Image Description via Annihilation of Essential Structures

TR88-001

REVISED

December 1988

Stephen M. Pizer[1,3]
John M. Gauch[1]
Lawrence M. Lifshitz[4]
William R. Oliver[1,2]

The University of North Carolina at Chapel Hill
Department of Computer Science
CB#3175, Sitterson Hall
Chapel Hill, NC 27599-3175

[1]*Departments of Computer Science,* [2] *Pathology, and* [3] *Radiology.*
[4] *School of Medicine, University of Massachusetts, Worcester, MA*

Image Description via Annihilation of Essential Structures

Stephen M. Pizer[1,3], John M. Gauch[1],
Lawrence M. Lifshitz[4], William R. Oliver[1,2]
Departments of Computer Science[1], Pathology[2], and Radiology[3]
University of North Carolina
Chapel Hill, North Carolina, USA
and
School of Medicine[4]
University of Massachusetts
Worcester, Massachusetts, USA

Abstract

A fundamental approach for providing an image description in terms of visually sensible image regions is described. It involves a) the representation of the image by a structure that captures essential image information and then b) the definition of a hierarchy of components of that structure by the order of annihilation of those components as the image is continuously simplified by lowering the scale. The information-capturing "essential structure" is chosen so that image regions are associated with each structure component during the image simplification. To guarantee image simplification, successive Gaussian blurring is chosen as the means of scale lowering.

A number of candidates for essential structures are discussed. We argue that an essential structure that describes shape in both the spatial and intensity dimensions will produce an image description most likely to be useful for computer or human specification of image objects. In particular, we suggest that the intensity "ridge" and "course" curves defined by the locus of intensity level curve vertices, augmented by the pile of internal and external symmetric axes of these level curves, satisfies all desirable criteria for an essential structure. With such shape-based essential structures the approach of image description via annihilation under image simplification becomes a very attractive paradigm.

Introduction

Any process for the definition and labeling of objects appearing in images benefits from transforming the original image data into a description in terms of visually sensible regions. With such a description a source of intelligence, be it a human interacting with the display of the image or a computer program exhibiting artificial intelligence, has a good basis for fitting the image information to its model of the world in order to recognize an object.

In this paper we discuss approaches to producing such a useful image description. In particular, we describe the idea of generating a description by measuring an essential structure in the image and following it to annihilation as the resolution of the image is reduced. We give five properties which ensure that the image description is well behaved for image analysis. Several essential structures that we have recently investigated for image description are presented. The power of the overall approach is thereby illustrated, and the relative strengths of the structures for producing a useful description are compared.

Early Multiresolution Analysis

The most popular models of the human visual system (Robson [1983], Koenderink [1984], Wilson [1979], Ginsburg [1977]) recognize that it preprocesses the image by analyzing it simultaneously at multiple scales. In computer vision Crowley [1984] realized early that analysis at multiple scales could provide an important means of image description on which model-based pattern recognition could be based, and not just efficient analysis, as suggested by many (e.g., Burt [1983], Rosenfeld [1984]). Crowley based his analysis on various Difference Of Gaussians approximations to the Laplacian of the image. He followed peaks and ridges (or their negative counterparts) in this Laplacian image through many scales while keeping the energy of the blurred Laplacian operator constant. Describing the image involved locating the scale at which each peak appeared most strongly.

A somewhat more attractive idea is to take advantage of a blurring approach that simplifies the image and to define objects in terms of the *disappearance* of their features with simplification. The idea is that image objects are defined first by regions of large scale, with detail of these objects defined by regions of smaller scale. Regions of large scale are those that are retained as the image is simplified by reducing resolution (blurring), while small-scale regions disappear under less blurring. The description must also include the relation between small- and large-scale regions.

Witkin [1983], Yuille [1983], and Koenderink [1984, 1988] each suggested that Gaussian convolution was the best form of blurring, since it guaranteed image simplification with blurring, i.e., was the only form of blurring that did not allow the local creation of new values of any linear function of derivatives of the image as the blurring proceeded. Thus, for example, neither local image intensities (0th derivatives) nor Laplacian zeroes are created by this process. One of us [Lifshitz, 1987a] has shown that the required Gaussian blurring need be neither isotropic nor stationary for the simplification guarantee to be met, and he has suggested that variation of the parameters of the blurring Gaussian across the image could be used to reflect *a priori* or tentative knowledge about the scene.

Essential Structures and Their Annihilation

Using the notion of following image features through simplification, Koenderink [1984] suggested that the following of intensity extrema and of iso-intensity paths through Gaussian blurring could define sensible image regions: you followed each extremum to annihilation with a saddle point and defined the region as those locations whose iso-intensity paths ran into the path in scale space tracked by the extremum (the *extremal path*). Koenderink and we realized that this approach could be used to form an image description made from a hierarchy of these regions, where regions lower in the hierarchy were of smaller scale and blurred into their parent regions in the hierarchy.

We suggest that a most important feature of this approach was that image regions were defined by the annihilation of their extrema under blurring, or to take a more constructive point of view, by the creation of these extrema as deblurring was successively applied to the fully blurred image. In this paper we develop a generalization of this idea of creating an image description that is hierarchical by scale by following what we call *essential structures* to annihilation.

The concept is that an essential structure should be an image descriptor that has the following five properties:

1. It induces a subdivision of the image into regions.
2. It captures essential region properties, including the way intensity varies across it and the spatial properties of the region, i.e., its shape, and therefore the regions it induces are semantically sensible.
3. The structure relating image components does not change until a component annihilates.
4. It induces a hierarchy of regions by defining for each component the containing component into which it annihilates.
5. It is applicable for images of any spatial dimension.

Intensity Extrema and Iso-intensities as Essential Structures

Based on the idea of Koenderink [1984], we began by choosing as our essential structure the set of intensity extrema, augmented by iso-intensity contour segments [Lifshitz, 1987a,b]. As illustrated in Figure 1, the technique is to follow *extremal paths*, i.e. the tracks of each extremum as image blurring increases and the extremum intensity changes monotonically, until the extremum annihilates with a saddle point. *Iso-intensity paths*, i.e., paths connecting each image location to the closest location at the next higher scale with the same intensity, are also followed until they run into an extremal path.

This essential structure satisfies the following three properties of essential structure.

Region Definition. This method defines image regions by following iso-intensity paths, i.e., paths in scale space connecting each image location to the closest location at the next higher scale with the same intensity. These iso-intensity paths run into extremal paths as intensity changes monotonically along the extremal paths. The resulting association of image locations with an extremum forms image regions.

Region Hierarchy. The regions formed by linking iso-intensity paths to extremal paths are associated as subregions of other extremal regions according to the behavior of the iso-intensity paths which begin where an extremum annihilates. The extremal path into which this new iso-intensity path eventually links identifies the parent region for the subregion (for example, see Figure 1).

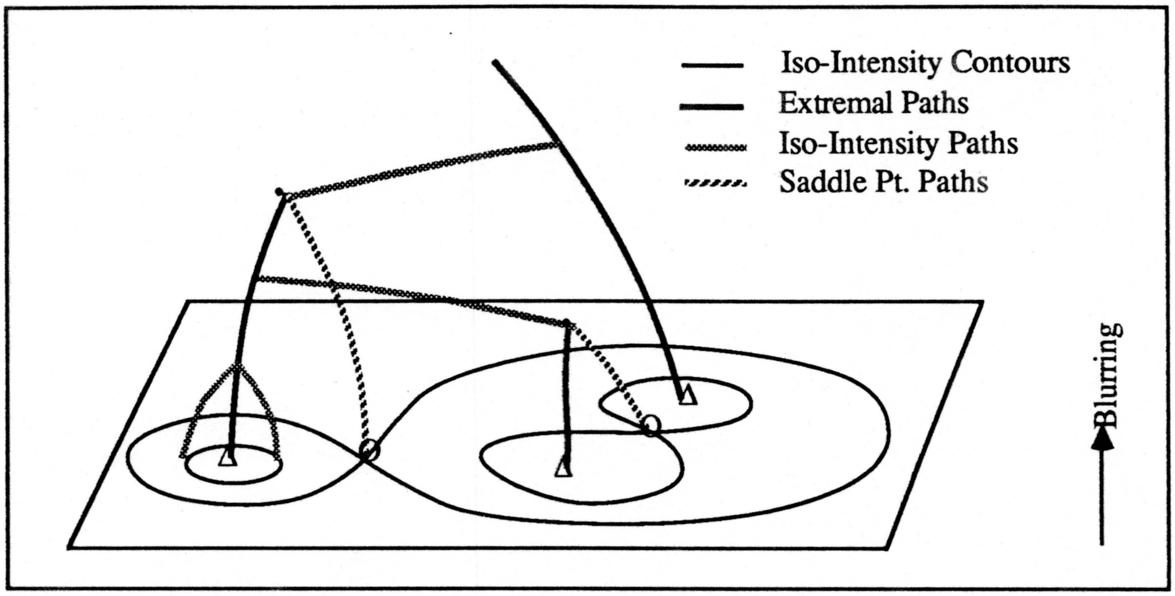

Figure 1: The behavior of extremal paths under resolution reduction. Note that maxima (Δ), and saddle points (O) move together and annihilate. The resulting non-extremal point (•) is then linked via an iso-intensity path to another extremal path.

Generalization to all dimensions. Lifshitz has demonstrated the method for 2D and 3D images and has described how the image representation can be extended to any number of dimensions.

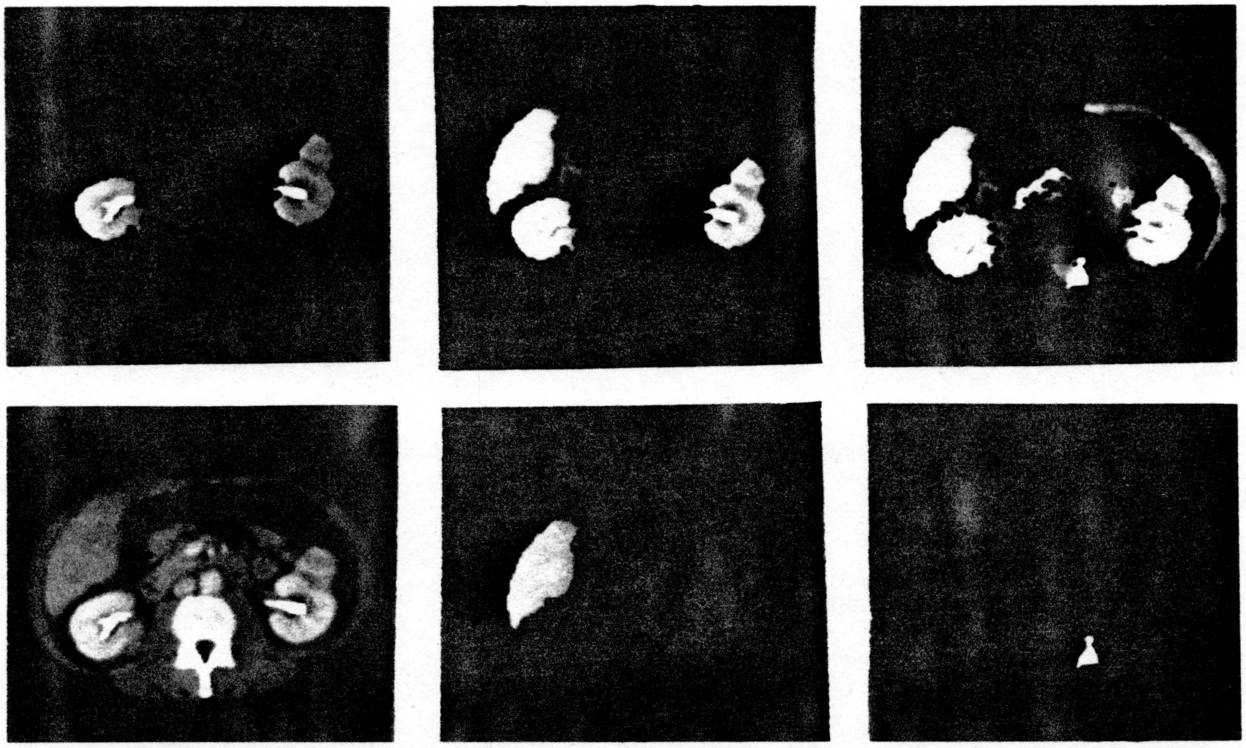

Figure 2: A CT slice through the upper abdomen together with a collection of anatomic regions automatically defined using Lifshitz's program.

However, this image description fails on two of the above-mentioned criteria.

Consistent Simplification. Extrema can misbehave under Gaussian blurring. Koenderink has shown that intensity extrema can be created out of nowhere by Gaussian blurring, for example when the image is made from two Gaussian peaks connected by a steep, narrow hill whose ridge monotonically falls from the higher of the peaks to the lower. The effect of intensity extremum creation is that regions can be associated with extrema that appear at scales greater than that of the original image rather than with just those in the original image. Lifshitz's results suggest that no difficulty arises from this fact.

Region Sensibleness. Applying these ideas to 2D and 3D medical images, we have shown that the regions in the description thus produced frequently form anatomic objects, or can be easily formed into such anatomic objects using the operations of union and difference (for example, see Figure 2). However, sometimes the resulting regions are not semantically sensible. The first type of misbehavior listed is of little consequence, but the second and third are bothersome.

1. Since an extremum annihilates with a saddle point, we would expect that the extremal region should include all of the image pixels inside of the iso-intensity contour surrounding the annihilating extremum at the intensity level at which the extremum annihilates. In [Lifshitz, 1987a] this is shown not always to be true, but this failure seems to occur only in unusual circumstances.

2. An extremal region may be made of disconnected components, with neither component itself being an extremal region. If these two components are separated by another region with which each component has a visually apparent edge, we would wish that each edge

would form the boundary of an extremal region, but this does not always happen (see Figure 3). This misbehavior results from the fact that the edges of regions may be heavily blurred before either of their iso-intensity sets run into extremal paths, so both sets can run into the same extremal path and thus form parts of the same region.

3. Regions which obviously hang together as a single object do not always combine into single regions in the tree. This problem is especially apparent with long branching regions, such as a blood vessel tree.

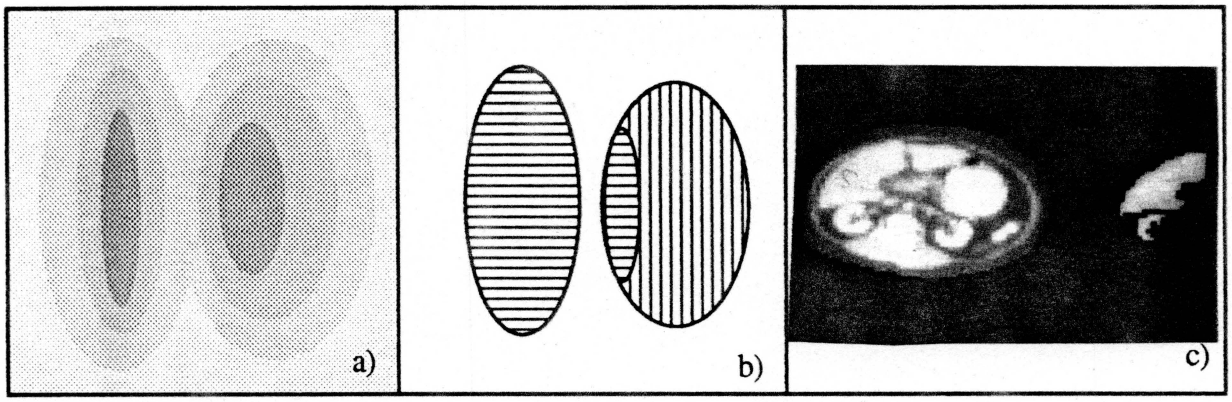

Figure 3: The halftone image in 3a depicts the situation where two dark regions in an image are separtated by a narrow lighter region. The two image segments in 3b illustrate how pixels linking to local extrema can sometimes result in unnatural image segments. In 3c we see an example where Lifshitz's program has identified part of a kidney and the whole liver as a single object in an abdominal CT image.

We suggest that these difficulties arise from the fact that the essential structure of intensity extrema and iso-intensites inadequately reflects shape and edges. Structures that are designed as shape descriptors would seem attractive candidates as essential structures. In the following we discuss a number of such candidates. We divide our discussion into those structures that describe shape via object boundaries and those that describe shape via the interior of the objects, which we call the *figure*.

Essential Structures via Boundary Shape

Richards and Hoffman [1985] have noted the importance of boundary curvature extrema in the perception of objects with these boundaries. They thus have defined boundaries in terms of pieces which range from a curvature minimum, through a curvature maximum, to another curvature minimum. They call these boundary pieces *codons*. One of us [Gauch, 1987] has noted that if we apply some form of resolution reduction to the boundary itself, these codons annihilate, one by one, and these annihilations can induce a hierarchy. For each such annihilation a maximum of boundary curvature annihilates with a minimum of boundary curvature, and we can take the remaining codon sharing the annihilated minimum to be the hierarchical parent of the annihilating codon.

A somewhat more attractive boundary-based description, produced by Leyton [1987], defines the boundary in terms of the growth of codons from one another by cardinal deformations, rather than the concatenation of codons suggested by Richards and Hoffman. The deformations take place at the maximum curvature points of the boundary. The shape description is then given by the deformation sequence rather than by a subdivision into regions. The difficulty is that the process is ambiguous; there are many sequences of cardinal deformations that can produce the same object.

By defining codons as components of those into which they annihilate under resolution reduction of the boundary, our multiresolution approach can select a particular sequence as the descriptor.

The difficulty of both of these approaches for image description is that they depend on having a predefined boundary. While there remain interesting proposals for ways in which such boundaries could be created [e.g., Grossberg, 1985], the question remains whether edges should not be the result of region definition rather than the progenitor of them. This turns our attention to figure-based methods. With such methods a pre-process of edge or contrast enhancement could be used to reflect edge information (cf. Crowley's work). Alternatively, measures of edge strength might be made part of figure-based essential structures.

Essential Structures via Figure Shape

Figure-based shape description is based on decomposition of the figure into cardinal regions. One of the most attractive approaches of this type has focused on axes of symmetry [Blum, 1978; Brady, 1984; Leyton,1987]. Of the various alternatives the symmetric, or medial, axis (SA) stands out by being a connected tree that by division at branch points induces a decomposition into regions. For objects without holes, each region has an unbranching axis and two associated boundary sides. When the SA component includes an SA endpoint, the two sides meet at a point of maximal boundary curvature. The axis is the center of a figure and its associated radius function specifies the locations inside the figure. The major weakness of the SA [Pizer, 1987a] was the fact that detail, or noise, destroyed the naturalness of the decomposition it induced. However, were resolution reduction to cause the SA branches corresponding to detail to shrink and annihilate early into the limbs off which they branch, the regions corresponding to these branches would be defined as subregions and, more important, the limbs would be restored to the natural correspondence with a single region rather than two regions interrupted by a detail. A natural hierarchy would result (see Figure 4).

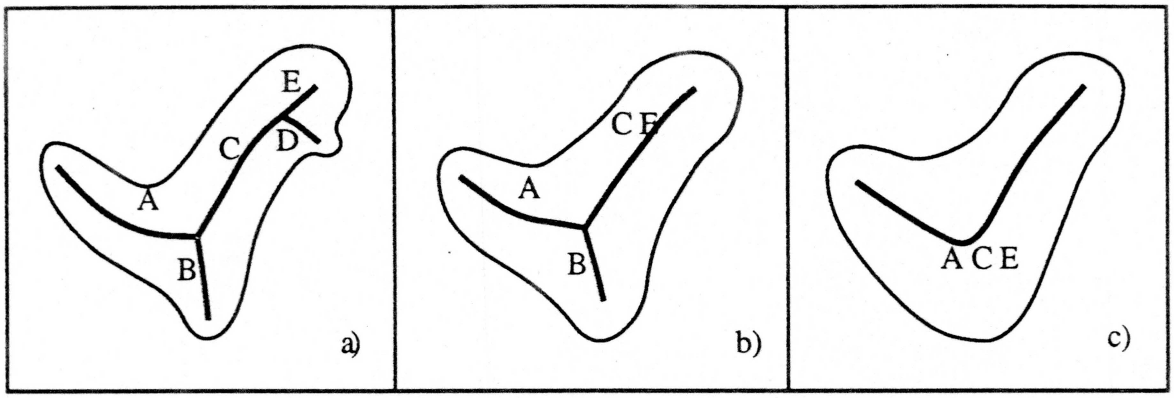

Figure 4: The desired effects of resolution reduction on the SA of figure 4a are shown in 4b and 4c. Upon annihilation of branch 'D', we would identify that branch as a subobject of the combined branch 'CE'. Similarly, branch 'B' would be determined to be a sub-branch of 'ACE'.

But what means of resolution reduction can achieve this behavior? We must avoid boundary-based approaches, partially to avoid the need for predefinition of image boundaries, partially to allow figures with nearby, disconnected pieces to be described as a single object, and partially to retain consistency with the figure-based SA. That is, resolution should be reduced by Gaussian convolution with the characteristic function specifying the figure [Koenderink, 1986a]. But Gaussian convolution turns a characteristic function into a function with many grey levels, i.e., an image. It is therefore necessary to define an axis that is sensible for an image and not just for an object. Such an axis would be applicable to characteristic functions as special cases of an image.

To accomplish this axis definition, we cannot take Koenderink's approach [1986a] of choosing an intensity threshold to make the image into a characteristic function. Such a choice causes SA branches to break into pieces as we blur the image (see criterion 3). Instead, we view the image as a terrain map, with intensity as height. Since the intensity dimension is incommensurate with the spatial dimensions, we must treat height specially and view the terrain map as being made up of a continuous pile of binary images, each corresponding to a (not necessarily planar or level) slice through the terrain, and having value 0 where there is air and value 1 where there is earth. The appropriate cut surfaces may be image dependent (see Figure 5); how to choose them is a subject of research, and we avoid that question for now by allowing only the selection of intensity level sets as the cut surfaces after preprocessing of the image (e.g., by a locally adaptive contrast enhancement [Pizer, 1987b]).

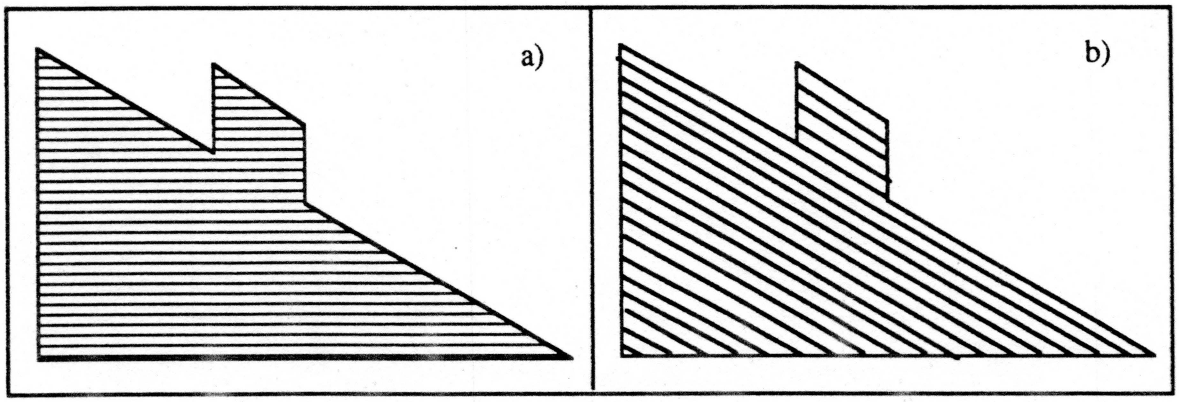

Figure 5: Intensity profile of an image with a raised area on a ramp. The iso-intensity slices in figure 5a may be less natural for describing the image than the inclined slices in 5b.

Symmetric Axis Pile and Vertex Curves as Essential Structures

The image I(x,y) has thus been characterized as a sequence of binary images, with height (intensity level L) parameterizing the selection of the slice. In order to capture both light objects on dark backgrounds and dark objects on light backgrounds, we specify two SAs on each level L of the 'terrain map': the SA of the earthen objects (x,y such that $I(x,y) \geq L$) and the SA of the air (x,y such that $I(x,y) \leq L$). Looking at these SAs as a function of level L produces what we call the symmetric axis pile (SAP). Gauch [1987] has shown that the SAP is continuous except at critical points of I(x,y). He shows that the SAP consists of a forest of branching sheets, each such branch characterizing shape in both space and intensity of a corresponding part of the image. Corresponding to each sheet in the SAP is a radius function, and a region image R defined by R(x,y) = the maximum intensity level for which the radius function of a sheet point at that level includes x,y (see Figure 6). Furthermore, branches in the SAP shrink to annihilation under Gaussian blurring of the image I (see Figure 7) inducing a hierarchy, or multiple hierarchies, on SAP sheets and their associated image regions. The question of how multiple hierarchies should be formed into one is still under investigation.

The SAP for an n-dimensional image is an n+1-dimensional forest of sheets, a prodigious object to follow through image blurring. However, Gauch has pointed out that since each SA is terminated by a boundary point of maximum curvature magnitude (a *vertex*), each SAP sheet is terminated by a curve of level curve vertices (a *vertex curve* -- see Figure 8). These vertex curves are simply tracks in the original image, corresponding to ridges or courses in the "terrain map". They can be followed through image blurring, and when a vertex curve annihilates, the SAP sheet that it terminates must also annihilate. Therefore, it is possible to compute the SAP only for the original image, and for each vertex curve annihilation to follow the corresponding SAP sheet to its branch curve. The sheet defines a region image R and specifies R as a subregion of the region image corresponding to the limb sheet into which it connects.

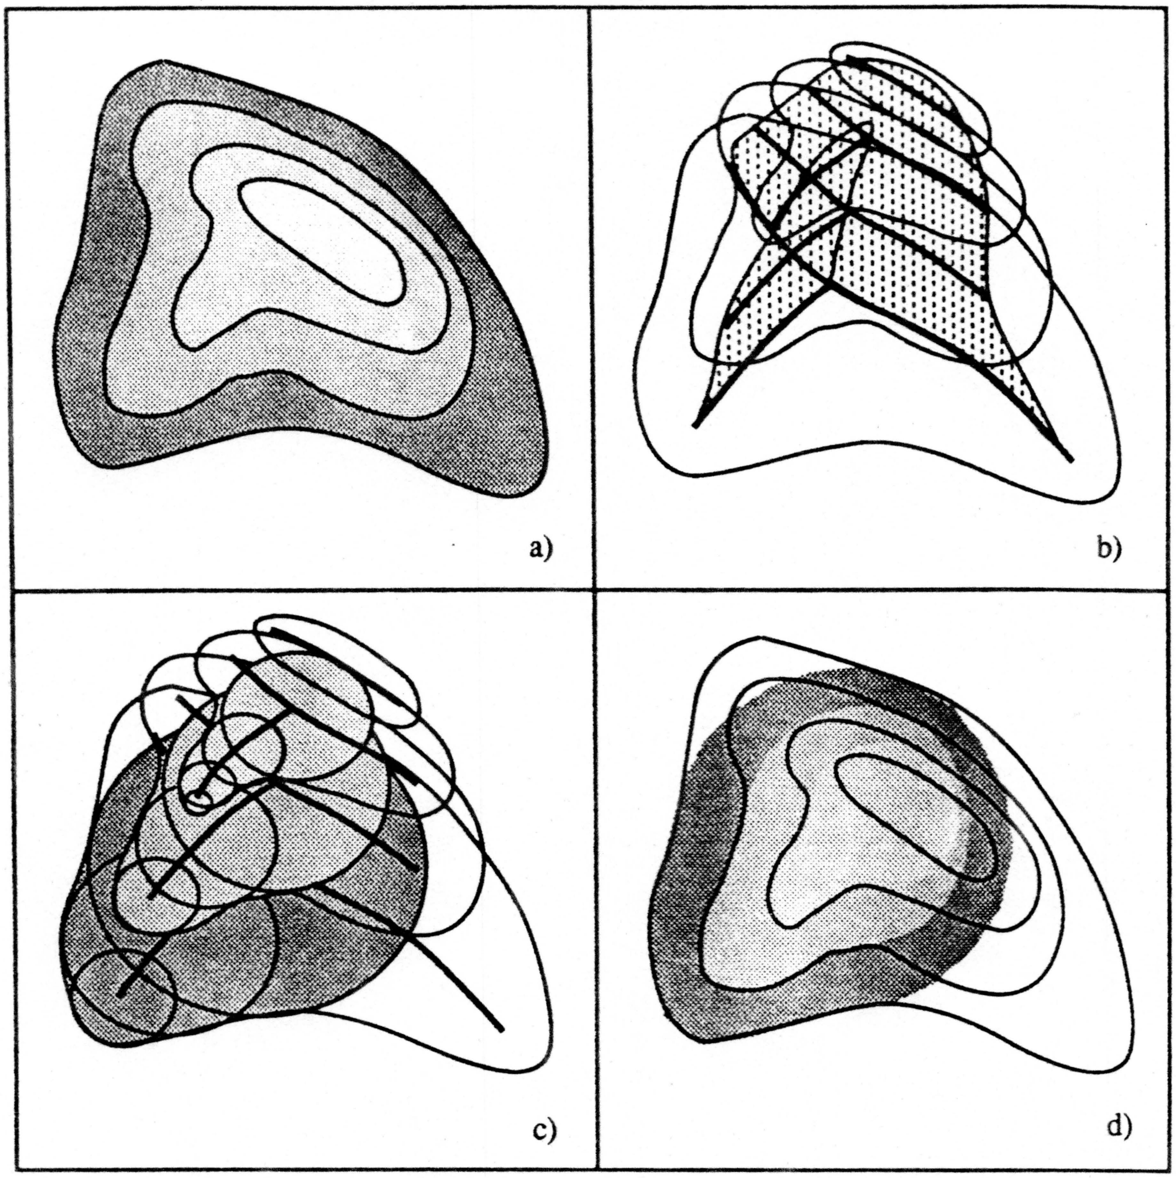

Figure 6: a) A simple grey-scale image represented by four level curves, b) its associated symmetric axis pile with the SA for each level shown in bold, c) the union of maximal circles centered on one branch of the SAP, d) the highest intensity at each point within this union yields an image associated with this SAP sheet.

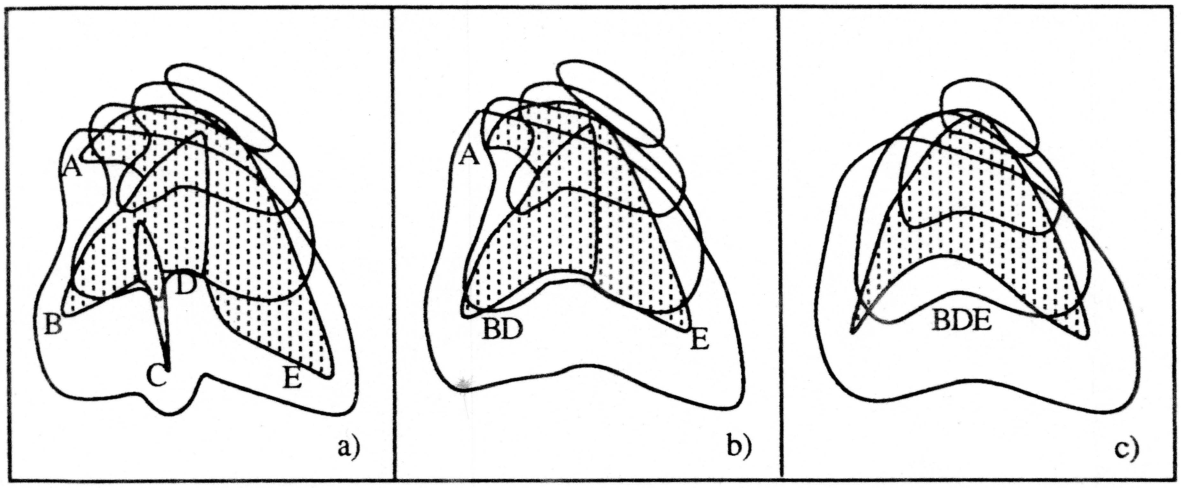

Figure 7: The effects of resolution reduction on the earthen SAP of figure 7a are shown in 7b and 7c. The branch sheets appear under the ridges of the image represented by the level curves. When branch 'C' annihilates, we identify it as a subobject of the combined branch 'BD'. Similarly, branch 'A' is determined to be a sub-branch of 'BDE'.

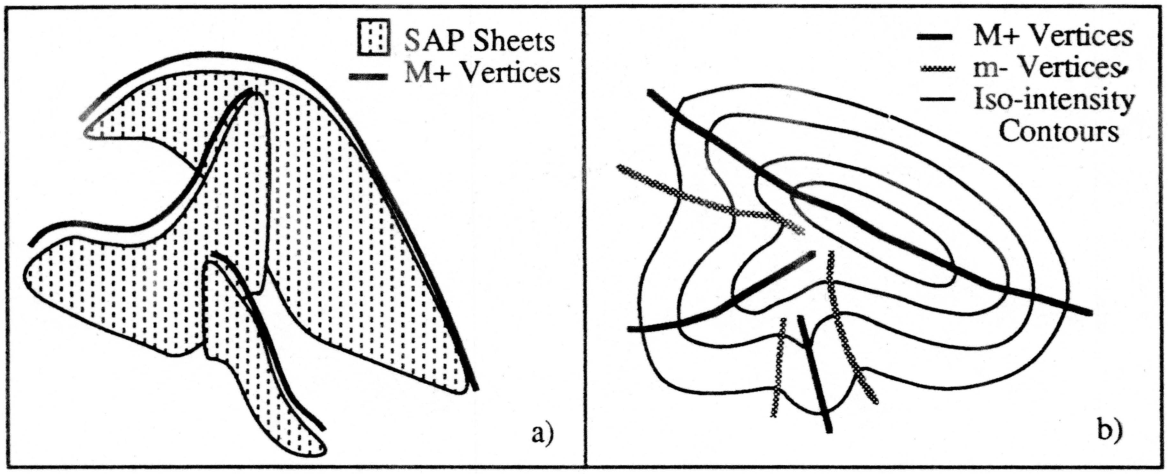

Figure 8: a) The relationship between the symmetric axis pile for an image and the vertex curves corresponding to the end curves of the 'earthen' SAP. For clarity, only the vertex curves corresponding to positive curvature maxima (M+), are shown. b) Iso-intensity contours and vertex curves corresponding to positive curvature maxima (M+) and negative curvature minima (m-).

The correspondence between vertex curves and image structures is illustrated in Figure 9. Here, we have applied various degrees of Gaussian blurring to a digital subtraction angiogram, an image of blood vessels (top row) and calculated the corresponding level curve curvature (second row) and the vertex curves (third row) for these images. In the curvature images magnitude of curvature is shown by the grey level; high positive curvature (M+) is shown in white, and low negative curvature (m-) in black. The grey points on these black and white curves correspond to saddle points in the image (where curvature is undefined). Only M+ and m- vertex curves can be the tops of SAP sheets, and these SAP sheets can help in following the vertex curves across saddle points. The vertex curves move continuously to annihilation and induce the hierarchical description described in the previous paragraph.

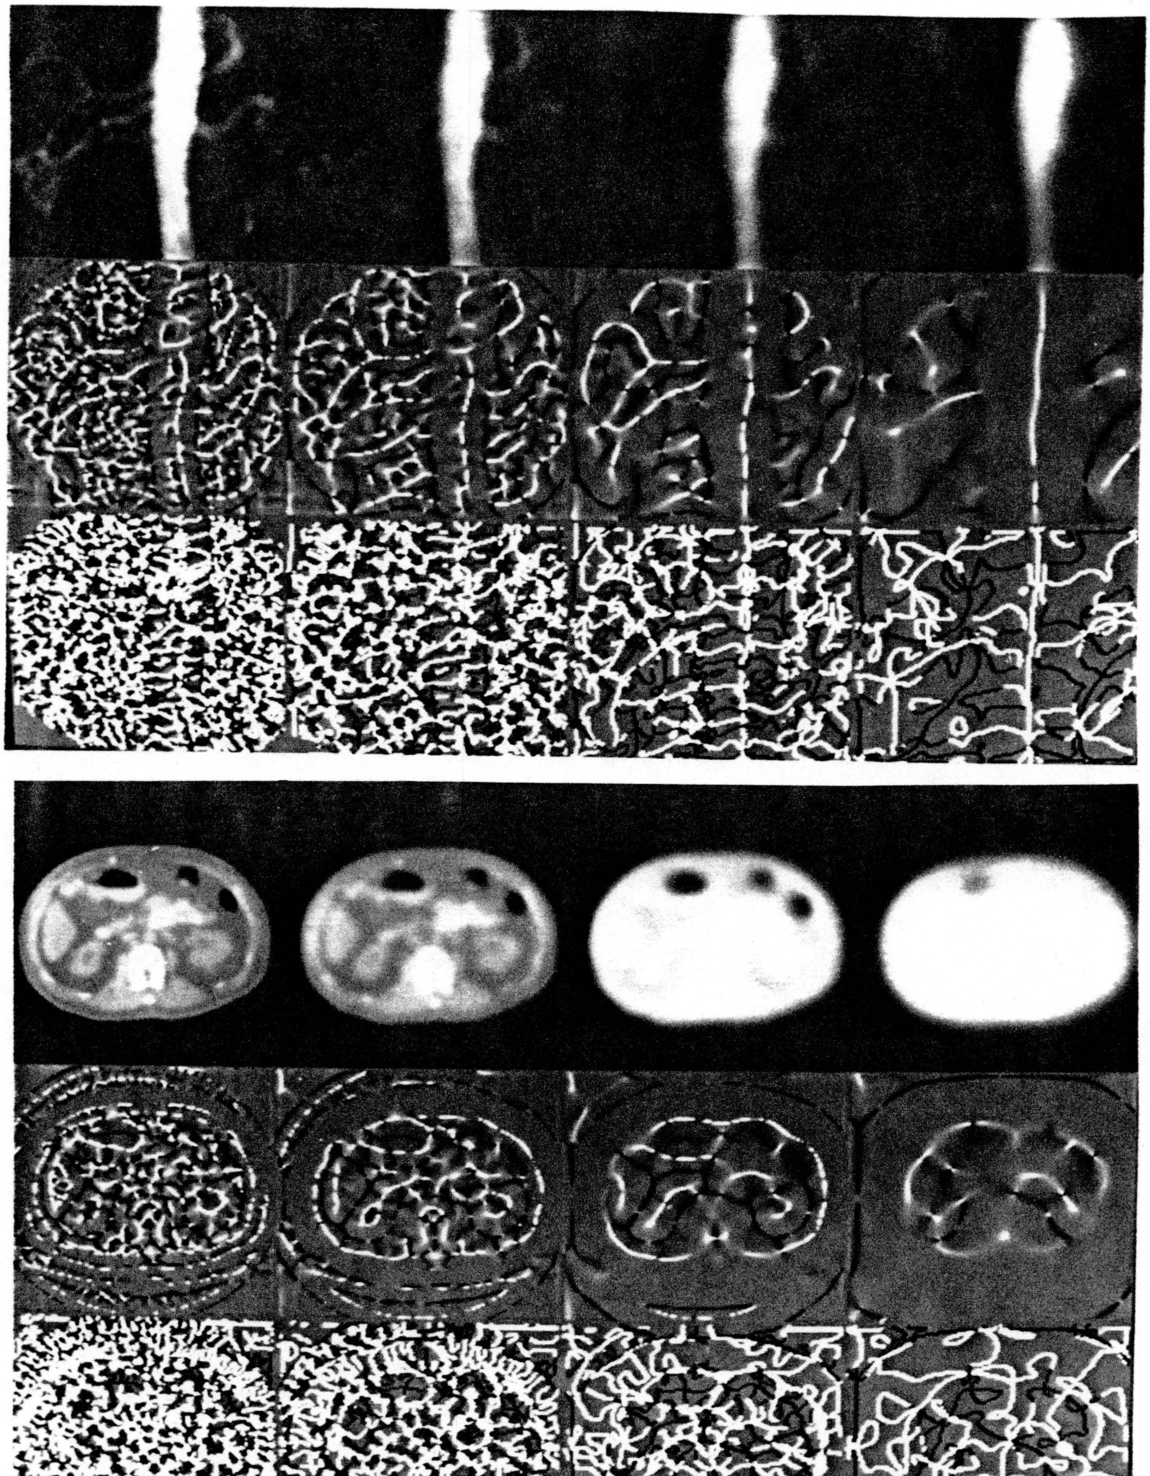

Figure 9: A sequence of blurred digital subtraction angiogram images (top row) with their corresponding level curve curvature (second row) and vertex curve images (third row). Rows four, five and six show analagous sequences for an abdomen CT image.

The level curve curvature K at each image point is computed as $K = v^t\, hessian(I)\, v$, where I is the image and v is the unit vector in the direction of the level curve tangent, $(-\partial I/\partial y, \partial I/\partial x)$. The K values were computed by the multiresolution n-jet approach of Koenderink [1986]. This approach involves computing $\partial^n I/\partial x^m \partial y^{n-m}$ for all n less than some limit, all m $\leq$ n, and all degrees of blurring. From these, many feature values, including level curve curvature, Laplacian value, etc. can be easily computed. Listing useful essential structures that can be computed in this way ought to be the subject of active research.

Vertex curves together with the SAP of the original image seem to satisfy all of the criteria specified for an essential structure.

Region definition. It induces a subdivision into regions that also carry information on intensity variation.

Region sensibleness. The fact that it is based on ridges and courses seems to allow object curving to be followed and objects not to break into unrelated pieces. Like the essential structure of intensity extrema augmented with iso-intensity contour segments, the vertex curve / SAP structure captures the behavior of critical points, but it is more oriented to a whole object rather than one point.

Consistent simplification. Under image blurring no new values of level curve curvature are created, but the topology of the associated vertex curves can change. These changes occur when saddle-extremum pairs annihilate (or form) and also when locally concave or convex regions on the side of hills and valleys are destroyed. By following the smooth evolution of vertex curves, the simplification of SAP structure can be deduced.

Region hierarchy. The hierarchy induced by image simplification involves only a selection among branch sheets of the SAP which are already in the form of a tree (or a forest of trees). Furthermore, the regions they induce are directly described in terms of intensity and spatial shape by the properties of the symmetric axis transform.

Generalization to all dimensions. The method seems extendable to higher dimensions, though details need to be investigated.

The vertex curves / SAP essential structure thus seems quite promising. However, the usefulness of this description and some of its mathematical properties are still under investigation. Furthermore, the dependence on intensity level curves seems unfortunate, and improved means of slicing the image surface need to be developed.

Discussion

Hierarchies generated by the annihilation of essential structures meeting the criteria listed in earlier sections seem to have great promise. Yet they have a few inherent difficulties. First, since they are defined only by image intensities, they cannot be expected always to reflect semantic information. Some means will be necessary either to edit the resulting descriptions to reflect such image understanding or to let such understanding affect the creation of the descriptions.

Second, an important weakness of the hierarchical form of description is its sensitivity to the order of essential structure annihilation. Similar images have descriptions that are made up of qualitatively different regions if their essential structure components annihilate in a different order. The difficulty arises from the association of regions only with annihilating structures. Consider the situation illustrated in Figure 10 when two adjoining essential structure components (e.g., intersecting SAP branches, or intensity extremal paths connected by an iso-intensity path) are both near annihilation. One component annihilates first, defining a region. However, the second component simply becomes part of a larger component, and no region is defined for it, only for the larger component of which it is a part. Despite the fact that a small change in the image might have associated a region with the second component (by changing the order of essential structure annihilation) that region appears nowhere in the description. Since a small change in the image can cause a large change in the description, equivalence classes of images, the essence of a useful object mode, are hard to create. It thus appears that some means needs to be found of forming a description in which candidates for hierarchical components are identified and given weights. A weight would be given a component according to its degree of retention under Gaussian blurring.

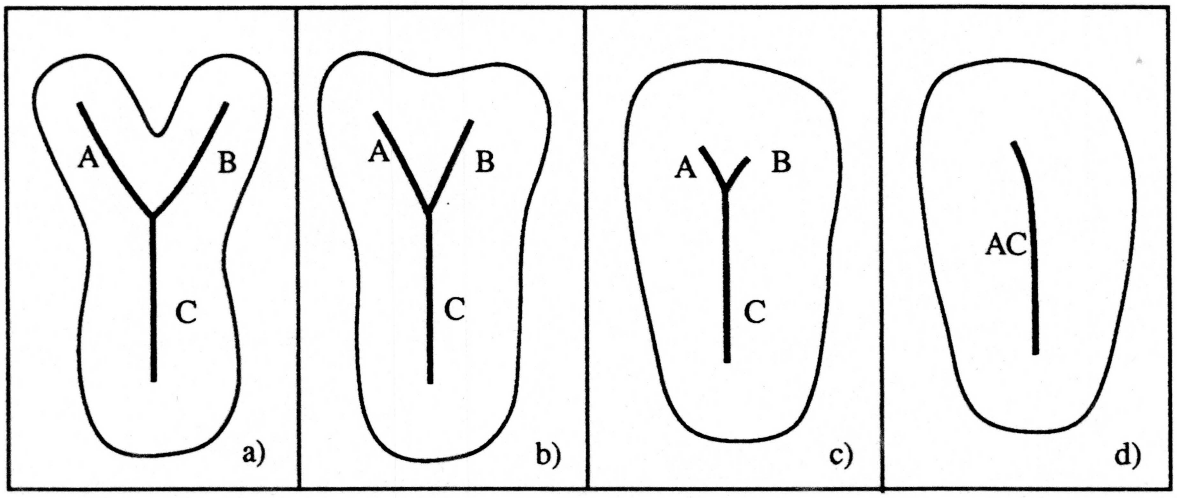

Figure 10: Effect of resolution reduction on figure 10a through stages b, c, and d illustrates the sensitivity of any structural description for such an object. In the case illustrated, region 'B' appears in the description, and region 'A' does not, being part of region 'AC'. If region 'B' is made slightly longer, resolution reduction will cause the branch 'A' to annihilate sooner, so region 'A' will appear in the description and region 'B' will not.

Third, while the approaches we have sketched appear to apply to images of any number of spatial dimensions, it is not yet clear how to extend them to vector-valued images or to images of space and time.

Summary

We have shown that describing images hierarchically by following essential structures to annihilation is attractive if the essential structures satisfy a number of criteria. We have seen that the idea can be applied to a wide range of essential structures. However, the vertex curve / SAP essential structure seems particularly attractive in meeting all of the criteria. Other structures based on geometrical features of the intensity surface might also have these strengths.

This paper has left many open directions for exploration, including how edges should be reflected, how cuts through terrain images should be made, how useful the vertex curve-based descriptions will be, and what other essential structures ought to be investigated. We are confident that such research will lead to the production of useful image descriptions.

Acknowledgements

We are grateful to Jan Koenderink for useful discussions, to Bart ter Haar Romeny for collaboration on the implementation of n-jet-based methods for computing essential structures, and to Sandra Bloomberg for programs for computing symmetric axes. We thank Sharon Laney for help in manuscript preparation and Bo Strain and Karen Curran for photography. The research reported here was carried out under the partial support of NIH grant # R01 CA39060. This manuscript was prepared for presentation at the Fourth Copenhagen Workshop in Computational Vision: Multi-Resolution Representation of Images, November 1987.

References

Blum, H., and Nagel, R.N., "Shape Description using Weighted Symmetric Axis Features", *Pattern Recognition*, Vol. 10, pp. 167-180, 1978.

Brady, M., and Asada, H., *Smoothed Local Symmetries and Their Implementation*, MIT A.I. Memo 757, 1984.

Burt, P.J., and Adelson, E.H., "The Laplacian pyramid as a compact image code", *IEEE Trans. on Communications*, Vol. 31, No. 4, pp. 532-540, 1983.

Crowley, J., and Parker, A., "A Representation for Shape Based on Peaks and Ridges in the Difference of Low-Pass Transform", *IEEE Trans. PAMI*, Vol. 6, No. 2, pp. 156-169, March 1984.

Gauch, J.M., Oliver, W.R., Pizer, S.M, "Multiresolution Shape Descriptions And Their Applications In Medical Imaging", Tech. Report 87-018, University of North Carolina - Chapel Hill, to appear in *Information Processing in Medical Imaging*, 10th Conf. Utrecht, Plenum, 1987.

Ginsburg, A., *Visual Information Processing Based on Spatial Filters Constrained by Biological Data*, Doctoral Dissertation, University of Cambridge, England, 1977.

Grossberg, S., and Mingolla, E., "Neural Dynamics of Perceptual Grouping: Textures, Boundaries, and Emergent Segmentations", *Perception and Psychophysics*, Vol. 38, No. 2, pp. 141-171, 1985

Koenderink, J.J., "The Structure of Images", *Biological Cybernetics*, Vol. 50, pp. 363-370, 1984.

Koenderink, J.J, and van Doorn, A.J., "Dynamic Shape", *Biological Cybernetics*, Vol. 53, pp. 383-396, 1986a.

Koenderink, J.J, and van Doorn, A.J., "Representation of Local Geometry in the Visual System", *Biological Cybernetics*, Vol. 53, November, 1986b.

Koenderink, J.J, "Operational Significance of Receptive Field Assemblies", *Biological Cybernetics*, Vol. 54, January, 1988.

Leyton, M., "A Process Grammar for Shape", to appear in *Artificial Intelligence*, 1987.

Lifshitz, L.M., *Image Segmentation using Global Knowledge and A Priori Information*, Ph.D. Dissertation, UNC Chapel Hill Technical Report 87-012, 1987a.

Lifshitz, L.M., "A Multiresolution Hierarchical Approach to Image Segmentation based on Intensity Extrema", Tech. Report 87-019, University of North Carolina - Chapel Hill, to appear in *Information Processing in Medical Imaging*, 10th Conf. Utrecht, Plenum, 1987.

Pizer, S.M., Oliver, W.R., Bloomberg, S.H., "Hierarchical Shape Description Via The Multiresolution Symmetric Axis Transform", *IEEE Trans. PAMI*, Vol. 9, No. 4, pp. 505-511, 1987a.

Pizer, S.M., Amburn, E.P., Austin, J.D., Cromarti, R., Geselowitz, A., Greer, T., ter Haar Romeny, B., Zimmerman, J.B., Zuiderveld, K., "Adaptive Histogram Equalization and Its Variations", *Computer Vision, Graphics and Image Processing*, Vol. 39, No. 3, pp. 355-368, 1987b.

Richards, W., and Hoffman, D.D, "Codon Constraints on Closed 2D Shapes", *Computer Vision and Image Processing*, Vol. 31, No. 2, pp. 156-177, 1985.

Robson, J., "Frequency Domain Visual Processing", *Physical and Biological Processing of Images*, OJ Braddick and AC Sleigh, ed., Springer-Verlag, 1983.

Rosenfeld, A., *Multiresolution Image Processing and Analysis*, Springer-Verlag, Berlin, 1984.

Wilson, H., and Bergen, J., "A Four-Mechanism Model for Threshold Spatial Vision", *Vision Research*, Vol. 19, pp. 19-32.

Witkin, A., "Scale-Space Filtering", *Proceedings of International Joint Conference on Artificial Intelligence*, 1983.

Yuille, A.L., and Poggio, T., *Scaling Theorems for Zero-Crossings*, MIT A.I. Memo 722, June, 1983.