

COMPUTER RECOGNITION AND ANALYSIS OF FREEZING
CELLS IN NOISY, CLUTTERED IMAGES

Thomas E. Dietz
L. S. Davis
K. R. Diller
J. K. Aggarwal

TR-82-2

July 1982

This research is supported in part by the National Science Foundation
under grant ECS-8021511.

TABLE OF CONTENTS

Chapter 1. Background	1
Chapter 2. System Input/Output	3
Chapter 3. Object Extraction	4
3.1. Edge Detection	5
3.2. Noisy Edge Removal	10
3.3. Shape Detection	12
3.4. Cell Boundary Tracing	16
Chapter 4. Feature Analysis	18
Chapter 5. Conclusion	19
Bibliography	20

Chapter 1

Background

The process of using a computer to analyze the morphology of the cells during the freezing and thawing can be divided into two basic steps: extracting the cell from the image background, and quantifying the two-dimensional structure of the cell. The cell extraction is essentially an image segmentation problem involving partitioning of an image into meaningful regions (in this case, the cell of interest is one region, everything else is another region).

Previous work in this area has been done primarily by researchers interested in automated cell classification. Typically there is reasonably good control over the appearance of the scene being analyzed: the cells are well separated on a featureless background and differential stains are used to give the cells predictable optical densities. Under these conditions, a combination of isodensity contour tracing and heuristic searching can effectively trace the boundaries of the cells of interest [3, 4].

In a problem similar to the current one, yeast cells were analyzed morphologically as they were frozen by Diller and Knox [7]. In these experiments, as in ours, the cells are viable and cannot be

stained. The image background is not featureless; rather, it becomes very cluttered with ice boundaries. Diller and Knox used a boundary-tracking approach to find the cells, but it was fairly sensitive to contrast fluctuations from image to image, and significant preprocessing of the image was required.

We will present and discuss a new system which can segment and analyze images of freezing granulocytes more effectively than has been possible previously. Scene segmentation is presented as a multi-phase process of detecting and enhancing edges, recognizing shapes, and tracing cell boundaries. Feature analysis of the extracted object is discussed.

The computational system presented has five main steps. First, there is the input and representation of the digitized image. Second, an edge detector and enhancement scheme is used to find the edges of objects in the digitized image. Third, we perform shape recognition by using a Hough transform to separate edges belonging to the cell from edges belonging to other objects in the image. Fourth, the outline of the cell is completed using a Heuristic Search boundary tracing algorithm. Finally, the two-dimensional features of the cell boundary are analyzed. The individual steps of the analysis procedure are described as follows.

Chapter 2

System Input/Output

Source images are obtained as 35-mm black-and-white negatives from a motor-driven camera mounted on the cryomicroscope [6]. The negatives are digitized using a bellows close-up attachment coupled to a video camera. The resulting 512 pixel x 512 pixel (a pixel is a picture element) digital images are reduced by 4-pixel averaging to 256 pixel x 256 pixel image frames; each pixel is represented as an eight bit (256) gray-level value (0=very black and 255=very white). Input to the image processing system consists of a sequence of digitized frames of the freeze-thaw sequence under consideration.

Output from the image processing system consists of visual image display as well as more traditional printed information. The image display is viewed on a Grinnell System 512 pixel x 512 pixel color monitor, and is comprised of the original black and white image frame with the traced boundaries of the cell outlined in red. The printed output consists of the results of the feature analysis: cross-sectional area, perimeter, and shape descriptions for the traced cell.

Chapter 3

Object Extraction

The first stage of the analysis is the recognition of a specific cell in the overall image. To this end the image is segmented into two regions; "cell" and "background". This process takes place in four phases, as explained in detail in the following sections: edge detection, noisy edge removal, shape recognition, and cell boundary tracing.

Digital images take up a tremendous amount of computer memory: a 256 x 256 pixel image requires 32K (32768) bytes of memory; associated data arrays needed for processing bring the figure up to 128K bytes, larger than the addressing capabilities of the computer which we are using (a PDP 11/34). Processing the entire image therefore means keeping most of the data on secondary (disk) storage and actively operating on only a small portion of it at any one time. Access times for disk storage media are typically 3 to 4 orders of magnitude slower than for primary memory; therefore processing images in this fashion is very slow. If we focus the system's attention on the neighborhood of the cell, we gain substantial processing speed for two reasons:

1. The size of the window and associated data arrays are small

enough to keep entirely in the core memory, saving data-access time.

2. The amount of computation needed is greatly reduced since the window is much smaller than the full image.

For this reason, the segmentation process for a particular cell is confined to a localized area of the image, centered approximately over the expected location of the cell. The size and location of this window is set by the operator in the first frame of a freezing sequence; in subsequent frames it may be set either by operator interaction or based on the size and location of the cell in the previous image frame, with allowances for possible cell motion and size changes.

3.1 Edge Detection

The process of cell recognition begins with the detection of edges in the image. A classical approach to edge detection is to compute the intensity gradient (the first derivative of image intensity) at each point in the image, and to designate as edges those points with large gradients. The Kirsch operator [8, 14] is used to compute the intensity gradient as follows: for each image point x , let a, b, c, \dots, h represent the image intensities at the points surrounding x , as shown in Figure 1. The gradient, $G(x)$, is then

$$\text{right} = a+b+c$$

$$\text{up} = c+d+e$$

$$\text{left} = e+f+g$$

$$\text{down} = g+h+a$$

$$G(x) = (\max(|\text{right-left}|, |\text{up-down}|))/3 \quad (1)$$

E	D	C
F	X	B
G	H	A

Figure 1: The Kirsch gradient locations around x.

Peaks in the first derivative are easily found by looking for changes in sign of the second derivative. This method is extremely noise-sensitive, however, as the slightest fluctuations in intensity will produce local peaks in the first derivative. This sensitivity to noise can be controlled by using the Laplacian operator to compute the two-dimensional second derivative over arbitrary sized neighborhoods around each pixel in the image [9, 14]. A square-wave approximation to the Laplacian is computed for each image pixel, and the result is scanned vertically and horizontally, marking as edges those pixels where the value of the Laplacian crosses zero (the second derivative changes sign).

If $S_k(x)$ = sum of image intensities in the $k \times k$ neighborhood of x , the Laplacian at x can be approximated by:

$$\text{Lap}(x) = [S_{k_1}/k_1^2] - [(S_{k_2} - S_{k_1}) / (k_2^2 - k_1^2)] \quad (2)$$




Figure 2: The Laplacian neighborhoods around x.

Figure 2 shows the relationship of the inner and outer neighborhoods (k_1, k_2 respectively) around point x . The choice of values for the k_1, k_2 parameters is an important one. Smaller neighborhoods will give finer detail on actual object boundaries, but also have a higher sensitivity to noise, with resultant false edges. Larger neighborhoods will distort the object boundaries and have a greater loss of fine detail. Figure 3 shows an image of frozen granulocytes; a 50 pixel x 50 pixel window containing a granulocyte is outlined in black. Figure 4 is an expanded view of the window image; the Laplacian zero-crossing edges (neighborhoods $k_1=3, k_2=7$) for the window are shown in Figure 5.

These processed images contain strictly binary edge-markers, only indicating either edge or no-edge. They lack valuable information about the edge, such as its strength and direction. The strength of an edge is simply the intensity gradient at the edge; the direction of an edge is the direction of maximum positive intensity change at the edge.




Figure 3: An image of frozen granulocytes




Figure 4: A 50 x 50 pixel window of the image.

In order to retain this useful information, we convert the binary edges, detected by the Laplacian zero-crossing method, to directed edges composed of a strength, direction pair. The strength (intensity gradient) is computed using the Kirsch operator (equation (1)); at the same time, the edge direction is determined by finding the brightest (i.e., highest gray level) 3-pixel bar of the four shown in equation (1), and associating it with one of the four directions shown in Figure 6.




Figure 5: The Laplacian zero-crossing edges.




Figure 6: The four defined edge directions.

For example, if 'right' is the brightest bar, the edge direction is 1; if 'up' is the brightest bar, the edge direction is 2, and so forth.

3.2 Noisy Edge Removal

The next step in extracting the cell from the background is to clean up the initial edges. There are two types of unwanted edges as typified by:

1. changes in intensity in the image that are detected as edges by the Laplacian zero-crossing algorithm, but which are not readily visible to the eye and do not represent "real" edges.
2. real edges that are part of objects other than the cell of interest (e.g., ice boundaries).

Edges of the first type are weak edges, that is, they have a low intensity gradient and they can be removed by an intensity gradient thresholding scheme. Because we desire a system that is insensitive to contrast fluctuations (from frame to frame), we need a relative rather than an absolute definition of "weak"; for each image we must choose a gradient threshold value below which edges are "weak". To determine the relative threshold level, the edges are histogrammed by the magnitude of their intensity gradient. A threshold value can be derived from this histogram based on a specification of what fraction of edges to reject; edge points with an intensity gradient less than the threshold value are discarded. The relative rejection parameter can be specified a priori; we have achieved good results with 0.6 (discard 60% of the initial edges). Figure 7 shows the intensity gradient histogram for the edges in Figure 5; T is the minimum gradient below which at least 60% of the

edges lie (the discrete 60th percentile). The edges remaining after thresholding at 60 percentile are shown in Figure 8.

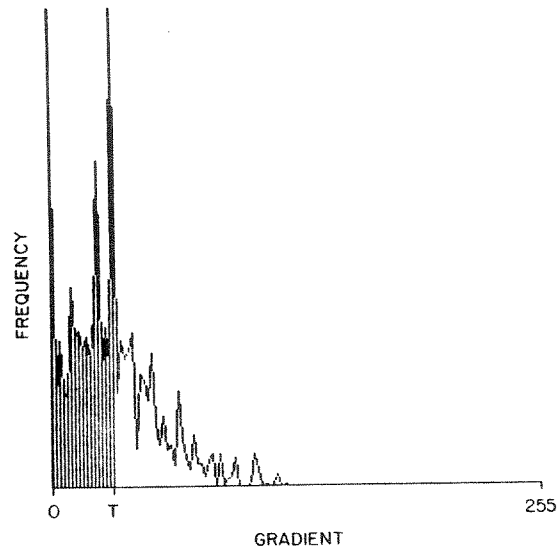


Figure 7: The intensity gradient histogram.




Figure 8: The edges after thresholding at T.

3.3 Shape Detection

In order to separate edges that belong to the cell boundary from edges that belong to other objects, techniques considerably more sophisticated than thresholding are required. We have found a modified version of the Hough transform [2, 5] to be useful in recognizing shapes in the enhanced (post gradient thresholding) edge image.

The Hough transform [1, 12, 13] develops a two-dimensional space (the Hough space) and labels each pixel with a value which is a measure of the likelihood of the cell centroid being located at the corresponding pixel in the image. The Hough space can be thought of as a two-dimensional histogram corresponding to the two-dimensional image window; each point h in the Hough space contains a count of the number of edges which, were they to lie on a cell boundary, could consider h to be the cell centroid. Not knowing which edges correspond to cell boundaries and which correspond to boundaries of other objects, we assume that each edge in the binary image could lie on the boundary of the cell. Then, for each edge all of the Hough space points that correspond to possible centroids for that edge are incremented by one. For example, if we are looking for a circle of radius r , any given edge x might lie on such a circle and would have its center r pixels away. In this case we would increment all of the Hough space points corresponding to image points in a circle of radius r (the "possible-center" line) about edge-point x ; each of these image points now has one more "vote" for its selection as the center of the circle (Figure 9). Edges which

actually do lie on a circle will have possible-center lines that are coincident at the actual circle center, building a peak in the Hough space (Figure 10). After all of the edges have been considered, the largest peak in the Hough space corresponds to the most-likely centroid of the object we are looking for: it has the most "votes" from edges that could be part of the object if its centroid were at that location.




Figure 9: "Possible-center" line for edge x.




Figure 10: Coincident "possible-center" lines.

Consider now a disc of uniform intensity having boundary edges that form a circle. The Hough transform technique can be refined further to take advantage of additional knowledge about the disc:

whether it is bright or dark against the background. If the disc is bright, then the intensity gradient at the boundary of the disc is highest toward the interior, and the direction of a boundary edge will point toward the half-plane that the center of the disc lies in. If the disc is dark, then the direction of a boundary edge will point away from the half-plane that the disc center lies in. We can now restrict the possible-center lines to lie in half-planes, resulting in possible-center semicircles. This constraint has two effects: it doubles the speed of the Hough transform algorithm, and it improves the accuracy of the Hough transform by eliminating the contribution of edges to centers for which they could not belong.

In the present application we are looking for cell boundaries which in general are roughly circular, but may have radii which change as a function of position. In order to allow for this geometric variation and still be able to identify the true boundaries, we have "blurred" the semicircular possible-center lines mentioned above into half-annuli possible-centroid masks, with inner and outer radii matching the expected minimum and maximum cell boundary radii.




Figure 11: Possible-centroid masks.

Figure 11 shows the half-annuli masks used for bright cells; r_1 and r_2 are the inner and outer radii, respectively; the arrows represent the direction of the edge over which the mask is centered. Using these half-annuli, the Hough transform will find round or deformed cells, as long as their radii lie within the limits of the shape annulus. The minimum and maximum radius limits of the cell, and whether it is bright or dark may be supplied by the operator, or determined from the previous frame in the sequence. Figure 12 shows the Hough space that is derived from the edges in Figure 8 by specifying a minimum radius of 6 pixels and a maximum radius of 10 pixels.




Figure 12: The Hough space.

3.4 Cell Boundary Tracing

The highest peak in the Hough space represents the most-likely cell centroid location. Using this probable centroid location, we can determine which edges contributed to the Hough space peak and develop a set of rough cell boundary edges (Figure 13). There still remain two types of errors which may be associated with this rough outline: discontinuities in the boundary that need to be closed, and edge points that do not belong to the cell boundary.




Figure 13: Rough outline of the cell.

A technique known as heuristic search [4, 11, 10] may be used to determine the final boundary of the cell from the rough outline. Heuristic search is a graph-searching technique that considers the set of possible boundary points as a directed graph and finds a path through the graph (develops a boundary) which minimizes a cost function. Terms in the cost function have minimal values associated with desirable boundary attributes, and weight the importance of one type of attribute

over another. Our boundary cost function includes terms which favor (in decreasing importance) boundaries which:

- follow edges in the rough outline.
- follow constant-radius arcs.
- follow isodensity contours in the image.

Heuristic information ("serving to aid discovery") is used during the search to estimate which path shows the most promise as being the minimal-cost path; this constrains the search for a cell boundary and improves the computational efficiency of the boundary-tracing process. The heuristic search technique closes discontinuities in the rough outline, rejects spurious edge points remaining in the rough outline, and yields a final best-boundary for the cell (Figure 14).




Figure 14: Final cell boundary.

Chapter 4

Feature Analysis

The final stage of the computational process is the analysis of the cell's morphological features. From the cell boundary we can determine the cross-sectional area, perimeter, and shape descriptions for the cell. If we assume a specific three dimensional morphology we can determine the cell volume as well.

The final boundary of the cell is an outline: it lies adjacent to the exterior of the cell. The perimeter is thus computed as the count of pixels immediately interior to the traced boundary; the cross-sectional area is the number of pixels enclosed by the traced boundary.

By tabulating the cell features for each frame in the sequence, we can develop a history of morphology features for the cell; from these we can compute additional measures such as total changes in perimeter, area, and volume, maximum boundary deformity, optical properties such as texture and contrast, and other feature changes germane to analysis and interpretation of the freezing process.

Chapter 5

Conclusion

The problem of analyzing the two-dimensional morphology of freezing cells has been shown to be one of extracting the cell boundary from the image, and analyzing the features of the cell based on the boundary.

In very noisy and/or cluttered images, simple thresholding and isodensity contour-tracing techniques fail, and even basic Hough transform techniques can fail to give satisfactory performance in the extraction problem. The modified Hough transform scheme presented here, together with the heuristic search boundary completion technique, has proved to be quite successful at finding cell boundaries in images very cluttered with ice boundaries.

Bibliography

- [1] Ballard, D.H., and Sklansky, J.
A ladder-structured decision tree for recognizing tumors in chest radiographs.
IEEE Trans. Computers C-25:503-513, 1976.
- [2] Ballard, D.H.
Generalizing the Hough transform to detecting arbitrary shapes.
Pattern Recognition 13:111-122, 1981.
- [3] Brenner, J.F., Gelsema, E.S., Necheles, T.F., Neurath, P.W., Selles, W.D., and Vastola, E.
Automated classification of normal and abnormal leukocytes.
J. Histochem. Cytochem. 22:697-706, 1974.
- [4] Brenner, J.F., Lester, J.M., and Selles, W.D.
Scene segmentation in automated histopathology: techniques evolved from cytology automation.
Pattern Recognition 13:65-77, 1981.
- [5] Davis, L.S.
Computing the spatial structure of cellular textures.
Computer Graphics Image Processing II:111-122, 1979.
- [6] Diller, K.R., and Cravalho, E.G.
A cryomicroscope for the study of freezing and thawing processes in biological cells.
Cryobiology 7:191-199, 1970.
- [7] Diller, K.R., and Knox, J.M.
Identification and tracking of blurred boundaries in cluttered pictures.
In Pattern Recognition Image Processing 81, pages 137-139. IEEE, 1981.
- [8] Kirsch, R.
Computer determination of the constituent structure of biological images.
Computers Biomedical Research 4:315-328, 1971.

- [9] Marr, D., and Hildreth, E.
Theory of edge detection.
Technical Report AI memo no. 518, MIT AI Laboratory, 1979.
- [10] Nilsson, N.J.
Problem-Solving Methods in Artificial Intelligence.
McGraw-Hill Book Co., New York, 1971, pages 43-71.
- [11] Nilsson, N.J.
Principles of Artificial Intelligence.
Tioga Publishing Co., Palo Alto, CA, 1980, pages 61-88.
- [12] Shapiro, S.D.
Properties of the transforms for the detection of curves in noisy pictures.
Computer Graphics Image Processing 8:219-236, 1978.
- [13] Sklansky, J.
On the Hough technique for curve detection.
IEEE Trans. Computers C-27:923-926, 1978.
- [14] Wechsler, H., and Sklansky, J.
Automatic detection of ribs in chest radiographs.
Pattern Recognition 9:21-30, 1977.