# Region-Based Grouping Operations for Locating and Describing Objects

Michael F. Kelly and Martin D. Levine

Center for Intelligent Machines & Dept. of Electrical Engineering
McGill University, Montréal, Québec, Canada, H3A 2A7

**E-mail: kelly *or* levine@cim.mcgill.edu**

## Abstract

A machine vision system has been developed based upon the need to extract *object* information from images. It is assumed that contrast edges capture all relevant object information. Principles which dictate how edge features may be grouped to infer objects are based upon detecting *symmetrical enclosing* edge configurations. Such configurations are detected using annular operators applied at multiple scales to edge data which have been extracted at multiple scales from a gray-scale image. The subsequent grouping of symmetry points results in a set of *parts* which make it possible to identify the *location* of objects within an image, and are used as a basis for constructing coarse *descriptors* for the objects found in the scene. Preliminary results are presented to illustrate the approach.

## 1 Introduction

Research in computational vision has often been directed towards the emulation of human visual capabilities. It can be argued that for humans, as well as for other biological organisms relying on sight, the visual sense is primarily used to extract information pertaining to *objects* in a given environment. In simple organisms this information may be used to indicate the presence or absence of an object, while in more complex visual systems, it provides a means for differentiating between a wide variety of objects. In this paper, we present research that is based on the assumption that the design of a vision system should be motivated primarily by the need to extract information about objects. This assumption provides the necessary constraints that may be used to guide the development of a a practical machine vision system capable of locating and describing objects in a scene. The resulting system provides significant benefits over established computational vision approaches which have, to date, been incapable of extracting such information directly from real images.

The proposed computational method is motivated by the fundamental need to locate objects within a scene. But what do we mean by the term object? Clearly there is an important distinction to be made between a physical object and the set of associated image data that results from visually sensing a physical object. If we consider a system where vision provides the only source of information about the world, then a physical object is interpreted and represented strictly based on this visual information. Since the only knowledge the system has is provided through visual means, all visual information pertaining to a given physical object is *equivalent* to that object as far as the system is concerned. In this paper the term object will be used to refer to that set of image data which provides information directly relating to physical objects in the scene.

To identify objects in a gray-scale image, we make several assumptions about the sources of object information and about the way that this information is distributed within an image. First, it will be assumed that all of the relevant object information is captured by the *contrast edges* (or simply edges) found in the image. Hence edges will be the only image features used to infer information about objects in the scene. In this paper we will be concerned with various *grouping* processes which allow object information to be inferred from a set of edge features. This grouping will require that certain assumptions be made about object structure so that it is possible to generally characterize how the edge features associated with an object are configured. Our second assumption is that all objects can be described in terms of two different types of contour geometries. We refer to these as *parts*. A given part may be characterized by a particular set of symmetric relationships which can be found using operators which detect *symmetric enclosure* within an image edge map. The first relationship is characterized by pairs of contours associated with object regions that are elongated. These are termed *limbs.* A second type of relationship is associated with regions which are less directed, and more globular. These regions are termed *blobs.* The distinction made between these two object primitives is related to the work of others [20], where objects are described in terms of three volumetric primitives - termed *sticks, plates,* and *blobs.* We note that there is a direct connection between these three-dimensional (3D) primitives and our two-dimensional (2D) ones, since generally 3D blobs project to 2D blobs and 3D sticks project to 2D limbs. As in [20] we make the assumption that a 2D decomposition in terms of blobs and limbs will be adequate for representing a wide range of common objects.

## 1.1 Motivation

To successfully detect objects based on a set of image features requires knowledge of what spatial edge patterns most likely correspond to an object. The process whereby a set of edge features is interpreted to infer the presence of objects may be described as *grouping,* whereby all edges associated with a given object are assigned to the same group. Studies examining perceptual grouping in humans have been made in the past, with perhaps the most famous examples coming from the Gestalt movement which proposed several principles that could be used to explain observed grouping phenomena [12]. One of the most general grouping principles depends upon relationships that exist between object features, particularly those relationships exhibiting a high degree of similarity or symmetry. More recent research has also suggested the importance of symmetry in determining the gaze fixation point for humans [7]. Since symmetry appears to be an important factor in the human perceptual grouping process, it seems natural to employ symmetric grouping principles when attempting to detect objects in a set of image features.

There are other reasons for choosing to group based upon symmetrical relationships. It is clear that for any object - both physical or within an image - there are always at least some symmetric structural relationships present. This fact has been exploited in past attempts to describe the two-dimensional shape of natural objects [2] where rotational symmetries between edge elements are represented in terms of a *symmetry curve.* Symmetry has also been used when developing tools for computational vision [6][18][11]. In all of these previous implementations it has been necessary to first identify or 'segment' objects in the image prior to symmetry analysis, and the analysis could only be performed on closed bounding contours. A further limitation relates to the sensitivity these previous methods have to small scale noise-like contour variations that usually are not relevant at higher scales or to the symbolic description of the object. We propose that grouping based on symmetric relationships may be used to segment objects directly from a set of edge features while at the same time providing the information necessary to construct useful representations for the segmented objects.

Edges that exhibit rotational symmetry about a point also appear to have a particularly high perceptual significance. For visual stimuli composed exclusively of edges, evidence suggests that humans tend to base their perception of inside and outside (or figure and ground) upon the degree to which a region is *enclosed* or *surrounded* by a set of edges [19][17]. As in the case of closed object contours, we will assume that an enclosing configuration can always be represented in terms of symmetrical relationships between various edge features. This sort of representation consists of a set of *symmetrically enclosed* points [9] that bears some resemblance to the set of points comprising the Medial Axis[2]. By using the principle of symmetry, it becomes possible to group discontinuous sets of edge features in order to infer the presence of enclosing structures.

In this paper, enclosing contour structures are detected using operators that cover a circular annular region in the image. Based on where in the annular band edges happen to fall, it is possible to empirically quantify the degree of enclosure and to classify it as either a blob or limb [9]. Enclosure information is represented by a set of spatially located *symmetry points.* Each point corresponds to the center of an annular operator at which the type and quantity of detected enclosure are recorded. To our knowledge, this application of annular operators has not been previously explored in computer vision[1].

## 1.2 Review of Related Methods

There are several examples of past research which are indirectly related to our work. Perhaps the closest is that of [22] in which a symmetry operator is used as a tool for finding 'interest points' within an image. This method requires no previous ob-

[1] Annular operators have been used in the past for edge detection [13], based upon comparing image data falling within a central circular region to that falling within a surrounding annular region. The operator which we propose here differs from [13] in that it uses only information which falls within the annular region.

ject segmentation. However, it does not lend itself to building shape representations since there is no principled means for choosing an appropriate range of scales without having some prior knowledge of the sizes of the projected objects in the image. Furthermore, with this method it is not possible to localize edge structures at a particular scale. Another related method relies on locating symmetrical structures while permitting objects to be located and coarsely described [16]. However, the authors do not attempt to deal with the issue of scale.

Other methods which attempt to explicitly use scale information fall under the category of 'scale space' [21]. Most of these are boundary-based approaches [1][14] (as opposed to the regional approach taken here) so that the term *scale* refers specifically to curvature properties of the boundary. With these multiscale methods it is unclear how representations may be constructed since again there are no principled means for determining which scale is appropriate for a given condition. In this paper, we suggest that the scale associated with enclosed regions provides one practical way of addressing this problem. These regions make it possible to establish the range of scales that are appropriate when analysing a given boundary segment.

There are a few region-based approaches which consider multiple scales directly. Some methods have employed multiscale (or resolution) skeletons [15][5], while another has been based upon locating isolated symmetric points within a continuous scale space [11]. For these multiscale region-based approaches, it is again unclear how one might use the scale space to develop practical shape representations for visual tasks. A main contribution of our work is the demonstration of how regional scale information may be used explicitly to locate and describe objects in an image.

### 1.3 Advantages

Our approach provides several advantages in overcoming problems which have traditionally encumbered previous computational methods. Firstly, no a priori segmentation of objects from ground is required. It is possible to locate points which are potentially on the inside of object regions and which, at the same time, can be used as the basis for a representation of that region. Secondly, by employing annular operators, there is a built-in tolerance which can compensate for traditionally undesirable edge data attributes such as contour gaps and small-scale contour variations. Thirdly, since symmetry is identified as a function of scale, it is possible to practically separate image structures occurring at different scales. One final advantage is that scale information is used *directly* in the construction of the object shape descriptors. It appears that this has not been done previously in computer vision. From a computational standpoint, the annular sampling structure is entirely parallel and lends itself to efficient hardware implementations.

## 2 Method

#### Overview

We propose to locate objects within an image and derive coarse object descriptions based solely on a set of extracted edge features, . The method can be described in terms of three computational stages: (i) the extraction of object *features* (edges) over a range of scales (see Section 2.1), (ii) the identification of object *parts* based on symmetric relationships between object features (see Section 2.2), and (iii) the grouping of parts so that object *descriptors* may be constructed (see Section 2.3). The resulting object descriptor is expressed as a graph which captures the spatial relationships between various blobs and limbs. These are ranked using a *perceptual significance* heuristic which measures their relative prominence as visual entities. The ranking is employed in the hierarchical construction of object descriptors, and the resulting graph permits individual objects to be located in terms of their constituent parts. Furthermore, the graph provides a coarse means of characterizing an object's structure.

### 2.1 Feature Extraction

Starting with an image, contrast edges are extracted over a range of scales (see Figure 1). At a particular scale, edges are detected by employing elongated first derivative operators having a form similar to those in [4]. Operators are located on a sampling grid which spans the image and whose spacing increases with scale. Having detected an edge at a particular spatial location, its quantized *orientation* value and contrast *magnitude* are recorded to be used in subsequent processing stages.

### 2.2 Feature Grouping

Edge features are grouped in two steps. The first involves the application of annular operators to detect symmetrical relationships within the feature set. This results in the identification of a set of enclosed *symmetry points.* In the second step, the symmetry points are grouped in order to infer the

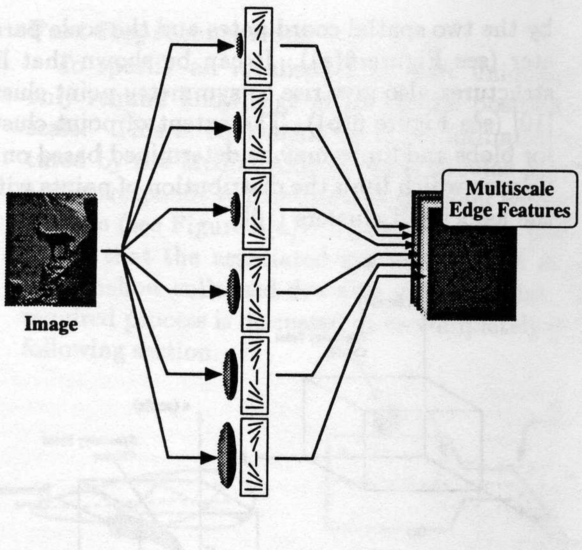


Figure 1: Edges are extracted using oriented operators applied over a range of scales with the same number of operator orientations chosen for each. The resulting object features are represented by a set of $N$ edge maps, each of which will be used by a specific set of annular operators at the next stage.

presence of blob and limb *parts*. Due to the brevity of this paper, it is possible to provide only a superficial description of the computational steps taken to implement feature grouping. A more detailed description is provided elsewhere [8][10].

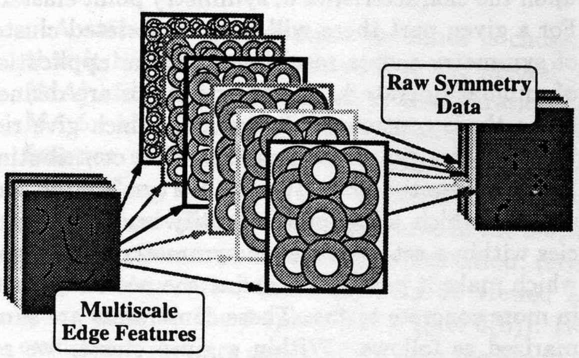


Figure 2: For each of the $N$ scales used to detect edges, there is a corresponding *layer* of annular operators. Each layer consists of a set of similar annular operators whose size is set in proportion to the edge operators of the associated feature map. The spatial extent of these operators is configured so as to cover the image plane, thereby enabling enclosing structures to be detected at any position [8].

## Identifying Symmetry Points

A set of annular operators is used to detect enclosure relationships between edge features (see Figure 2). Assuming that edges are found within a given annular region, for each edge segment at a particular location, it is possible to specify the annular tangent direction by projecting radially onto either of the annular bounding circles. If the edge orientation is consistent with the annular tangent direction, then the presence of an edge is recorded as a function of its angular position relative to the annular center [9]. This is done for each edge segment within the annular operator, and the resulting function, called an *angular profile*, is found for all operators in the sampling space. Two forms of enclosure are identified using these angular profiles (see Figure 3). Based on the angular profile, it is

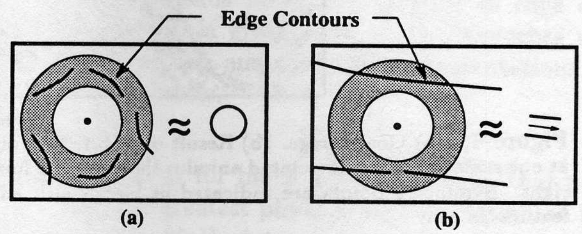


Figure 3: (a) Blob enclosure is characterized by a contour configuration within the annulus which does not favour a particular direction and where gaps in the angular profile are relatively small. Such a configuration may be viewed as approximating a closed circular contour. (b) For limbs there is a direction associated with the enclosing contour due to significant breaks present in the angular profile. Each break angle may be assigned a direction relative to the annular center. The direction of the limb enclosure (shown by an arrow in the figure, modulo $\pi$ radians) is then represented by the average of these two directions.

also possible to measure the degree of enclosure associated with a particular contour configuration [9]. This value is subsequently used to quantify the perceptual significance of blobs.

A sampling strategy has been designed so that symmetric configurations of edges can be detected using annular operators which, for a given scale, span the image plane [8]. Each of the $N$ sets of symmetry data which result from annular sampling corresponds to one of the $N$ layers in Figure 2. An example of symmetry data extracted by one annular layer is shown in Figure 4. In the next stage, this raw symmetry data is grouped in order to infer object parts.

## Parts

Before considering the second stage of feature grouping, we examine how object structure is manifested within a set of extracted symmetry points. The concepts of blob and limb enclosure have been introduced above as a means of representing the regional characteristics of a contour by a set of isolated points. It is possible, however, to extend these concepts so that *groups* of symmetry points sharing common characteristics may be identified. These groups will be referred to as *parts*. In accordance

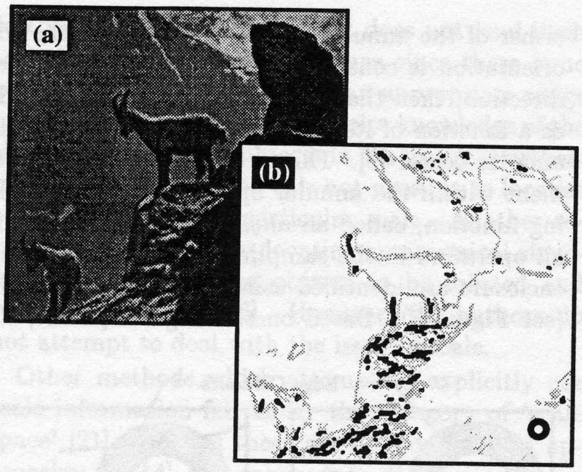


Figure 4: **(a) Goat image. (b) Result of feature grouping at one scale, with the associated annulus shown at the lower right. Symmetry points are indicated in black, with edge features in gray.**

with the symmetry extraction process, we choose to make a distinction between two different types of parts - namely blobs and limbs. A blob is defined by a contour configuration giving rise to blob enclosure, whereas a limb is defined by a contour configuration giving rise to a *set* of symmetry points which manifest limb enclosure. Prototypical examples of *idealized* blob and limb enclosing structures [10] are shown in Figure 5. Based on these examples it is possible to intuitively understand the essential nature of what constitutes a part. Details required to fully define blob and limb parts will become evident during our discussion of symmetry point clusters and part representations.

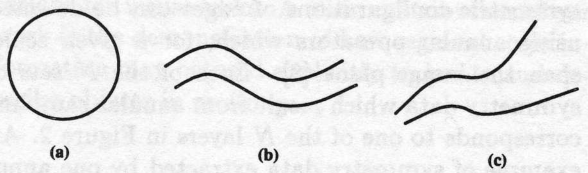


Figure 5: **(a) A circle represents the idealized blob enclosing structure. (b) A 'worm-shaped' contour represents the idealized limb enclosing structure. (c) It is also possible that the sides of the limb will converge.**

### Symmetry Point Clusters and Consistency

There is a degree of *redundancy* associated with the application of annular symmetry operators [8]. Thus for an idealized blob structure, there will be several operators that detect blob enclosure within at least one layer. Similar symmetry point distributions will also be found in neighboring layers. We find that an idealized blob contour produces a *cluster* of symmetry points within the three-dimensional space (hereafter referred to as the $\mathcal{S}$-space) spanned by the two spatial coordinates and the scale parameter (see Figure 6(a)). It can be shown that limb structures also give rise to symmetry point clusters [10] (see Figure 6(b)). The extent of point clusters for blobs and limbs may be determined based on envelopes which limit the distribution of points within the associated clusters [10].

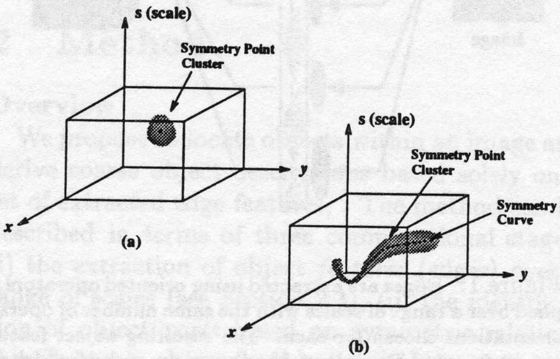


Figure 6: **(a) A symmetry point cluster for a blob (shaded) is shown within the three-dimensional $\mathcal{S}$-space. A single point located towards the center of the cluster may be used to represent the blob . (b) A symmetry point cluster for a limb (shaded). A *symmetry curve* (dashed) located towards the center of the cluster is used to represent the limb.**

It is possible to refine our part definition based upon the characteristics of symmetry point clusters. For a given part there will be an associated cluster of symmetry points resulting from the application of annular operators. Blobs and limbs are defined to be those contour configurations which give rise to symmetry point clusters where all contributing points are mutually *consistent*. We can impose constraints which allow us to identify such consistencies within a set of extracted symmetry points, and which make it possible to define the notion of parts in more concrete terms. These constraints are summarized as follows. Within a given cluster we require that all points manifest only one type of enclosure (either blob or limb, exclusively). In both cases, spatial and scalar continuity constraints are imposed, and in the case of limbs, a further smoothness constraint is applied to limb orientation values.

With these constraints it is hypothetically feasible to devise a computational scheme to perform symmetry point grouping. However, it is not immediately clear how the resulting groups could be used to provide useful representations of the original data. Before considering the process used to group symmetry points, it is necessary to examine the type of information considered to be useful for representing blobs and limbs.

## Part Representations

To specify an idealized blob part uniquely, we only require knowledge of its spatial position and scale. This is equivalent to the information captured by the Medial Axis [2] for a circular contour, and corresponds to locating a single point within the $\mathcal{S}$-space (see Figure 6(a)). To obtain such a blob requires that the associated symmetry point cluster be somehow collapsed down to a single point. The required process is discussed more completely in the following section.

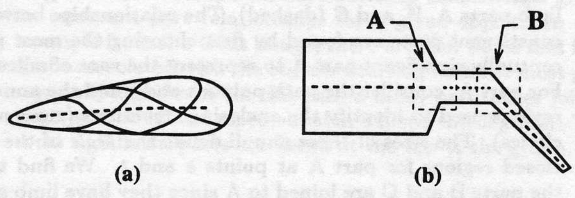


Figure 7: **A *symmetry axis* represents the projection of a symmetry curve onto the spatial plane. (a) A symmetry axis (dashed) may be extracted in settings where the contour has internal structure, thereby capturing information pertaining to the overall shape of the outer boundary. (b) An object contour which is decomposed into three limbs indicated by three symmetry curves (dashed lines). In grouping symmetry points, several constraints are imposed. This grouping produces a result which differs substantially from the Medial Axis Transform [2]. Constraints are based upon: scalar continuity (A), orientational continuity (smoothness) (B), and spatial continuity (A and B).**

A suitable representation for limbs is chosen so that it bears close resemblance to the Medial Axis (MA) of the associated contour configuration. The MA for a limb consists of a curve and an associated function which defines the regional size or *scale* at points along the curve. This information can be expressed equivalently as a space curve in the $\mathcal{S}$-space. We have chosen a particular representation, termed a *symmetry curve*, that may also be viewed as a space curve in the $\mathcal{S}$-space (see Figure 6(b)). However, the symmetry curve differs from the MA in two fundamental ways. First, a symmetry curve may be extracted in cases where there is internal structure (similar to the SLS of [3]; see Figure 7(a)). Second, since grouping is performed by applying a set of consistency constraints, the symmetry curve is structurally different from the MA (see Figure 7(b)). The symmetry curve is approximated by a set of points defining a *path* within the discretely sampled $\mathcal{S}$-space. The process for extracting these limb part representations is discussed more completely in the following section.

## Extracting Part Representations

With the requirements for blob and limb representations in mind, we now consider a process whereby part representations are extracted directly from a set of symmetry data. The part extraction process depends upon the identification of *paths* within the $\mathcal{S}$-space. Based on the set of consistency constraints (see Figure 7(b)) a valid path is found by performing a constrained search within the set of extracted symmetry points. Part representations are extracted in a sequential order based upon their *physical support* [10]. For both blobs and limbs, this measure depends upon the quantity of edge data that supports a given part. The approaches used to extract blob and limb part representations are described below.

### BLOBS

The search for blobs is sequential, so that parts having the greatest physical support are found first (see Figure 8(a)). A two stage process is used to collapse a cluster of consistent symmetry points down to a single representative point. In the first stage, a path consisting of mutually consistent symmetry points is found. For a given symmetry point cluster, the consistency constraints previously described (see above) make it possible to define a unique path using only those points which demonstrate a local maximum in the degree of enclosure [9] within a particular layer. Paths which correspond to individual blob structures will traverse the $\mathcal{S}$-space only along the scale dimension (due to sampling redundancies) with little variation along either of the spatial dimensions. Thus, points in the resulting locus will approximate a line segment in $\mathcal{S}$-space. In the second stage, the path point which exhibits a maximum degree of enclosure is chosen to represent the entire part.

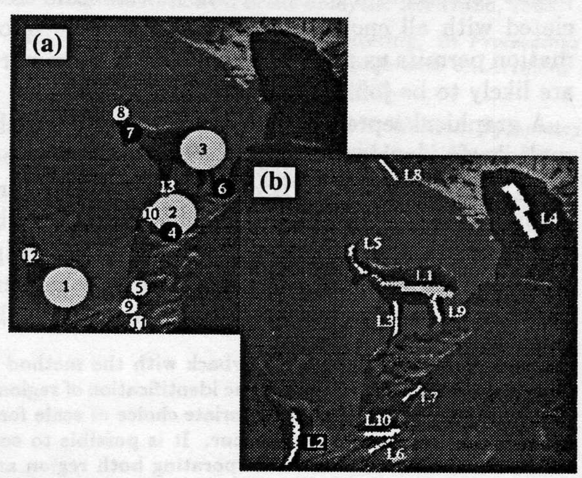


Figure 8: **Blobs (a) and limbs (b) extracted from the Goat image are numbered in order of decreasing physical support (1 being the greatest).**

LIMBS

Within the $\mathcal{S}$-space, paths corresponding to limbs are constructed using a graph search. While a particular path is being constructed, grouping constraints (see Figure 7 and Symmetry Point Clusters and Consistency section above) are imposed to ensure that all path points remain within the same limb cluster. Optimization is performed to find the part with the maximum physical support value over the set of all possible paths. Having found the corresponding path, all related cluster points (those consistent with the optimal part) are removed, and the search is reinitialized to extract the next most supported part (see Figure 8(b)).

## 2.3 Part Grouping

The grouping of symmetry points results in a set of ranked blobs and limbs. The final computational stage may be viewed in terms of yet another grouping process. This time parts are grouped together to form descriptors for objects in the scene. The result of this grouping is a *graph*, where nodes correspond to object parts and edges represent the spatial relationships between the various parts.

Starting with a set of extracted object parts, grouping begins with the part having the highest *perceptual significance* ranking. The perceptual significance of parts is computed by comparing the set of extracted parts based on quantities such as enclosed area, length, and average edge magnitude. Since the grouping method does not explicitly employ boundary-based information, we have no knowledge of high curvature points (or corners) that are often used to indicate the presence of part junctions[2]. For any given part, the set of symmetry points represents a record of the spatial scale associated with all enclosed regions. This scale information permits us to predict situations where parts are likely to be joined (see Figure 9).

A graphical representation is used to capture the spatial relationships between the constituent object parts (see Figure 10). For a given primary part, other *secondary parts* that are joined to it may be represented in terms of their spatial relationships with the primary part. These are expressed by identifying *landmark* positions along a limb axis, the

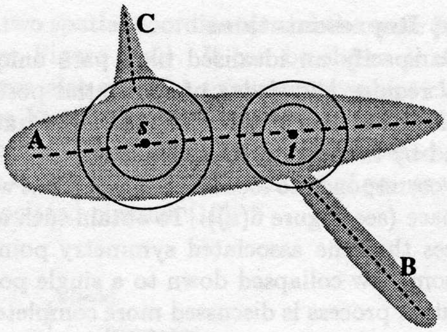


Figure 9: This object (filled) is decomposed into three limb parts A, B, and C (dashed). The relationships between constituent parts are found by first choosing the most perceptually significant part A to represent the root of a graph. For part A, consider the path points $s$ and $t$ and the annular regions used to identify the enclosing structures (concentric circles). The sizes of these annuli define the *scale* of the enclosed regions for part A at points $s$ and $t$. We find that the parts B and C are joined to A since they have limb axes which end within the annular regions at $s$ and $t$. Thus, for part A, the scale at which a join occurs is consistent with the scale of the enclosed region at that point. The points $s$ and $t$ are found and used as landmark positions to locate the parts B and C along the path representing part A. These particular points were found firstly because a join can be detected at these locations, and secondly because they are spatially closest to the joining part.

*scale* at which a join occurs, or the *angular position* of a join relative to a limb axis or to other joins.

# 3 Summary

We have developed a machine vision system based upon the premise that visual sensing is primarily a means for acquiring object information. To accomplish this we have proposed a means for feature grouping using principles found in studies of human perceptual grouping processes.

Grouping is first performed by detecting *symmetrical enclosing* configurations of edge features using annular symmetry operators. By explicitly using scale in this process, symmetric relationships can be extracted in cases which require the identification of gross object structures, independent of internal or external structure (or texture). Following this, symmetry points are then grouped in order to infer blobs and limbs by finding paths within the discretely sampled $\mathcal{S}$-space. Finally, a graph representation for objects is constructed automatically based upon spatial and scalar relationships between parts.

Starting with the 370 × 390 pixel Goat image shown above, all computations required to produce graphical results for the entire image execute within

[2]While this is an obvious drawback with the method in its current form, we believe that the identification of regional scale must occur before an appropriate choice of scale for a boundary-based approach can occur. It is possible to conceive of a hybrid approach, incorporating both region and boundary features, where the identification of enclosed regions would allow for the establishment of *stable reference points* and would also make it possible to select an *appropriate scale* before attempting curvature analysis.

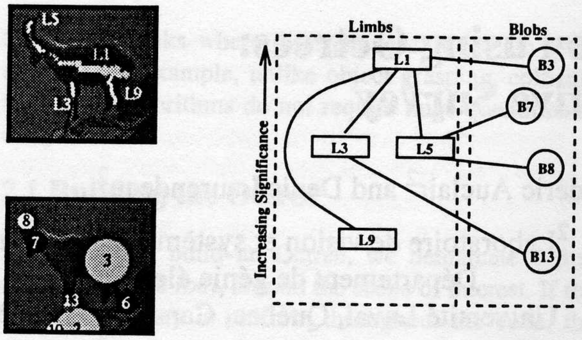


Figure 10: The system automatically constructs a graph to represent part relationships for a given object. In this example from the Goat image, there are four blobs and four limbs associated with the object. Starting with the most perceptually significant (in this case a limb) part, other joined parts are found and represented in a hierarchical graph. The most significant parts are found at the top of the graph.

approximately ten minutes when implemented on a SPARC 10 workstation. Due to the inherent parallelism of the annular processing, it is likely that an implementation using parallel hardware would lead to substantial improvements in execution time. The resulting system could potentially provide the basis for developing a real-time vision system for applications such as autonomous robotics. The results presented in this paper demonstrate that our method can provide a means of visually locating and deriving coarse representations for objects in a scene.

## Acknowledgements

The authors would like to thank the Canadian Institute for Advanced Research and PRECARN for its support. This work was partially supported by a Natural Sciences and Engineering Research Council of Canada Strategic Grant and an FCAR Grant from the Province of Quebec.

# References

[1] H. Asada and M. Brady. The curvature primal sketch. Technical Report 785, MIT, 1984.

[2] H. Blum. Biological shape and visual science (part I). *Journal of Theoretical Biology*, 38:205–287, 1973.

[3] M. Brady and H. Asada. Smoothed local symmetries and their implementation. Technical Report A.I.Memo 757, Massachusetts Institute of Technology AI Lab, 1984.

[4] J. Canny. A computational approach to edge detection. *IEEE Trans. PAMI*, 8(6):679–697, 1986.

[5] A. R. Dill, M. D. Levine, and P. B. Noble. Multiple resolution skeletons. *IEEE Trans. PAMI*, 9(4):495–504, 1987.

[6] T. Kanade. Recovery of the three-dimensional shape of an object from a single view. *Artificial Intelligence*, 17:409–460, 1981

[7] L. Kaufmann and W. Richards. Spontaneous fixation tendencies for visual forms. *Perception and Psychophysics*, 5(2):85–88, 1969.

[8] M. F. Kelly and M. D. Levine. A sampling strategy using multi-scale annular operators. Technical Report TR-CIM-93-20, Center for Intelligent Machines, McGill University, Montreal, Canada, 1993.

[9] M. F. Kelly and M. D. Levine. The symmetric enclosure of points by planar curves. Technical Report TR-CIM-93-1, Center for Intelligent Machines, McGill University, Montreal, Canada, 1993.

[10] M. F. Kelly and M. D. Levine. From symmetry to representation. Technical Report TR-CIM-94-12, Center for Intelligent Machines, McGill University, Montreal, Canada, 1994.

[11] B. B. Kimia, A. Tannenbaum, and S. W. Zucker. Towards a computational theory of shape: An overview. In *Proceedings of the European Conference on Computer Vision*, pages 402–407, 1990.

[12] W. Köhler. *Gestalt Psychology.* New American Library, New York, 1947.

[13] D. Marr and E. Hildreth. Theory of edge detection. *Proceedings of the Royal Society (London), Ser. B*, 207:187–217, 1980.

[14] F. Mokhtarian and A. Mackworth. Scale-based description and recognition of planar curves and two-dimensional shapes. *IEEE Trans. PAMI*, 8(1):34–43, 1986.

[15] S. M. Pizer, W. R. Oliver, and S. H. Bloomberg. Hierarchical shape description via the multiresolution symmetric axis transform. *IEEE Trans. PAMI*, 9(4):505–511, 1987.

[16] K. Rao and R. Nevatia. Describing and segmenting scenes from imperfect and incomplete data. *CVGIP: Image Understanding*, 57(1):1–23, 1993.

[17] I. Rock. *The Logic of Perception.* MIT Press, Cambridge Mass., 1983.

[18] A. Rosenfeld. Axial representations of shape. *Computer Graphics and Image Processing*, 33:156–173, 1986.

[19] E. Rubin. *Visuell wahrgenommene Figuren.* Copenhagen, 1921.

[20] L. Shapiro, J. Moriarty, P. Mulgaonkar, and R. Haralick. Sticks, plates and blobs: a three-dimensional object representation for scene analysis. *AAAI*, 80, 1980.

[21] A. P. Witkin. Scale space filtering. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pages 1019–1022, 1983.

[22] Y. Yeshurun, D. Reisfeld, and H. Wolfson. Symmetry: a context free cue for foveated vision. In *Neural Networks for Perception*, volume I Human and Machine Perception, pages 477–491. Academic Press Inc., 1992.