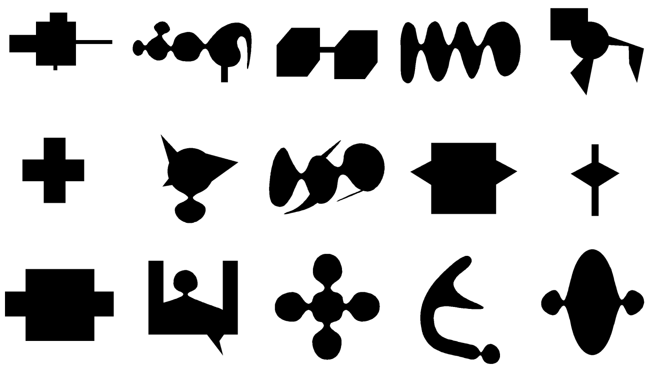

Figure 9: The 15 nonsense shapes.

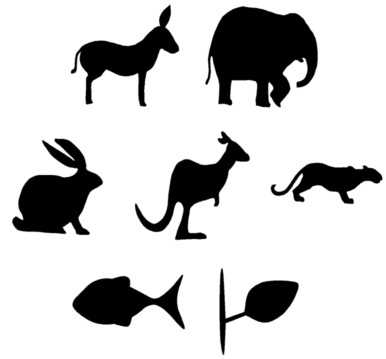

Figure 10: The 7 biological shapes.

across subjects.

Experiment Two: Inter-Subject Consistency. The goal of the second experiment was to measure consistency across subjects. The results from the first trial of the five subjects in experiment one were pooled with those from nine additional subjects who performed the experiment once, as described by the procedure above. The results are presented in Figure 12. There is a high degree of consistency for the nonsense shapes; 81% of all the perceived part-lines were agreed upon by at least 11 out of 14 subjects. Evidence for inter-subject consistency for the biological shapes is strong as well; 60% of all the perceived part-lines were agreed upon by at least 11 out of 14 subjects. It appears that the partitioning strategies used by different subjects are indeed similar. The results of this experiment, coupled with the high level of intra-subject consistency, shown in experiment one, provide evidence for a perceptual notion of parts in the human visual system. The next experiment compares parts perceived by our subjects with parts computed by our algorithm.

Experiment Three: A Comparison of Perceived and Computed Parts. The goal of the third experiment was to determine the degree of similarity between perceived and computed (or proposed) part-lines. We first examined the percentage of perceived part-lines that were also proposed as necks or limbs, Figure 13. It is apparent that a high percentage of perceived part-lines were in fact proposed; 90% for the nonsense shapes and 67% for the biological shapes. This provides strong evidence in favor of our two proposed parts types. We then examined the percentage of part-lines proposed as necks or limbs that were also perceived, Figure 14. For the nonsense shapes 58% of the proposed part-lines were perceived, for the biological shapes the percentage was 49%. It is evident that whereas a significant proportion of proposed part-lines were also perceived, an appreciable proportion were not; we address this in the discussion. Finally, we examined the position and orientation errors between proposed and perceived part-lines that coincided, Figures 15 and 16. 87% of the position errors for the nonsense shapes and 76% for the biological shapes were less than 5 pixels (1% of the image size); 96% of the orientation errors for the nonsense shapes and 83% for the biological shapes were less than 10 degrees. There is striking evidence that the part-lines were placed precisely at proposed necks and limbs, not merely in their vicinity. Figure 17 is a powerful illustration of the high degree of

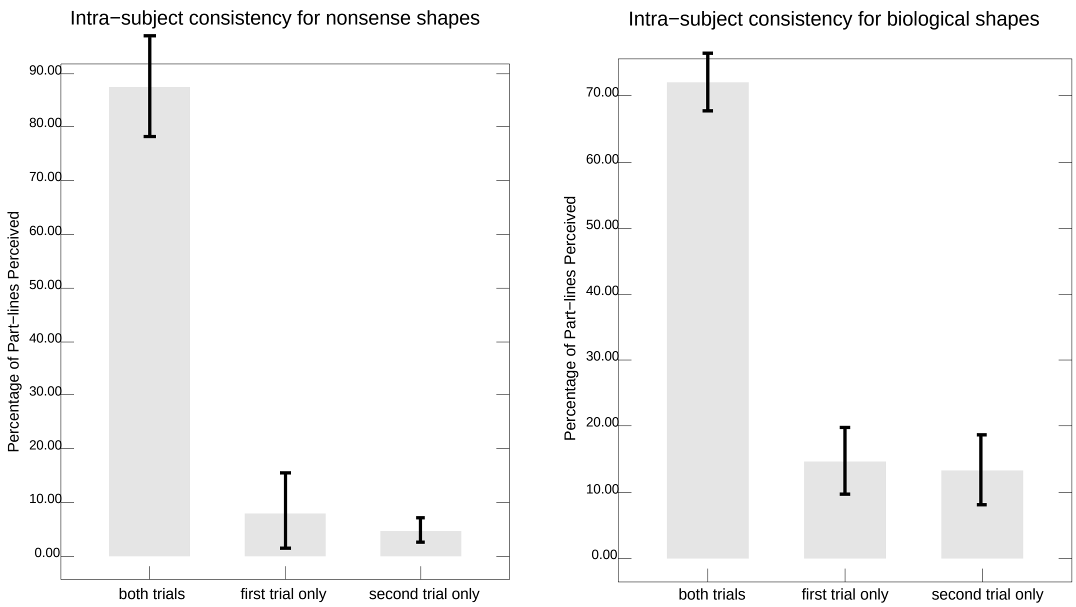

Figure 11: Intra-subject consistency: The percentages of all perceived part-lines that are drawn for both trials, or drawn for one trial only are plotted, averaging over five subjects, for the nonsense shapes (left) and the biological shapes (right).

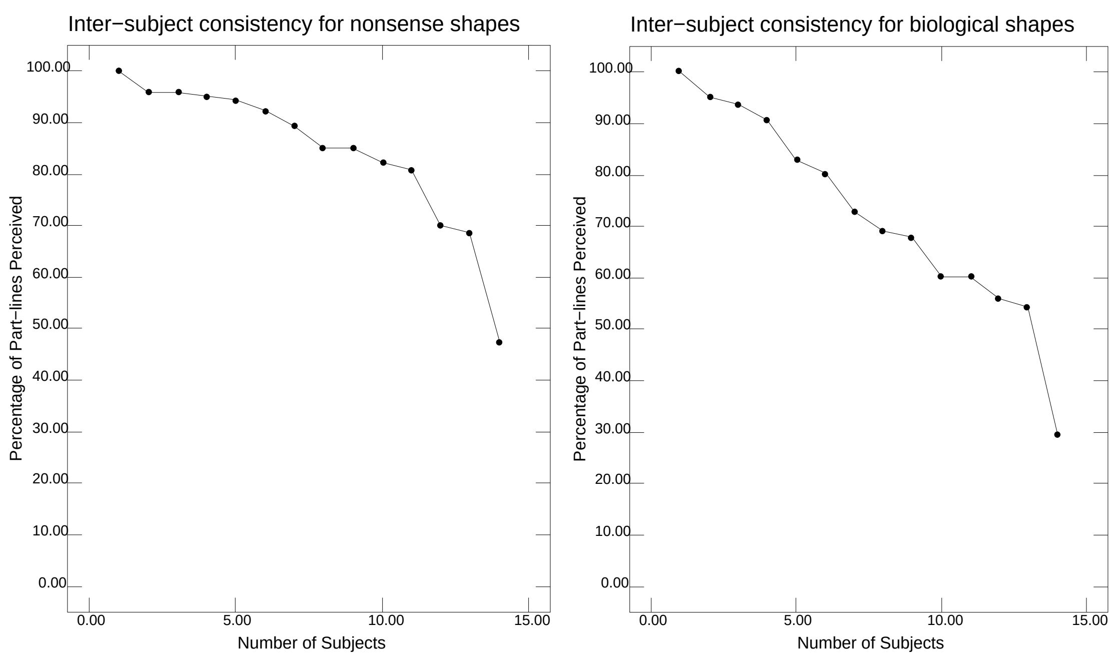

Figure 12: Inter-subject consistency: The part-lines drawn for a shape are divided into 14 (possibly overlapping) groups, the first containing those drawn by *at least* one subject, the second those drawn by at least two subjects, and so on. The number of part-lines in each group is stored as a percentage of the total number of part-lines drawn for the shape; the results are averaged over the nonsense shapes (left) and the biological shapes (right).

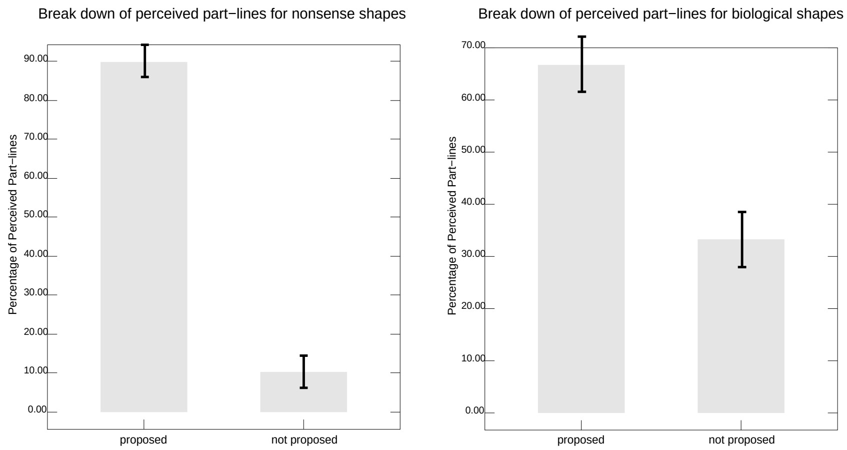

Figure 13: Proportions of perceived part-lines that were proposed, or were not proposed, averaged across all subjects for the nonsense shapes (left) and the biological shapes (right).

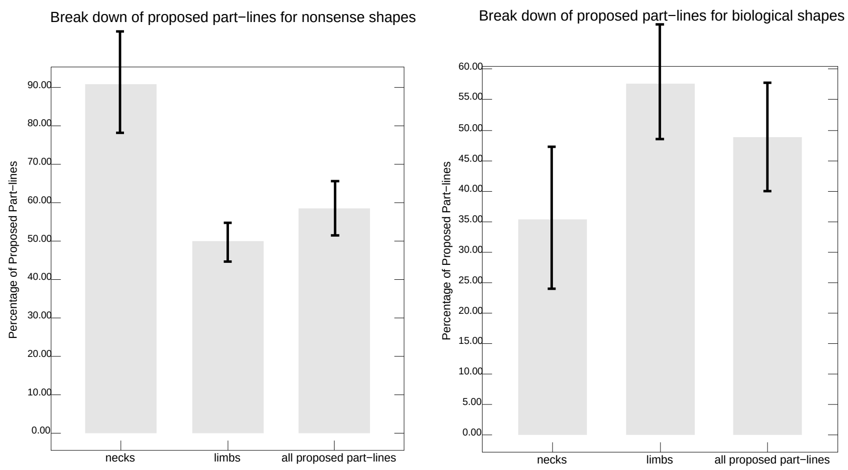

Figure 14: Proportions of proposed necks, proposed limbs and proposed part-lines (necks and limbs) that were perceived, averaged across all subjects for the nonsense shapes (left) and the biological shapes (right).

overlap between perceived and proposed part-lines. Each box depicts the original shape (left), the parts computed by applying our algorithm (middle), and the parts perceived by a majority of the 14 subjects (right).

Discussion We address three questions: (1) Why is the percentage of perceived part-lines that are also proposed lower for the biological shapes than for the nonsense shapes? First, cognitive knowledge may override perceptual evidence, *e.g.*, it is observed that some subjects break off components such as the feet of the animal shapes, even when perceptual evidence is weak or absent, leading to an "over-segmentation" of these in comparison to the nonsense shapes. Second, for some limbs perceptual evidence has been partially removed by the process of bending, leading to "trunks" or "tails", Figure 17 (E, F, G). In such situations, whereas one end-point is clearly defined, the placement of the second one is rather arbitrary. (2) Why is the percentage of proposed part-lines that are perceived lower than the percentage of perceived

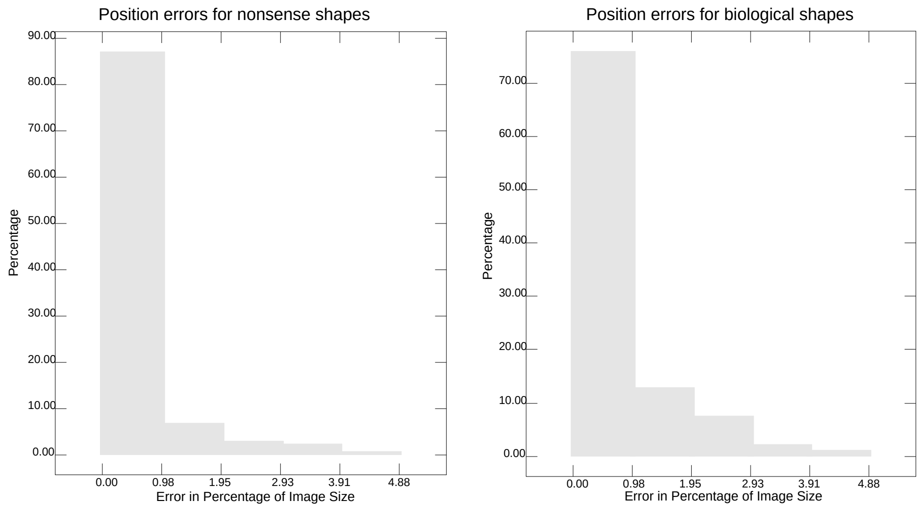

Figure 15: Position errors between perceived and proposed part-lines, for the nonsense shapes (left), and the biological shapes (right). The errors were extremely small, 87% of the position errors for the nonsense shapes and 76% of the position errors for the biological shapes were less than 5 pixels (1% of the image size).

part-lines that are proposed? Recall that no restriction was placed on the number of components to be broken off; each subject was required to use his/her best judgment. Whereas a large majority of the part-lines drawn do in fact correspond to proposed part-lines, subjects differ in their choice of which components to break off, in most cases *selecting only a subset* of all the necks and limbs. (3) What factors account for perceived part-lines that are not proposed? First, on occasion, subjects expressed their frustration at being limited to the use of a single straight line, their percept being that of a curve. Second, several subjects remarked that some shapes were more readily perceived in 3D, *e.g.*, as "two cubes connected by a bar", Figure 9 (top middle). Third, symmetry and regularity tend to influence a subject's behavior, *e.g.*, a subject who broke off one foot of the elephant often broke off the other two as well, despite the lack of perceptual evidence for parts. Lastly, the outliers, *e.g.*, part-lines perceived by only one subject, have not been excluded from the analysis.

4 Figure/Ground Segregation: The Role of Parts

In discussing the relationship between parts and figure/ground segregation, we are faced with somewhat of a chicken-and-egg problem. Whereas thus far we have assumed the availability of a 2D form, because intensity, color, and texture are usually not uniform across an object, segregating it from its background and from other objects in its vicinity is not an easy task. How can partitioning proceed when figure/ground separation has not taken place? We suggest a view of parts as an *intermediate representation* that allows for the flow of bottom-up as well as top-down information. Since part computations are local, edges of the appropriate polarity can interact to form necks and limbs *prior* to obtaining a separation of the object, *e.g.*, the arrows in Figure 18 (right) indicate a pairing of edges of the appropriate

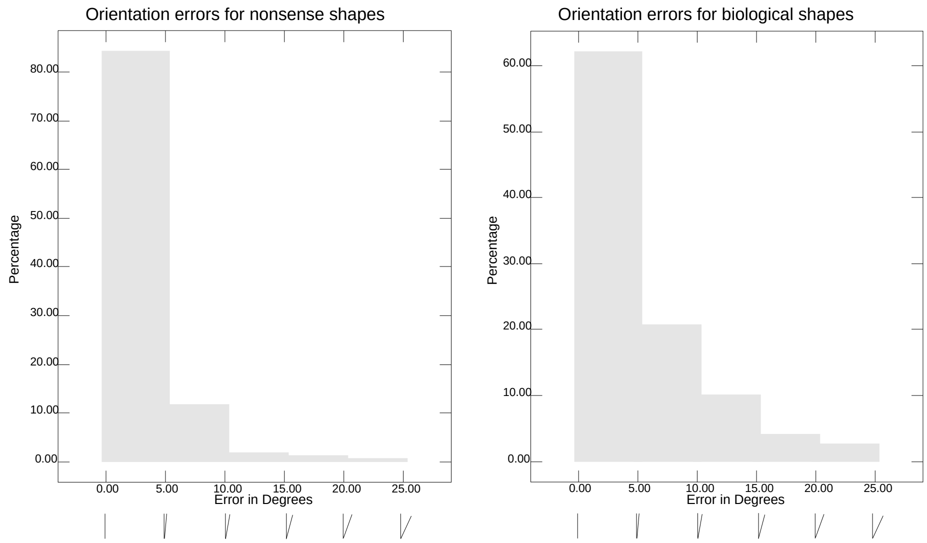

Figure 16: Orientation errors between perceived and proposed part-lines, for the nonsense shapes (left), and the biological shapes (right). The errors were small, 96% of the orientation errors for the nonsense shapes and 83% of the orientation errors for the biological shapes were less than 10 degrees.

polarity that provide evidence for the existence of a limb. This leads to the notion of a "PARTS RECEPTIVE FIELD", *i.e.*, a local operator that looks for pairs of edges of the appropriate polarity that provide evidence for the existence of limbs or necks, Figure 18 (left). This constitutes the bottom-up flow of information, *i.e.*, from local edge hypotheses to more global part hypotheses. Now part hypotheses can in turn play an integral role in the figure/ground segregation process through the top-down flow of information. A combination of likely parts can trigger an object hypothesis, followed by a figure/ground separation hypothesis for the image. During this stage, pairs of limbs that correspond to occluded contours ("hidden limbs" [13, 14]) can be used to group together portions of the scene which they connect, *e.g.*, two sides of an occluded branch, Figure 18 (right). Such a notion of parts may be key in resolving the bottom-up/top-down bottleneck of recognition.

5 Parts and Subjective Contours

There is an interesting, speculative connection between our partitioning proposal and the perception of certain subjective contours. Observe the Kanizsa illustrations in Figure 19 (top), each of which is perceived as a white object resting against the background. The edges, however, do not correspond to actual changes in intensity and hence are subjective or "filled in" by our visual system. One explanation lies in interpreting each illustration as the outcome of occlusion, where the occluding object happens to share the color of part of its background (white), although it appears to be brighter. It is then possible to *recover the subjective contours as limbs* arising from occlusion. To illustrate, the parts receptive fields depicted on the contrast reversed versions of the original illustrations, Figure 19 (middle), recover limbs which do indeed

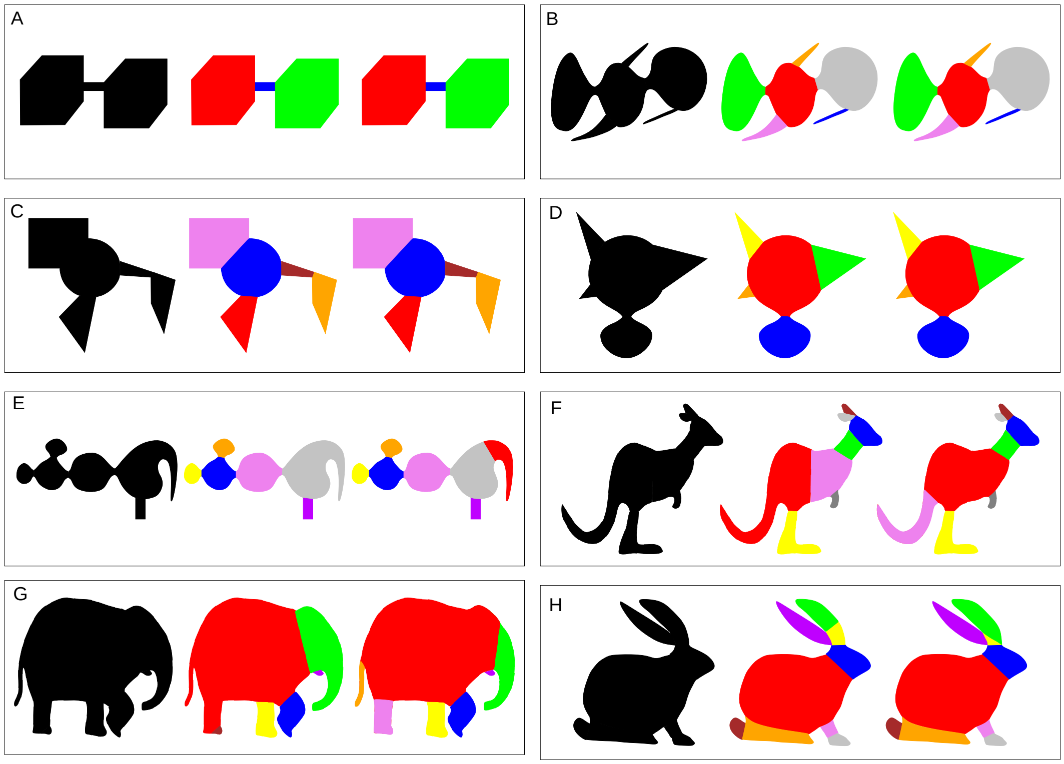

Figure 17: This Figure compares computed parts and perceived parts for a subset of the shapes used in our experiments. Each box depicts the original shape (left), the parts computed by applying our algorithm (middle), and the parts perceived by a majority of the 14 subjects (right). Note that for shapes (A) through (D), the computed and perceived parts are in exact agreement. Shapes (E), (F), and (G) illustrate discrepancies that occur due to the existence of bent limbs, *e.g.*, those manifested as the kangaroo's tail and the elephant's trunk.

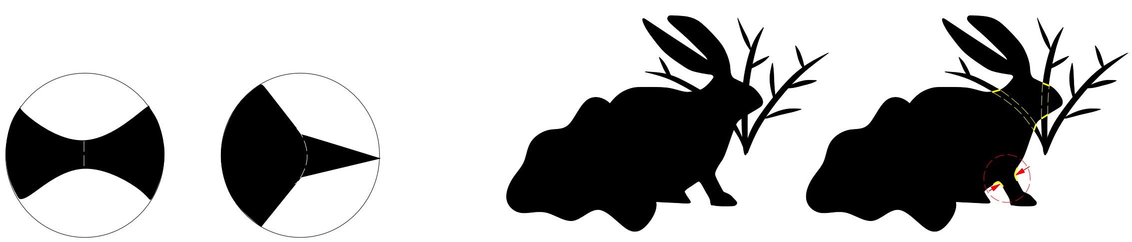

Figure 18: This Figure depicts the application of a *"parts receptive field"* (LEFT) to a natural scene (RIGHT). It is clear that the separation of the rabbit's shape (figure) from those of other objects (ground) is not an easy task, especially when no color, texture or additional visual information is available. Nevertheless, local edge hypotheses can interact to form more global part hypotheses *prior* to obtaining an actual figure/ground segregation. The arrows indicate a pairing of edges of the appropriate polarity that provide evidence for the existence of a limb. In addition, pairs of "hidden limbs" [13, 14] (dashed lines) signal single objects under occlusion; these limbs can be used to group the corresponding parts on either side.

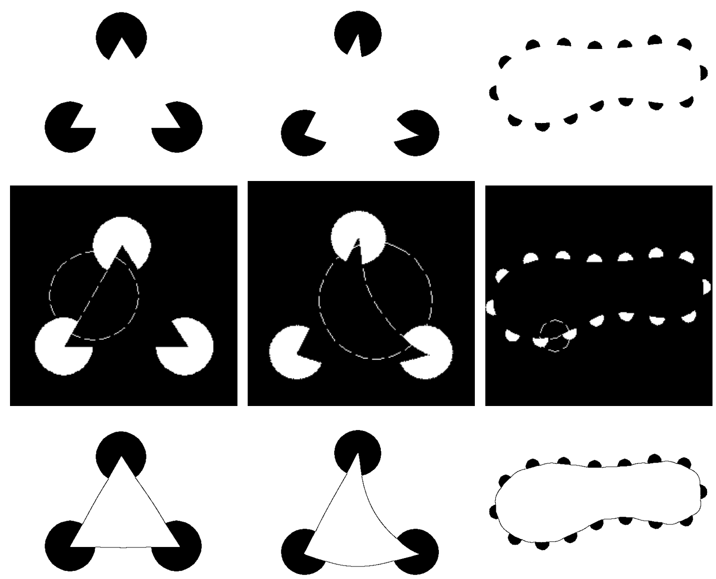

Figure 19: TOP: The Kanizsa illustrations are perceived as white objects resting on top of the background. MIDDLE: The limbs are computed by a parts receptive field. BOTTOM: Superimposition of the results on the original illustrations. The limbs coincide with the subjective contours!

correspond to the subjective contours, Figure 19 (bottom). Edges alone, however, are not sufficient to explain such illusions, Figure 20. We have recently learned of similar ideas [5].

Acknowledgements The authors thank David Cooper, Allan Dobbins, James Elder, Danny Keren, Frédéric Leymarie, Francene Reichel, Jayashree Subrahmonia, Bill Warren and Steve Zucker for discussions and comments. We thank our subjects for participating in the experiments. The animal shapes were kindly provided by Fridtjof Stein of the University of Southern California. This research was supported by NSF grant CDA-9305630.

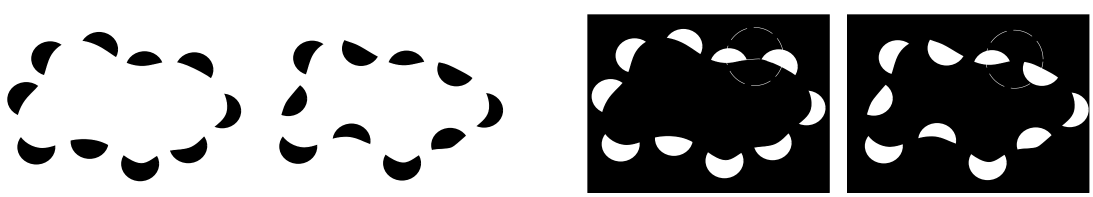

Figure 20: Edges alone cannot explain the illusory contours in this Kanizsa Figure [4]. LEFT: The figure on the right was created by reversing the polarity of each alternate edge of the figure on the left, while maintaining its location. The illusion is severely weakened. RIGHT: Computed limbs correspond to the illusory contours (LEFT). For the reversed polarity figure, whereas the boundary pieces continue to show "good continuation", the corresponding regions do not lie on the same side of the part-line, therefore no limbs result (RIGHT).

References

[1] I. Biederman. Recognition by components. *Psych. Review*, 94:115–147, 1987.

[2] M. L. Braunstein, D. D. Hoffman, and A. Saidpour. Parts of visual objects: an experimental test of the minima rule. *Perception*, 18:817–826, 1989.

[3] D. D. Hoffman and W. A. Richards. Parts of recognition. *Cognition*, 18:65–96, 1985.

[4] G. Kanizsa. *Organization in Vision: essays on Gestalt perception.* Praeger, 1979.

[5] P. J. Kellman and T. F. Shipley. A theory of visual interpolation in object perception. *Cognitive Psychology*, 23:141–221, 1991.

[6] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. The shape triangle: Parts, protrusions, and bends. Technical Report TR-92-15, McGill University Research Center for Intelligent Machines, 1992.

[7] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Shapes, shocks, and deformations, I: The components of shape and the reaction-diffusion space. *International Journal of Computer Vision*, To Appear, 1992.

[8] J. J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13:321–330, 1984.

[9] J. J. Koenderink and A. J. van Doorn. The shape of smooth objects and the way contours end. *Perception*, 11:129–137, 1982.

[10] M. Leyton. Inferring causal history from shape. *Cognitive Science*, 13:357–387, 1989.

[11] M. Leyton. *Symmetry, Causality, Mind.* MIT press, April 1992.

[12] P. Parent and S. W. Zucker. Trace inference, curvature consistency and curve detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(8):823–839, August 1989.

[13] K. Siddiqi and B. B. Kimia. Parts of visual form: Computational aspects. *IEEE Trans. Pattern Analysis and Machine Intelligence*, To Appear, 1994.

[14] K. Siddiqi, K. J. Tresness, and B. B. Kimia. Parts of visual form: Ecological and psychophysical aspects. Technical Report LEMS 104, LEMS, Brown University, June 1992.