Parts of Visual Form:
Computational Aspects

Kaleem Siddiqi

Division of Engineering
Brown University
Providence, RI 02912

Benjamin B. Kimia

Division of Engineering
Brown University
kimia@lems.brown.edu

Abstract

Underlying recognition is an organization of objects and their parts into classes and hierarchies. A representation of parts for recognition requires that parts be invariant to rigid transformations, robust in the presence of occlusions, stable with changes in viewing geometry, and be arranged in a hierarchy. These constraints are captured in a general framework using notions of a PART-LINE *and a* PARTITIONING SCHEME. *A proposed general principle of "form from function" motivates a particular partitioning scheme involving two types of parts,* NECK-BASED *and* LIMB-BASED, *whose psychophysical relevance was demonstrated in [35]. Neck-based parts arise from narrowings in shape, or the local minima in distance between two points on the boundary, while limb-based parts arise from a* pair *of negative curvature extrema which have "co-circular" tangents. In this paper, we present computational support for the limb-based and neck-based parts by showing that they are invariant, robust, stable and yield a hierarchy of parts. Examples illustrate that the resulting decompositions are robust in the presense of occlusion and noise for a range of man-made and natural objects, and lead to natural and intuitive parts which can be used for recognition.*

Introduction Underlying recognition is an organization of objects and their parts into classes and hierarchies. In computer vision, the notion of recognition based on "parts" has become increasingly popular. The task of forming high-level object-centered models from low-level image-based features requires intermediate representations and parts serve such a role for robust recognition. Perhaps the most compelling support for this idea is based on recognition in the presence of occlusion: while local features are sensitive to noise and other variations, global structures are susceptible to occlusion by other objects. However, the stable computation of a few parts can lead to recognition that is robust in the presence of occlusions. In addition, objects are often composed of moving, or growing parts: while the description of each part remains intact, the relationships between parts change.

How should parts of a shape be computed? Previous approaches have focused either on the decomposition of its interior region or on the segmentation of its boundary. Blum's idea of shape as a collection of overlapping disks leads to partitioning of the interior shape along branch points of the symmetric axis [6, 7]. Other region-based parts include Shapiro and Haralick's maximally convex parts [33], a description based of best combination of primitives such as generalized cylinders [5, 24, 25], and superquadrics [27, 3], or the *simplest* description in some language [19, 28]. In contrast, contour-based segmentations use boundary features such as high-curvature points [9] whose salience is demonstrated in [2, 4]. Koenderink [18] noted that three-dimensional shapes are perceived as the composition of elliptic regions with the hyperbolic patches as glue, as evidenced in the works of art, and advocated a decomposition along parabolic lines, or inflection points of the two-dimensional shape. Primitive-based boundary partitioning schemes approximate the boundary as the best combination of primitives such as polygons [29], or primitive curvature changes, *e.g.*, cranks, bends, bumps, and ends [1].

Figure 1: Hoffman-Richards's theory of curve partitioning segments a contour at negative curvature minima.

A significant departure from the traditional techniques is the boundary-based method proposed by Hoffman and Richards [11] who advocate that *part decomposition should precede part description.* They contrast primitive-based approaches with boundary-based methods and propose a theory of parts which relies not on the shape of parts, as captured by primitives, but rather on general principles underlying their formations, or "regularities of nature". The *transversality principle*, as an example of such a regularity, asserts that "when two arbitrarily shaped surfaces are made to interpenetrate they always meet in a contour of concave discontinuity of their tangent planes". Singularity regularity, or "lawful properties of the singularities of the retinal projection", is a second example. These two regularities lead to a partitioning rule for plane curves: "divide a plane curve into parts at negative minima of curvature." Figure 1 illustrates the successful application of this rule to two examples, a face and a hand, which produces intuitive results. This theory explains several figure-ground reversal illusions successfully; see [8] for further psychophysical evidence. This partitioning proposal leads to a representation of the shape boundary based on *codons*, pieces of the boundary bounded by negative curvature minima; the curvature maxima and

1063-6919/93 $03.00 © 1993 IEEE

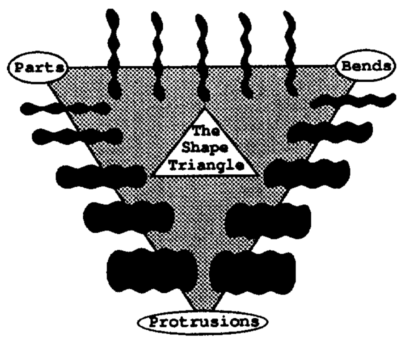

Figure 2: The *shape triangle* represents three cooperative/competitive processes acting on shape, namely, parts, protrusions, and bends. Each process is biased to see the shape in its own terms. The sides of the triangle represent continuums of shapes whose extremes correspond to the node and whose perception varies from one description to another. For a more detailed description see [16].

zeros are then used to classify each piece as one of the six possible types [31, 30, 32].

Leyton [20, 22, 21, 23] contrasts Hoffman and Richards' notion of parts with one based on a *process* in which a part is not a rigid segment. Rather, Leyton's process-analysis views a part as a causal explanation, or as a consequence of historical processes, which is based on a biologically-relevant Symmetry-Curvature duality principle [23]. To merge these extreme views, Kimia, Tannenbaum and Zucker [16] propose a continuum connecting two extremes, and capturing the distinction between, (1) the *parts* extreme where objects have clearly defined and distinct segments, *e.g.*, resulting from object composition, and (2) the *protrusions* extreme, where an object is best described as another whose boundary has been deformed, *e.g.*, due to growth. In [16], this continuum is viewed in the context of the *shape triangle* representing three cooperative/competitive processes acting on shape: *parts*, *protrusions*, and *bends*, Figure 2. The three sides of the triangle represent three separate continuums: the parts-protrusion continuum, the parts-bends continuum and the protrusions-bends continuum, capturing the tension between object composition, boundary deformation and region deformation. In this context, the proposal presented here may be viewed as a study focused on a single node of the shape triangle, namely the parts process.

Returning to the theory proposed by Hoffman and Richards, the distinction they present between the primitive-based and boundary-based methods can be viewed, alternatively, as a distinction between those primitive-based methods acting on the interior region and primitiveless, general-purpose boundary-based techniques. In other words, we propose that the boundary/region and primitive/primitiveless distinctions are orthogonal and not necessarily in conflict [17]. We are in general agreement with Hoffman and Richards on the limited capability of primitive-based techniques, whether they are applied to the boundary or the interior. However, we argue for a role not only for the boundary, but also for the interior of the shape [17]. For example, by stretching a small portion of the contour of a shape,

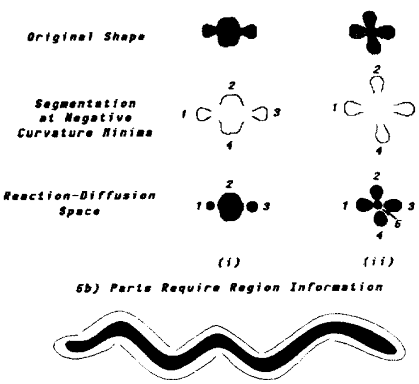

Figure 3: TOP: Partitioning of a two-dimensional shape requires not only boundary, but also region information.

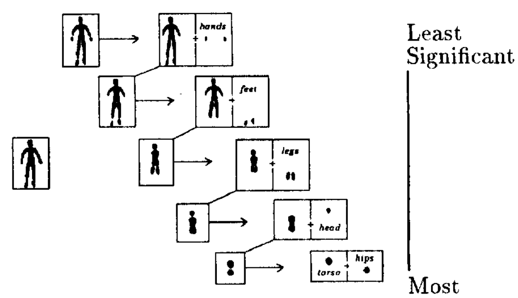

Figure 4: The hierarchy of parts for the doll obtained from the evolution of shocks in the entropy scale space [17].

while retaining it's remaining contours, one can produce new shapes whose parts change drastically [17], Figure 3 (top). As another example, the interaction of negative curvature minima through the interior region can produce "bending", rather than parts, Figure 3 (bottom). In addition, Kimia, Tannenbaum, and Zucker [17] discovered parts which are not derived from curvature extrema of the boundary, but were based on "necks" of the objects, Figures 4 and 7.

We now proceed to develop a formal framework to investigate the suitability of our partitioning scheme for recognition.

A Framework for Partitioning Shape In this section we develop a framework for studying partitioning schemes for decomposing visual form. Natural constraints from recognition demand that (1) a partitioning scheme be *invariant* to rigid transformations, be *robust* to local deformations and be *stable* with global transformations; (2) a partitioning scheme arrange parts in a hierarchy so that it is suitable for recognition of objects organized in a taxonomy, recognition in the presence of resolution effects, and for efficient recognition. We define notions of a *part-line* and a *partitioning scheme* and then formally develop the above constraints in this framework towards a general-purpose shape decomposition scheme.

A significant aspect of a partitioning scheme is the way decomposed parts behave under the various changes that may arise in the visual projection. Earlier approaches to partitioning have demanded that part computations be *reliable*, or invariant with time and viewing geometry [24, 11]. However, observe that three kinds of changes can occur in the visual world and parts should behave appropriately with these changes. First, we consider changes in the visual image formation that map to a combination of translations, rotations, and scalings of the two dimensional shape. It is clear that parts of the transformed shape must be exactly the transformed parts of the original shape. This constraint of *invariance* leads to computations that are *intrinsic* to the shape. The second class of changes are localized to certain portions of the shape, for example those due to partial occlusion, movement of parts, *etc.* It is natural to demand that such *local* deformations of the shape do not affect those parts that are remote to the change. This constraint of *robustness* leads to computations that are *local* to the part under consideration. Finally, changes in shape can be *global*, *e.g.*, due to changes in viewpoint, or viewing direction, growth, *etc.* When such changes to the two-dimensional shape are slight, we expect that the change in the part structure is also slight. This *stability* constraint leads to nonbinary part computations that are graded using a measure of strength.

A general-purpose decomposition of shape should be based on the interaction between two parts, rather than on their shapes. It is appropriate, then, to focus on the interface between parts:

Definition 1 *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.*

This definition is general in that it allows for schemes that favor boundary features as well as those that are based on regional properties. A collection of proposed part-lines which divide a shape into connected parts form a partitioning scheme:

Definition 2 *A* PARTITIONING SCHEME *is a mapping of a connected region in the image to a finite set of connected regions separated by part-lines.*

We are now in a position to formally develop the notions of invariance, robustness, and stability:

Definition 3 *A partitioning scheme is* INVARIANT *if the part-lines of a shape that is transformed by a combination of translations, rotations, and scalings, are transformed in exactly the same manner.*

For example, a translated shape's part-lines should translate by the same amount.

Second, we demand that variations in the scene and viewing geometry, in particular the occlusion of an object by another, cause no change to part-structures remote to such change. To illustrate this point further, observe that the "part-structure" of the visible portion of a doll, e.g. head, hands, torso, etc. remains invariant to the occlusion or movement of the feet, as in Figure 5. Let the size of a part-line be the distance between its two end points and define the LOCAL-NEIGHBORHOOD of a part-line as a disk centered on the part-line, whose radius is a constant proportion of the length of the part-line. A part-line is CONTAINED in some region if its local-neighborhood is contained in the region. Then,

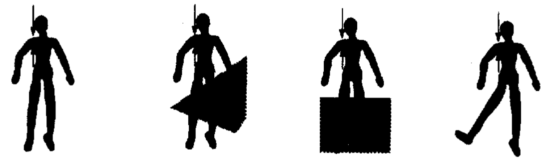

Figure 5: The part structure of the original image (left) should not be affected in areas remote to occlusion, movement of parts, *etc.*

Definition 4 *A partitioning scheme is* ROBUST *if for any two shapes, S_1 and S_2, which are exactly the same in some neighborhood, N, the part-lines $\{P_i\}$ contained in N for S_1 and S_2 are exactly equivalent.*

Informally, robust schemes must therefore determine part-lines "locally", namely, based on the information restricted to a local neighborhood of the part-line. However, since the part-line size may vary over the shape, this neighborhood also varies in size.

The third constraint relates global changes in the shape to the structure of its parts. We require that slight changes in the boundary of the object, which may arise from changes in viewpoint and viewing direction, deformation of the boundary, *etc.*, lead to slight changes in the parts and their relations. Informally, two similar objects must have similar parts. Formally,

Definition 5 *A partitioning scheme is* STABLE *if slight deformations of the boundary of a shape cause only slight changes in its part-lines.*

Specifically, slight deformations in the shape are measured by metrics for curves, *e.g.*, the Hausdorff metric, while changes in part-lines include changes in the part-line center, size, angle, its disappearance, or appearance.

In addition to the above three constraints that relate transformations of the shape to transformations of its parts, two factors constrain parts to be organized in a hierarchy. First, for recognition to be stable with changes in viewing distance and the accompanying changes in resolution, the parts obtained at one resolution ought to have some correspondence with those at another, leading to a *coarse-to-fine* representation of parts. Second, it is also well-known that efficient matching and storage requires a coarse-to-fine representation, which when applied to parts, suggests a hierarchical representation of parts.

A Partitioning Proposal Objects in the three-dimensional world project to two-dimensional entities, and in the process, much information about their part structure is lost. Our task is to define a two-dimensional notion of parts, based on the clues remaining in this lower-dimensional space, that leads to the recovery of three-dimensional parts. The recovery of structure, however, is not only a function of the information in the retinal image, but also depends on the assumptions one makes about the properties of the visual world [10, 13]. To recover parts, we rely mainly on two properties: one concerns the nature of objects and the other concerns the nature of visual projection. These assumptions lead to a notion of parts classified into *limb-based* and *neck-based* parts [35]. While a summary of the motivation for this definition is given below, our chief objective in this paper is to establish computational support for it

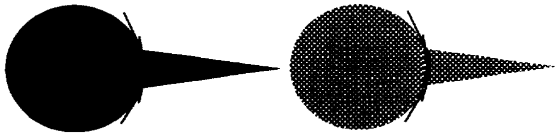

Figure 6: Limb-based parts are the result of partitioning at part-lines whose end points are negative curvature minima, *and* whose end-point's tangents continue smoothly from one to one another (co-circular tangents [26].)

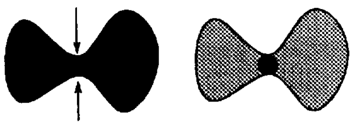

Figure 7: Neck-based parts are the result of partitioning at the narrowest regions, namely, at part-lines whose length is a local minimum, *and* which are the diameter of an inscribed circle. These necks correspond to the second-order shocks of [17].

as is related to the recognition of objects in a realistic situation involving partial occlusion, movement of parts, *etc.*

An object's form, as manifested in a two-dimensional image, is a function of both the *interaction* among objects in the three-dimensional world and the *projection* of objects onto the retinal image. First, as a result of *interaction* in an environment of objects, biological and man-made objects whether by evolution or by design, give rise to parts which specialize in their *function.* When objects specialize functions independently in two different, but connected regions, this results in sharp changes in the three dimensional surface of the objects, *e.g.*, the join between the beak of a bird and its head. The projection of these sharp changes yields a pair of high curvature points, where the tangent at one point smoothly continues to the tangent at the other point. 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 6. In contrast, two moving symbiotic regions require space for articulation, leading to a narrowing of the join between parts, *e.g.*, the tail of a fish; this provides the second type of parts, the *neck-based* parts. The two arguments together constitute the "form-from-function principle" which maintains that form results from the *interaction* among objects. Second, this principle is complemented by an assumption about the nature of *projection* of objects: a part of an object may partially occlude another, giving rise to a pair of negative curvature discontinuities where the tangent at one discontinuity "smoothly" continues to the other discontinuity, Figure 8. Thus, the projection of parts also leads to limb-based parts, but now with concave discontinuities in contrast to negative curvature extrema. Formally,

Definition 6 (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 6.*

Definition 7 (Neck) *A* NECK *is a part-line which is also a local minimum of the diameter of an inscribed circle, Figure 7.*

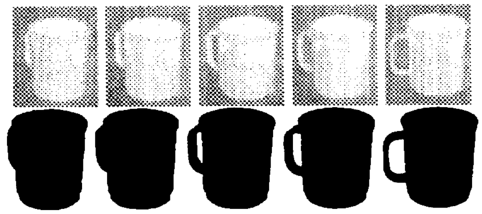

Figure 8: The projection of a mug leads to limb-based parts.

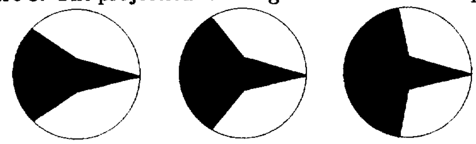

Figure 9: Salience of a limb increases as total curvature associated with the limb tangents on the supporting side of the limb decreases.

For a more detailed motivation of this definition and an extensive psychophysical validation of it see [35]. Intuitively, note that neck-based parts utilize *local region* information by minimizing the diameter of an inscribed circle. In contrast, limb-based parts utilize both *local boundary* information by minimizing negative curvature of their part-line's end points and *region* information by maximizing a measure of co-circularity, or continuation, of boundary tangents.

We now turn to properties of this partitioning scheme in relation to invariance, robustness, and stability. First, since the computation of necks and limbs, based on curvature extrema, co-circularity, and minimization of distance functions, is invariant to rotations, translations, and scaling, the resulting part-lines are *invariant* to such transformations. Second, observe that necks and limbs are computed based on information that is constrained to the *local-neighborhood* of the part-line. As such, variations in "remote" portions of the shape do not affect the resulting part-lines. In contrast, for partitioning schemes with a priori defined primitives, each part-line is the result of an optimization to describe shape as the best fitting combination of primitives. As such, globally-optimized primitive-based parts are a result of a global computation and not robust in the above sense.

Remark 1
The "Limbs-and-Necks" partitioning scheme is invariant.

Remark 2 *The "Limbs-and-Necks" partitioning scheme is robust.*

Next, the issue of stability for our scheme is more subtle and requires a further development. Note that so far the notions of necks and limbs are binary in that a part-line is either a neck or not; is either a limb or not. However, consider the behavior of parts on the

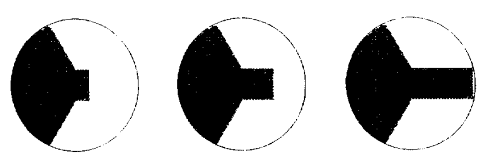

Figure 10: Salience of a limb increases as extent increases.

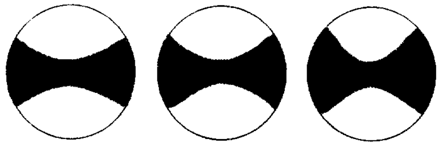

Figure 11: Salience of a neck increases as curvature disparity increases.

continuums of deformations represented by the sides of the shape triangle, Figure 2. Observe how the perception of the shapes depicted on the *part-protrusion* axis changes from a "sausage with parts" to a "lasagna noodle with no parts" [16]. A rigid binary inference of parts would lead to instability in that the part-based representation of the shape would change abruptly with a slight change at some point on this axis. A similar argument holds for the *part-bend* axis. Recall that our goal is the recovery of three-dimensional parts through intermediate two-dimensional parts. Since evidence for three-dimensional parts is not conclusive and is indeed often degraded, it is no surprise that a graded partitioning scheme, which necessarily associates a measure of **salience** with each part-line, is required.

Measures of salience for limbs and necks were psychophysically motivated. For limbs, two factors affect salience: *total curvature* and *extent*. Total curvature is the actual amount a tangent at one curvature extrema has to bend to align with the tangent at the next extrema. This measure describes how likely it is for the contour to continue to the contour across the limb, Figure 9. Extent is a measure of how likely it is that each side is an independent part, as indicated by the mass across the part-line, Figure 10. Similarly, the *curvature disparity* across the part-line indicates the salience of a neck, Figure 11, a measure of how likely it is to be a non-accidental visual structure. It is important to observe that salience ought to be independent of the absolute size of the part-line, since a magnification of the shape should leave part-structure invariant. Now, with the above measures of salience, the limb-based and neck-based part computations become stable: small changes in the shape induce small changes in the salience of parts. This measure of salience allows a smooth and stable traversal of the sides of the shape triangle.

Remark 3 *The "Limbs-and-Necks" partitioning scheme is stable.*

Having addressed the issue of how parts behave under various deformations, we move on to study the behaviour of part-structure under changes in resolution, for efficiency, and for suitability for a taxonomical representation of objects. First, it is important that the part-structure remain stable with changes in resolution. Note that when the object is viewed from increasingly larger distances, its visual details are increasingly more ambiguous. As the viewer moves closer to or away from the object, the part-structure must remain stable: detail may be added or removed to the parts-based representation, but it should not change drastically. What is needed is a notion of **scale** for parts that is compatible with changes in viewing distance. It is important to note that this notion of scale is independent of the notion of salience; all parts of Figure 12 are equally salient, but they occur at different scales. The hierarchy of parts

Figure 12: Parts of equal salience may occur at different scales.

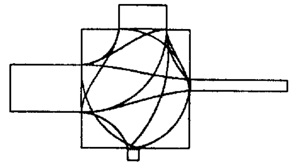

Figure 13: This figure illustrates how local part computations leads to global conflicts between part-line hypotheses.

in scale allows for efficient matching, and a taxonomical representation of objects.

The notion of scale for parts derives from a notion of smoothing for shape. As shapes are smoothed, small details, including parts, are removed while significant features remain, resulting in a hierarchy for parts. A smoothing space for shape, the *entropy scale-space*, was introduced in [15], Figure 18. Two processes underlie the smoothing in this space: *reaction* rigidly breaks off parts while *diffusion* melts the boundary into a smoother shape. There is an intimate connection between our partitioning scheme and the reaction-diffusion process of the entropy scale space. Neck-based part-lines are in one-to-one correspondence with negative reaction *second order shocks* formed by the reaction-diffusion process, while limb-based part-lines are the result of rarefaction waves corresponding to a pair of co-circular positive reaction *first order shocks*. What is important here is that parts are sorted into a hierarchy of scale using this smoothing process: reaction removes the information locally, while diffusion operates globally [15].

Thus far, we have discussed a notion of parts based on parts-lines with two independent measures of salience and scale. Since the computations have been local to each part, part hypotheses can globally be in conflict, *e.g.*, two crossing part-lines, or two part hypothesis using the same tangent. The conflict may be resolved by considering a **dynamic partitioning** selection procedure where the strongest part-line is selected at the coarsest scale and the corresponding part is then "broken off." The part that is broken off removes all conflicting part-lines and the procedure is then repeated by considering the next most salient part. The process is iterated for all scales until all part-lines are accounted for as either being selected, or discarded. To illustrate this, consider the Kanizsa Figure 14 which is perceived as two shapes occluding one another so that neither is entirely above the other. Our partitioning scheme finds eight candidate limbs of about equal salience, but which differ in scale. The smaller part-lines (solid), however, disappear due to the smoothing process, and at the same time remove evidence for the existence of the larger part-lines (dashed). Note that the discarded part-lines revive if the competing part-lines are removed by partial occlusion, Figure 14 right. In addition, while the accepted limbs are the occluding contours, the discarded ones are the occluded contours. As a result, a pair of these *hidden limbs* whose end-points are common with visible limbs

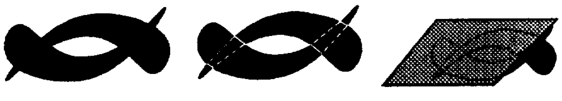

Figure 14: This Kanizsa figure is seen as two shapes occluding one another. The solid lines are the selected part-lines, while the dashed-lines indicate a grouping of two parts based on part-lines not selected. Note that the hidden limbs can become selected part-lines if the object is occluded (right).

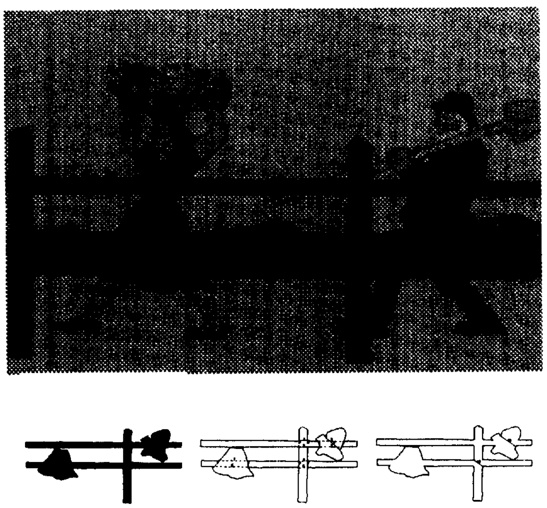

Figure 15: This figure taken from [14] is an example of how a local partitioning scheme may be in conflict with more global information. It is clear from the position of the feet and the occlusion of various body parts that the man and the woman are behind the fence. However, our partitioning scheme applied to the areas formed by the the man's jacket and the woman's skirt, which are merged with the fence, yields short vertical part-lines perpendicular to the fence. The discarded part-lines are long horizontal lines parallel to the fence and are considered hidden limbs. Together, they imply that the man's jacket, and the woman's skirt, occlude the fence, and hence must be in front of it!

signal a single object in occlusion, Figures 14, and 16. It is intriguing to note that our partitioning scheme is consistent with our visual perception of parts in Figure 15, where local evidence for parts overrides our cognitive knowledge of the visual scene [14].

Examples Figure 16 demonstrates the process of computing parts for the shape of an F16 fighter jet. Specifically, note how salience prunes the conflicting limbs and necks, leading to a natural intuitive decomposition and a recovery of the hidden parts, *e.g.*, the rockets. Figure 17 illustrates the effectiveness of the process, in particular its robustness in the presence of occlusions for the plane and tank shapes. Figure 18 illustrates that the computation of parts across scale for four similar pears [30] is similar at the coarse levels. Figure 19 illustrates the computation of parts for a tank in the presence of considerable occlusion and clutter. Details of implementation can be found in [34].

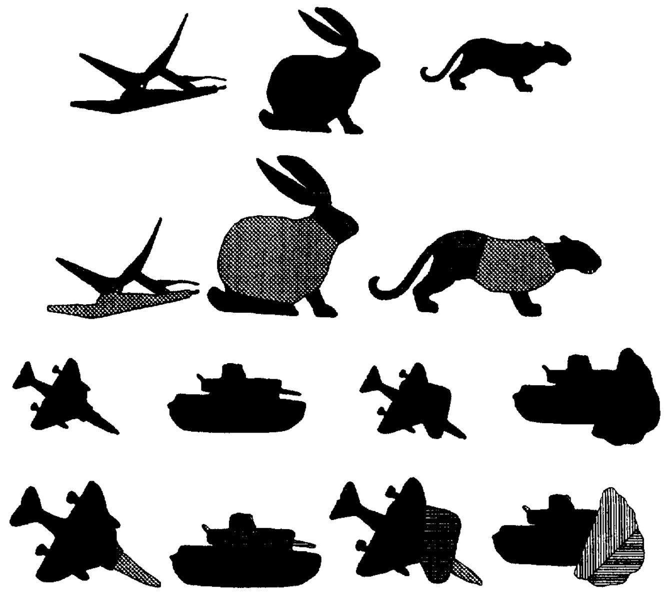

Figure 17: Examples of our partitioning scheme applied to a variety of shapes. While partial occlusion affects the occluded parts, it does not interfere with the recovery of the visible ones.

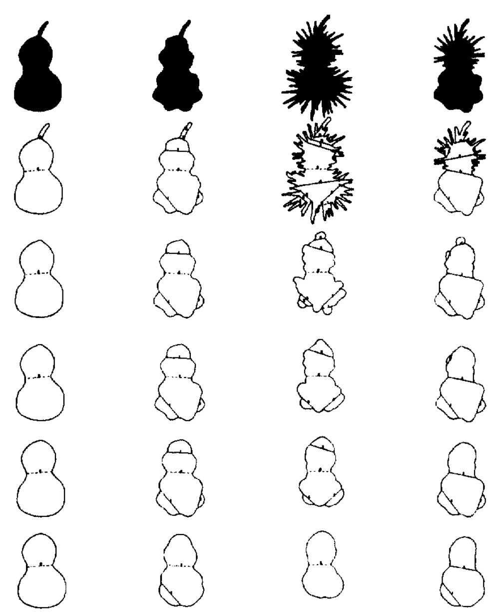

Figure 18: Bottom to Top: Necks and Limbs computed from coarse to fine scale for three different pears. The top row consists of the original images.

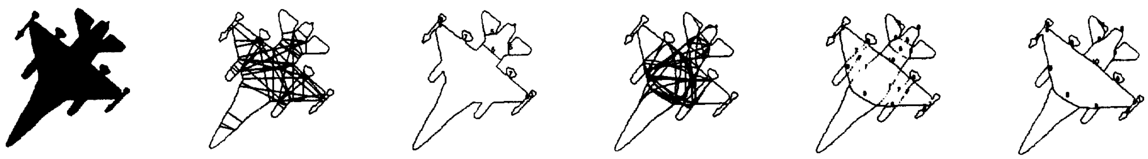

Figure 16: Necks and Limbs for a F16, from left to right: the original image; all candidate necks; selected necks with their strengths; all candidate limbs; selected limbs with their strengths, hidden limbs are shown as dashed lines; merged necks (light lines) and limbs (dark lines).

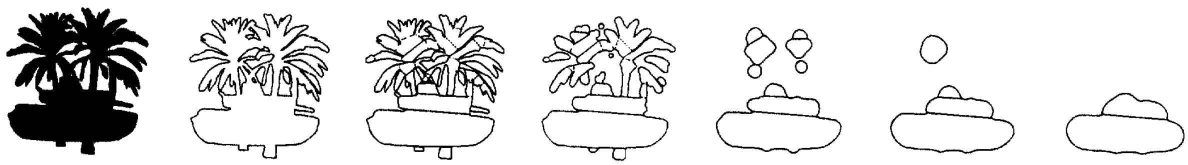

Figure 19: This figure illustrates that partitioning across scale brings out the coarse level structure of a tank which is blended-in with considerable occlusion and clutter.

Acknowledgements The authors thank David Cooper, James Elder, Jayashree Subrahmonia, and Steve Zucker for discussions and comments. We are thankful to Fridtjof Stein of University of Southern California for providing us with some of the images used here.

References

[1] H. Asada and M. Brady. The curvature primal sketch. *IEEE PAMI*, 8:2–14, 1983.

[2] F. Attneave. Some informational aspects of visual perception. *Psych. Review*, 61:183–193, 1954.

[3] R. Bajcsy and F. Solina. Three-dimensional object representation revisited. In ICCV1987 [12], pages 231–240.

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

[5] T. O. Binford. Visual perception by computer. In *IEEE Conference on Systems and Control*, December 1971.

[6] H. Blum. Biological shape and visual science. *J. Theor. Biol.*, 38:205–287, 1973.

[7] H. Blum and R. N. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10:167–180, 1978.

[8] 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.

[9] H. Freeman. Shape description via the use of critical points. *Pattern Recognition*, 10:159–166, 1978.

[10] J. Gibson. *The Perception of the Visual World.* Riverside Press, 1950.

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

[12] *First International Conference on Computer Vision, (London, England, June 8-11, 1987)*, Washington, DC., 1987. IEEE Computer Society Press.

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

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

[15] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Entropy scale-space. In C. Arcelli, editor, *Visual Form: Analysis and Recognition*, pages 333–344, New York, May 1991. Plenum Press.

[16] 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.

[17] 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*, Submitted, 1992.

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

[19] Y. G. Leclerc. Constructing simple stable descriptions for image partitioning. *International Journal of Computer Vision*, 3:73–102, 1989.

[20] M. Leyton. Symmetry-curvature duality. *Computer Vision, Graphics, and Image processing*, 38:327–341, 1987.

[21] M. Leyton. A process grammar for shape. *AI*, 34:213–247, 1988.

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

[23] M. Leyton. *Symmetry, causality, Mind.* MIT press, April 1992.

[24] D. Marr and K. H. Nishihara. Representation and recognition of the spatial organization of three dimensional structure. *Proceedings of the Royal Society of London*, B200:269–294, 1978.

[25] R. Nevatia and T. O. Binford. Description and recognition of curved objects. *Artificial Intelligence*, 8:77–98, 1977.

[26] 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.

[27] A. Pentland. Recognition by parts. In ICCV1987 [12].

[28] A. Pentland. Part segmentation for object recognition. *Neural Computation*, 1:82–91, 1989.

[29] U. Ramer. An iterative procedure for the polygonal approximation of plane curves. *Computer Vision, Graphics and Image Processing*, 1(3):244–256, 1972.

[30] W. Richards, B. Dawson, and D. Whittington. Encoding contour shape by curvature extrema. *JOSA A*, 3(9):1483–1489, 1986.

[31] W. Richards and D. D. Hoffman. Codon constraints on closed 2d shapes. *CVGIP*, 31(2):156–177, 1985.

[32] W. Richards, J. Koenderink, and D. Hoffman. Inferring three dimensional shapes from two-dimensional silhouettes. *JOSA A*, 4(7):1168–1175, 1987.

[33] L. Shapiro and R. Haralick. Decomposition of two-dimensional shapes by graph-theoretic clustering. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 1:10–20, 1979.

[34] K. Siddiqi and B. B. Kimia. Parts of visual form: Computational aspects. Technical Report LEMS 114, LEMS, Brown University, November 1992.

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