AD-A043 410 MARYLAND UNIV COLLEGE PARK COMPUTER SCIENCE CENTER F/G 9/2
THRESHOLD EVALUATION TECHNIQUES.(U)
APR 77 J S WESZKA F44620-72-C-0062
UNCLASSIFIED TR-531 AFOSR-TR-77-0980 NL

1 OF 2
AD A043410

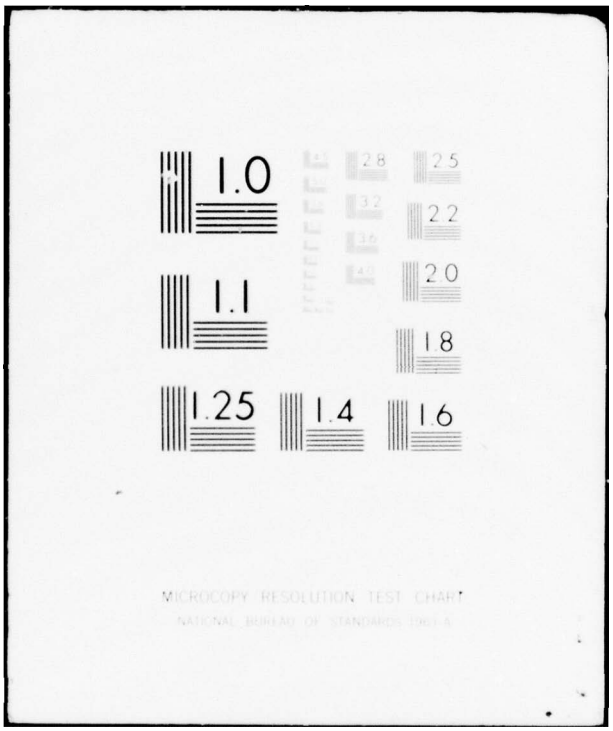
1.0
1.1
1.25
1.4
1.6
2.8
2.5
3.2
2.2
3.6
2.0
4.0
1.8
MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

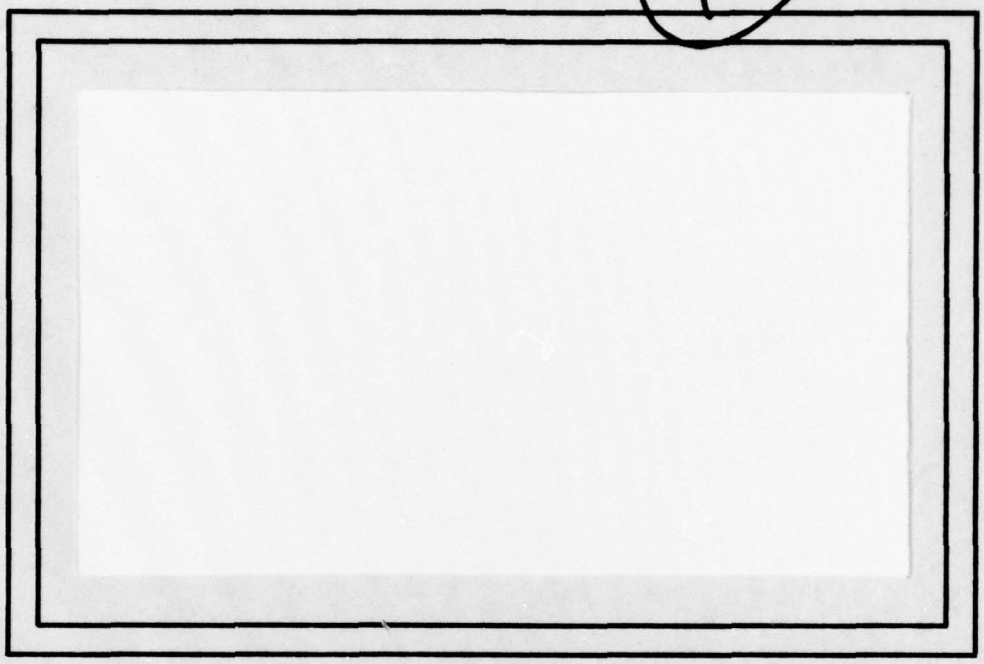

COMPUTER SCIENCE
TECHNICAL REPORT SERIES

UNIVERSITY OF MARYLAND
COLLEGE PARK, MARYLAND
20742

(18) (19) AFOSR-TR-77-0980
(14) TR-531
(15) F44620-72C-0062
(11) April 1977
(12) 108p.

(6) THRESHOLD EVALUATION TECHNIQUES.

(10) Joan S. Weszka
Computer Science Center
University of Maryland
College Park, MD 20742

(9) Technical rept.,

ABSTRACT

Threshold selection techniques have been used as a basic tool in image segmentation, but little work has been done on the problem of evaluating a threshold of an image. This paper addresses the problem of threshold evaluation and proposes two methods for measuring the "goodness" of a thresholded image, one based on a busyness criterion and the other based on a discrepancy or error criterion. These evaluation techniques are applied to both synthetic and real images and are shown to be useful in facilitating threshold selection. In fact, both methods usually result in similar or identical thresholds which yield good segmentations of the images.

The support of the Directorate of Mathematical and Information Sciences, U. S. Air Force Office of Scientific Research, under Contract F44620-72C-0062, is gratefully acknowledged.

ACKNOWLEDGEMENTS

I would like to express my sincere appreciation to Professor Azriel Rosenfeld, my advisor, whose help was invaluable throughout the writing of this paper. I would also like to thank Professor Laveen Kanal and Professor Ashok Agrawala for their comments on the paper. Ms. Shelly Rowe for typing the manuscript, and Ms. Melodie Wolfe for her assistance in preparing the figures. Finally, I am especially grateful to Larry, my spouse, for his patience and encouragement.

Table of Contents

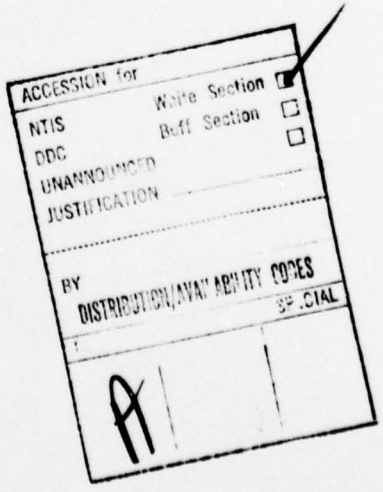

List of Tables

List of Figures

CHAPTER I

INTRODUCTION

The use of thresholding as a tool in image segmentation has been extensively studied, and a variety of techniques have been proposed for automatic threshold selection. A review of these techniques is presented in Appendix A.

This paper addresses the problem of evaluating the "goodness" of a threshold applied to an image and compares two criteria for evaluating thresholds. Such threshold evaluation criteria can be used to evaluate threshold selection techniques, by evaluating the thresholds obtained using these techniques. Of course, a threshold evaluation function can also be used for threshold selection, by applying the function to a range of thresholdings of an image and choosing the one yielding the best evaluation.

In using a threshold evaluation function for threshold selection, we should ideally select the threshold corresponding to the minimum cost. As we will see in later sections, minimizing an evaluation function does not always result in reasonable thresholds. However, even in such cases, the use of an evaluation function does provide a method of threshold selection that makes reasonable thresholds easier to detect.

We will examine two methods of evaluating and selecting thresholds, one based on a busyness criterion and the other based on a discrepancy or error criterion. By applying these methods to both synthetic and real images

we will find that they usually result in similar or identical thresholds which yield good segmentations of the images. A comparison of the computational cost of the two methods is also given.

Throughout the paper we will, whenever possible, relate our work on threshold evaluation and selection to image models in order to gain a better understanding of the meaning of a good image segmentation.

CHAPTER II

THRESHOLD EVALUATION CRITERIA

To evaluate a given thresholding of an image, one can take an approach analogous to that used by Martelli and Montanari [1] for evaluating smoothings of images. They defined a cost function which was a weighted combination of the following components:

1) A discrepancy measure, based on the difference between the original and smoothed pictures. The measure proposed was the sum of the squared differences between gray levels of corresponding points in the original and smoothed picture.

2) A busyness or roughness measure based on the computation of a local property on the smoothed picture. For example, the sum of absolute values of a difference operator such as the gradient or Laplacian could be used.

Both of these components of the goodness of a smoothed picture are potentially relevant to thresold evaluation. The structure of a thresholded image should not differ from that of the original image with respect to the number, sizes and shapes of objects. Thus it should be possible to formulate a criterion of goodness for thresholding analogous to the discrepancy criterion for smoothing. The busyness criterion is certainly applicable to tests of threshold goodness if we adopt the point of view that a good threshold is one which minimizes the amount of noise

or roughness in the resulting image. We will now examine the applicability of discrepancy and busyness criteria to threshold evaluation and discuss methods of computing these measures on a thresholded image.

We will not consider how the discrepancy and busyness measures might be combined into a single evaluation function. In any case, as we shall see, the two measures tend to give the same or similar thresholds when they are used for threshold selection.

2.1 Discrepancy Measures

It is not obvious how the definition of the discrepancy measured used to evaluate smoothing could be extended to threshold evaluation. One cannot simply measure the difference between the thresholded and unthresholded images, since the output levels used for display of the thresholded image are usually chosen arbitrarily, and this difference is sensitive to their choice. For example, if the output levels were fixed at 0 and 63, then for an image with gray levels in the range 0 to 9, a threshold of 10 or greater would always minimize discrepancy by producing a thresholded picture consisting of all 0's. Clearly, this would not be desirable. A careful choice of output levels is necessary in order to use this measure to arrive at meaningful thresholds. Since a threshold evaluation criterion based on the output levels used for display of the thresholded image does not correspond to any of our intuitive notions of how to choose a good threshold, we will not pursue this topic here. A further discussion of this type of discrepancy measure is contained in Appendix B.

An alternative approach to measuring the "discrepancy" of a thresholded image is in terms of classification error [2]. Let us suppose that the image consists of dark (i.e., high gray level) objects on a light background, where the distribution of gray levels in the objects is normal with mean μ_1 and standard deviation σ_1, and that the objects occupy

fraction θ of the picture points. Then the conditional probability density of the gray levels of object points is

$$p(z|\omega_1) = \frac{1}{\sqrt{2\pi}\ \sigma_1} \exp[-(z-\mu_1)^2/2\sigma_1^2]$$

Similarly, if the distribution of gray levels in the background is normal with mean μ_2 and standard deviation σ_2, then we have

$$p(z|\omega_2) = \frac{1}{\sqrt{2\pi}\ \sigma_2} \exp[-(z-\mu_2)^2/2\sigma_2^2]$$

Thus, the picture has overall gray level probability density

$$\theta p(z|\omega_1) + (1-\theta)p(z|\omega_2)$$

We can now take as a criterion of threshold goodness the number of points misclassified by the given threshold. Suppose that a threshold t is chosen such that all points with gray level $x > t$ are classified as object points. Then the probability of misclassifying an object point as a background point is

$$\int_{-\infty}^{t} p(z|\omega_1)\,dz$$

Similarly, the probability of misclassifying a background point as an object point is

$$\int_{t}^{\infty} p(z|\omega_2)\,dz$$

The overall misclassification probability is then

$$P(\text{error}|t) = \theta\left[\int_{-\infty}^{t} p(z|\omega_1)\,dz\right] + (1-\theta)\left[\int_{t}^{\infty} p(z|\omega_2)\,dz\right] \qquad (1)$$

Minimizing (1) enables us to obtain a minimimum error threshold for the image.

A closed form for t_m, the minimum error threshold can be obtained by setting $\frac{dP(\text{error}|t)}{dz}$ to zero and solving for t. Using this equation requires knowing the parameters and a priori probabilities of the object and background distributions. These parameters can be obtained by fitting normal curves to the gray level histogram. The algorithm we will use to do the fitting [3] is iterative, with each successive iteration approaching a better approximation according to some likelihood function. This procedure will converge toward a locally optimal fit but is not guaranteed to find a global optimum. Thus, the parameters obtained are dependent on the initial inputs for these parameters. In addition, it may be difficult to obtain good fits since the histogram may be badly truncated at either end of the gray scale.

2.2 Busyness Measures

A threshold evaluation criterion based on the amount of busyness, noise or roughness in the thresholded image seems to capture many of our intuitive notions of what constitutes a good threshold. Isolated noise points both in the object and in the background of the thresholded image are usually undesirable, as are larger holes in objects and clusters of noise points in the background, since the objects and background should ideally occupy disjoint gray level ranges and should have simple shapes. In short, we would like our thresholded images to look smooth rather than busy. We propose to embody these notions of a good threshold in an automatic threshold selection procedure based on a busyness measure computed on the thresholded image.

One method of computing the amount of busyness corresponding to a given thresholding of an image is based on the gray level co-occurrence matrix [4]. Let $M_{(1,0)}$, $M_{(0,1)}$, $M_{(1,1)}$ and $M_{(1,-1)}$ be the joint probability matrices for gray levels occurring in relative positions (1,0), (0,1), (1,1) and (1,-1), respectively. Represent by M the average of these four matrices, each of which is symmetric about the main diagonal. Suppose that a given threshold t maps all gray levels greater than t into the object and all other levels into the background. This mapping defines a partition of matrix M into three non-overlapping areas:

1) Matrix elements representing co-occurrences of

gray levels in the object, i.e., those M(i,j)s such that $i > t$ and $j > t$ (shaded area A of Fig. 1).

2) Matrix elements representing co-occurrences of gray levels in the background, i.e., those M(i,j)s such that $i \leq t$ and $j \leq t$ (shaded area B of Fig. 1).

3) Matrix elements representing co-occurrences of object gray levels with background gray levels, i.e., M(i,j)s such that $i \leq t$ and $j > t$ (shaded area C of Fig. 1)or $i > t$ and $j \leq t$ (shaded area C').

Note that since matrix M is symmetric about the main diagonal we can restrict our consideration to those matrix elements M(i,j) such that $i \leq j$.

Given a threshold t of an image, the measure of busyness C(t) which will be used throughout this paper is computed by summing those entries of the co-occurrence matrix representing the percentage of object-background adjacencies (i.e., the entries in area C of Fig. 1). If C(t) is relatively high for a given threshold we would expect the thresholded image to contain a large number of noise points and/or jagged edges. Conversely, a relatively low C(t) would indicate that the threshold chosen results in a smooth picture. C(t) will be zero if all gray levels are mapped into the same output level. Since this is obviously not desirable, the minimization of C(t) must

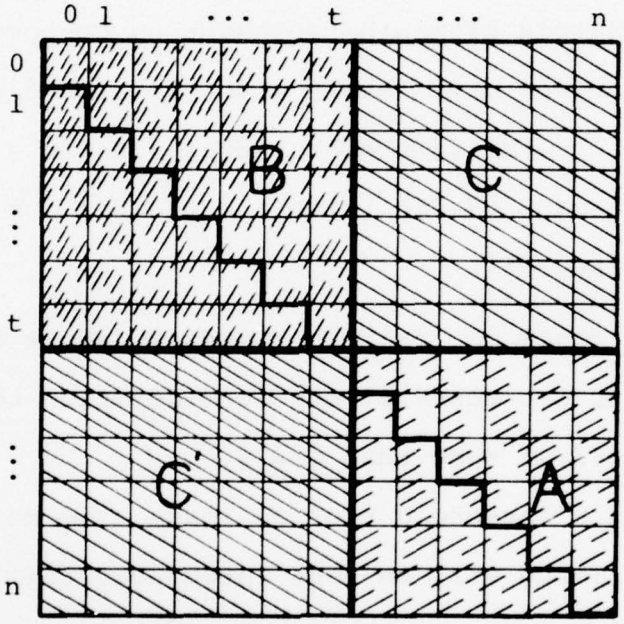

Fig. 1. Partition of an nxn co-occurrence matrix for a given threshold t, where n is the number of gray levels in the image.

be done subject to constraints on the allowable range of the threshold. These constraints will be discussed further in Section 2.3.

An alternative measure of the amount of busyness in a thresholded image could be computed by applying a 4 or 8 neighbor Laplacian to that image and using the average (or sum) of the absolute Laplacian values as a busyness measure. Note that salt and pepper noise points and jagged edges both contribute high absolute Laplacian values. The Laplacian-based busyness measure and the co-occurrence-based measure C(t) are basically equivalent. This can be shown by considering the contributions that different neighborhood configurations of 1's and 0's (in the thresholded image) make to both the Laplacian and co-occurrence matrices. The Laplacian value at a point in a thresholded image is proportional to the number of neighbors which differ from that point. For example, the value of the 4-neighbor Laplacian at an isolated noise point is 4. Similarly, if we compute a co-occurrence matrix based on pairs of points which are horizontally and vertically adjacent, this too examines the four neighbors of each point, and in region C, it counts the number of these neighbors which differ from the point in the thresholded image. The co-occurrence matrix is usually normalized so that the entries sum to one. Analogously, the average rather than the sum of the Laplacian values may be used. Thus the two methods of measuring busyness may yield results which differ by a constant factor.

The co-occurrence-based busyness measure has been chosen in this study for reasons of computational simplicity. Busyness measures for a range of thresholds of an image can be efficiently computed from its gray level co-occurrence matrix. The matrix can be computed in O(r) operations where r is proportional to the picture area. For successive thresholds, C(t), the busyness measure for a threshold t, can be efficiently computed by updating the previously computed busyness measure C(t-1) with the appropriate row and column entires as follows:

$$C(t) = C(t-1) - \sum_{i=1}^{t-1} M(i,t) + \sum_{j=t+1}^{n} M(t,j)$$

where n is the number of gray levels in the image (and the dimension of M). The range of thresholds over which we would compute the busyness measure could be limited to a narrow range by examining the gray level histogram.

A potential advantage of using the Laplacian measure is that the Laplacian values correspond to various neighborhood configurations -- for example, isolated noise points, line ends, lines, and edges, which could be lighter or darker than their surrounds. This information could be valuable in evaluating a given threshold of an image with respect to various types of busyness in the thresholded image, which cannot be distinguished by analyzing the co-occurrence matrix. Use of more refined measures of busyness, possibly based on the Laplacian, should be considered in future studies.

2.3 Threshold Selection Criteria

We have discussed in Sections 2.1 and 2.2 measures for evaluating the "goodness" of a thresholded image based on discrepancy and busyness criteria. Using either of these approaches we would minimize the associated evaluation function in order to select the best threshold. In this section we will show that constraints must be imposed on the allowable threshold range in order to guarantee that the threshold segments the image into two regions (in accordance with our image model).

In Section 2.1 we defined the discrepancy measure as the sum of the probabilities of the two types of misclassification errors: object point misclassifications and background point misclassifications. We have made the assumption that both of these types of errors are equally costly, and that a zero cost is assigned to a correct classification. Our decision rule is:

Decide ω_1 if $f(x) > t$;

Otherwise decide ω_2

where $f(x)$ is the gray level of a point x, and t is the threshold. To minimize the average probability of error we would choose t so that our decision rule always chooses the class ω_i whose a posteriori probability $P(\omega_i|x)$ is maximum. This decision rule is Bayesian [5].

To find the value t which minimizes the average probability of error we can differentiate the expression for $P(\text{error}|t)$ which was given previously in (1):

$$\theta\left[\int_{-\infty}^{t} p(z|\omega_1)dz\right] + (1-\theta)\left[\int_{t}^{\infty} p(z|\omega_2)dz\right]$$

and set the result equal to zero to obtain the following equation:

$$\theta p(t|\omega_1) = (1-\theta)\ p(t|\omega_2) \qquad (2)$$

Since pairs of normal curves whose variances are unequal intersect at two points, satisfying (2) does not result in a unique choice of t. In general, one of the intersection points (t_1) minimizes P(error|t) and the other (t_2) maximizes this probability [6]. Fig. 2 shows examples of intersections of pairs of normal curves. In Fig. 2a, the minimum error threshold t_1 lies between the object and background means. In Fig. 2b, both t_1 and t_2 lie outside the range between the means. In this latter case, if we threshold at t_1 (to obtain minimum error), the thresholded image would consist of almost all 0's. Since we have made the assumption that our images consist of two types of regions (ω_1 and ω_2), any threshold resulting in a single region is not satisfactory.

We will therefore impose on our minimum error threshold evaluation function the constraint that t lie between the object and background means, i.e., we will select t such that (a) $\mu_2 \leq t \leq \mu_1$ and (b) P(error|t) is minimal. Any threshold satisfying (a) and (b) will be referred to as a discrepancy threshold. The restricted range for t ensures that, if the means of the object and background distributions are equal to their medians (as is the case, e.g., for

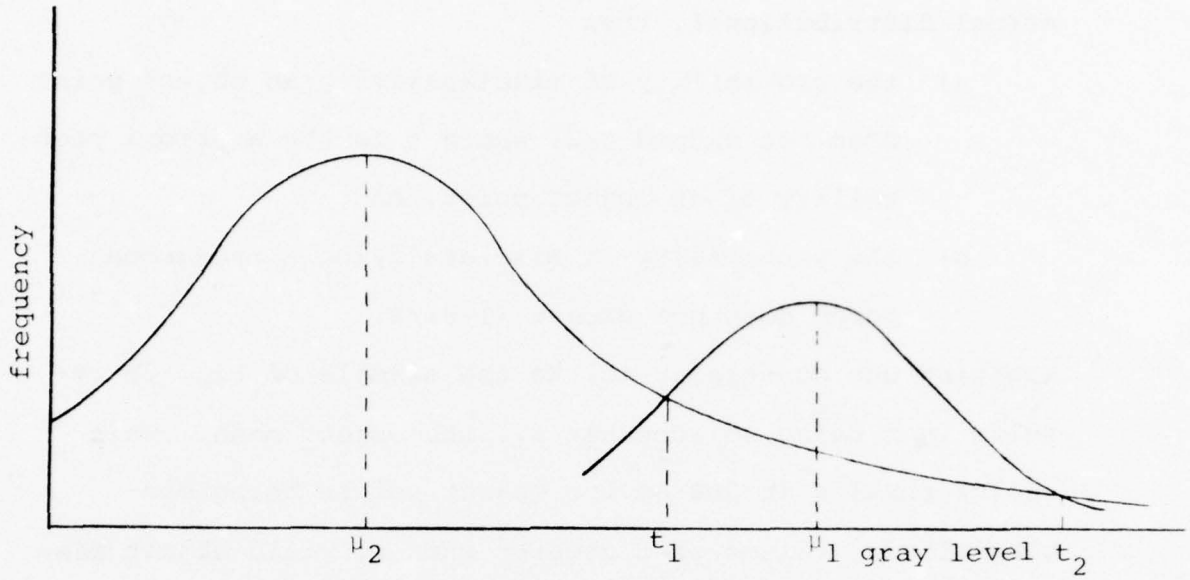

a

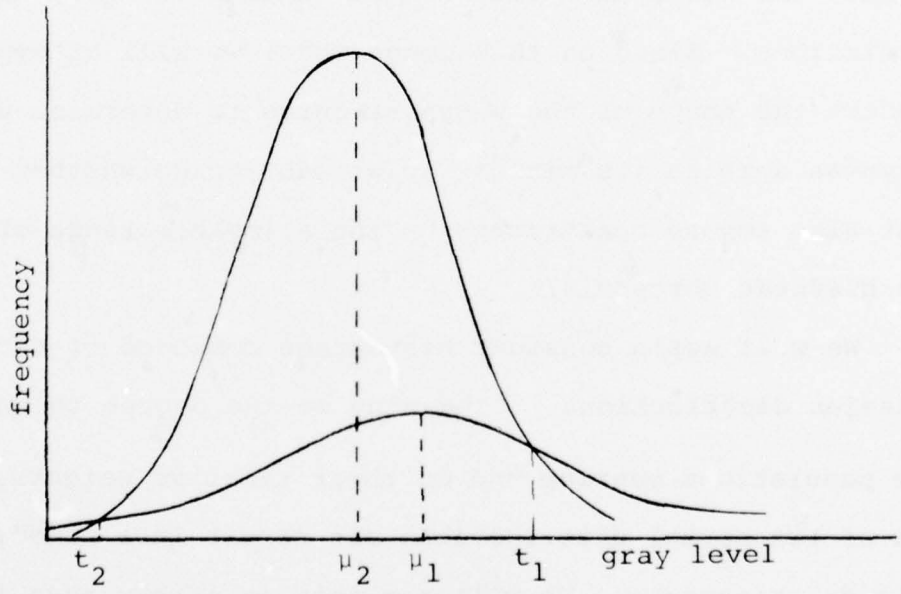

b

Fig. 2. Minimum error thresholds (t_1) for pairs of normal distributions.

normal distributions), then

a) the probability of misclassifying an object point does not exceed $\theta/2$, where θ is the a priori probability of an object point, and

b) the probability of misclassifying a background point does not exceed $(1-\theta)/2$.

Applying our constraint (a) to the example of Fig. 2b results in t being selected at μ_1, the object mean. This choice results in 50% of the object points being misclassified. Values of t greater than μ_1 yield object misclassification rates greater than 50%.

In our previous discussion of busyness (Section 2.2) we made the assumption that images contain two gray level populations. Based on this image model we will attempt to predict the shape of the busyness curve to determine where busyness attains its minima so we can decide whether we must also impose constraints on the allowable range of the busyness threshold.

We will again consider histograms composed of two Gaussian distributions. Depending on the degree to which the populations overlap and on their relative weights, the sum of two normal distributions can result in a bimodal or unimodal histogram. We will see that in either case it is reasonable to suppose that the busyness curve will resemble the image's gray level histogram.

Consider first the case of a bimodal histogram. If we threshold at the mode of the background distribution we would expect busyness to attain a relative maximum. This

occurs because the probability of pairs of background points occurring on opposite sides of t is greatest at the mode which is the mean of the background population in our image model. If we choose t at the mode, then p_b, the probability that a background gray level is less than t, is .5, and $p_b(1-p_b)$, the probability that pairs of background points are on opposite sides of t, is maximum, so busyness should also be maximum. Analogously, at the peak of the object distribution, busyness should also attain a relative maximum. Similarly, in the valley on the gray level histogram, the busyness curve should also have a valley. This occurs for the following reason: We are assuming that the gray levels in the valley lie on or near the borders of objects, and that their frequency of occurrence is low. Thus, when t is chosen in the valley, the busyness measure varies in proportion to the perimeter of the object which grows or shrinks in size depending on where in the valley t lies. (Note that values of t chosen to the left of the valley bottom should result in larger objects than values of t to the right of the valley bottom.) Few non-border points contribute to busyness since for thresholds in the valley the object and background interiors are assumed to be relatively noise-free. Hence, thresholds chosen in the valley result in much smaller busyness values than those which cause the object or background interior to break up (e.g., when t is chosen at either mode).

We can also view the busyness curve as a sum of two

types of busyness, contributed by the object and by the background, respectively. The curve corresponding to busyness in the object resembles the object histogram, and the curve corresponding to busyness in the background resembles the background histogram. So, for a bimodal histogram we would expect a bimodal busyness curve achieving a minimum between the peaks.

Since we know that busyness is zero for choices of t at the extreme ends of the grayscale (resulting in one case in all 1's, and the other in all 0's), the threshold at the valley bottom on a bimodal busyness curve does not correspond to an absolute minimization of the busyness function. However, if we constrain t to lie between the means (as we did for discrepancy), and then minimize busyness, the valley bottom would be chosen as our busyness threshold. The rationale for this constraint on t is basically the same as previously discussed. We are not willing to misclassify more than half of the object points as background points or vice versa.

We will now discuss the choice of a busyness threshold for a unimodal busyness curve. When the image contains a small object on a large background (θ close to 0), and the means of these regions are not too far apart, we may expect situations such as shown in Fig. 3. Here the histogram contains a peak corresponding to the background, but the object does not produce a peak, since the rate of falloff of the background population outweighs the rise correspond-

a

b

Fig. 3. Histograms containing peaks and shoulders.

ing to the object population. However, the presence of the object will at least produce a change in the slope of the histogram; thus, in Fig. 3, this slope flattens out (or at least becomes less steeply negative) as the object mean is approached. We will refer to such a slope change as a shoulder on the histogram. The busyness curve should also display such a shoulder, for the reason discussed previously. We will assume that the shoulder on a unimodal histogram, and on the corresponding busyness curve, lies to the right of the peak. The left edge of the shoulder is a point where an abrupt change in slope occurs. Beginning at this point the loss of busyness due to leaving the background population is influenced by the gain in busyness due to entering the object population. A reasonable choice for the threhsold t would thus be the shoulder's left edge. Since the curve becomes less steep on the shoulder the incremental decrease in busyness for each successive threshold t to the right of the shoulder's left edge is less than the incremental increase in busyness resulting from decreasing the threshold (i.e., moving to the left of the shoulder). The shoulder edge, like the valley bottom on a bimodal curve, corresponds to a change in slope on the busyness curve. At the shoulder edge the slope becomes less negative but does not change in sign as it does at the valley bottom. In the discussions which follow we will refer to thresholds chosen at the valley bottom on a bimodal busyness curve or at the shoulder on a unimodal busyness curve as busyness thresholds.

CHAPTER III

THRESHOLD EVALUATION FOR SIMPLE CLASSES OF IMAGES

To gain a better understanding of the busyness and discrepancy criteria for threshold evaluation, we can study busyness values for simple classes of images. Suppose that an image consists of dark objects on a light background, with sharp edges between them. We shall investigate some simple models for the populations of gray levels in the object and in the background. This will involve assumptions about the distributions of gray levels and of co-occurrences of gray level pairs. Based on such models we can predict the structure of the co-occurrence matrix which determines the busyness measures for given thresholds. We will also compare busyness thresholds with discrepancy thresholds.

3.1 Images with Uncorrelated Gray Levels

Let us make the same assumption as in Section 2.1 regarding the Gaussian distributions of gray levels in the object and background, and let us further assume that the gray levels of adjacent pixels are uncorrelated. (These assumptions will be refined later.) We can specify the structure of the co-occurrence matrix and the resulting busyness measure for a given threshold t. In modeling the co-occurrence matrix we shall treat it as an infinite matrix with zero entries outside the finite submatrix M.

The Gaussian model for the image gray levels gives rise to three bivariate probability densities in the co-occurrence matrix. These represent the co-occurrences of (1) object points with object points (population A_1), (2) background points with background points (population A_2), and (3) object points with background points (population A_3), or vice versa. These densities are defined by the exponentials

$$p_1(x,y) = \frac{1}{2\pi\sigma_1^2} \exp\{-[(x-\mu_1)^2 + (y-\mu_1)^2]/2\sigma_1^2\} ;$$

$$p_2(x,y) = \frac{1}{2\pi\sigma_2^2} \exp\{-[(x-\mu_2)^2 + (y-\mu_2)^2]/2\sigma_2^2\} ; \text{ and}$$

$$p_3(x,y) = \frac{1}{2\pi\sigma_1\sigma_2} \exp\{-[(x-\mu_1)^2/2\sigma_1^2 + (y-\mu_2)^2/2\sigma_2^2]\}$$

respectively. The relative positions of these populations in the co-occurrence matrix is shown in Fig. 4.

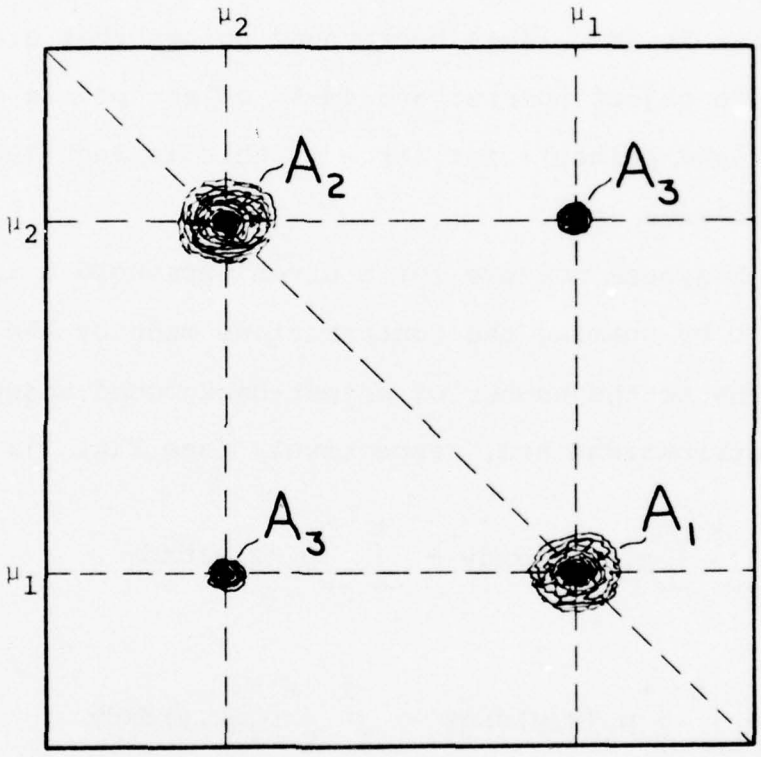

Fig. 4. Sketch of a co-occurrence matrix for images with uncorrelated gray levels.

The relative weights of these three subpopulations depend on the relative areas of the object, the background, and the object-background border. For large, compact objects, the border will have much less area than the objects. For example, if the object is an m by m solid square, the border area is 8m (4m+4 background points that are adjacent to object points, and 4m-4 object points adjacent to background points); for large m, this is much less than the object area (m^2).

The busyness measure for a given threshold t is arrived at by summing the contributions made by the three populations to the number of object-background adjacencies. These contributions are, respectively (see Fig. 5):

$$B_1(t) = \int_{-\infty}^{\infty}\int_{-\infty}^{t} p_1(x,y)\,dxdy - \int_{-\infty}^{t}\int_{-\infty}^{t} p_1(x,y)\,dxdy$$

$$+ \int_{-\infty}^{t}\int_{-\infty}^{\infty} p_1(x,y)\,dxdy - \int_{-\infty}^{t}\int_{-\infty}^{t} p_1(x,y)\,dxdy$$

$$= \int_{-\infty}^{\infty}\int_{-\infty}^{t} p_1(x,y)\,dxdy + \int_{-\infty}^{t}\int_{-\infty}^{\infty} p_1(x,y)\,dxdy -$$

$$2\int_{-\infty}^{t}\int_{-\infty}^{t} p_1(x,y)\,dxdy$$

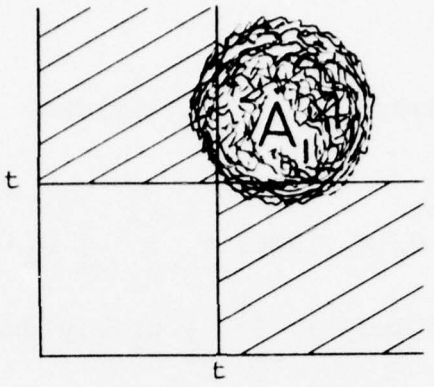

a

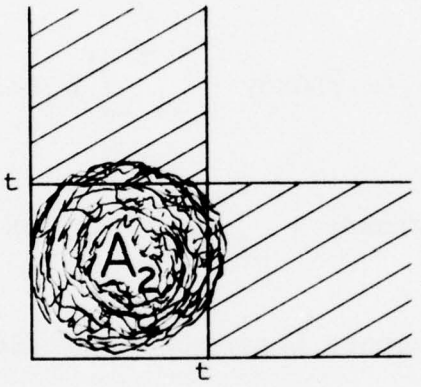

b

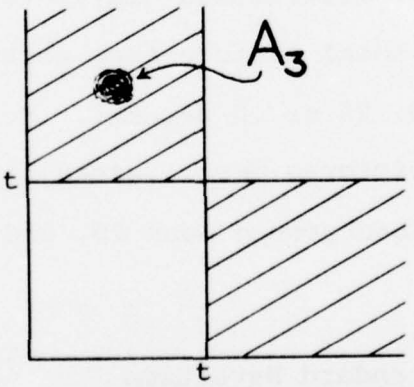

c

Fig. 5. Busyness contributions by the three bivariate probability densities in the co-occurrence matrix. 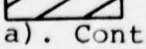 indicates contributions to busyness.

a). Contributions to busyness made by population A_1.
b). Contributions to busyness made by population A_2.
c). Contirbutions to busyness made by population A_3.

$$B_2(t) = \int_{-\infty}^{\infty}\int_{-\infty}^{t} p_2(x,y)\,dxdy - \int_{-\infty}^{t}\int_{-\infty}^{t} p_2(x,y)\,dxdy$$

$$+ \int_{-\infty}^{t}\int_{-\infty}^{\infty} p_2(x,y)\,dxdy - \int_{-\infty}^{t}\int_{-\infty}^{t} p_2(x,y)\,dxdy$$

$$= \int_{-\infty}^{\infty}\int_{-\infty}^{t} p_2(x,y)\,dxdy + \int_{-\infty}^{t}\int_{-\infty}^{\infty} p_2(x,y)\,dxdy - 2\int_{-\infty}^{t}\int_{-\infty}^{t} p_2(x,y)\,dxdy$$

$$B_3(t) = \int_{-\infty}^{\infty}\int_{-\infty}^{t} p_3(x,y)\,dxdy - \int_{-\infty}^{t}\int_{-\infty}^{t} p_3(x,y)\,dxdy$$

$$+ \int_{-\infty}^{t}\int_{-\infty}^{\infty} p_3(x,y)\,dxdy - \int_{-\infty}^{t}\int_{-\infty}^{t} p_3(x,y)\,dxdy$$

$$= \int_{-\infty}^{\infty}\int_{-\infty}^{t} p_3(x,y)\,dxdy + \int_{-\infty}^{t}\int_{-\infty}^{\infty} p_3(x,y)\,dxdy - 2\int_{-\infty}^{t}\int_{-\infty}^{t} p_3(x,y)\,dxdy$$

In order to evaluate busyness thresholds for this simple class of images, a set of synthetic images was generated. Each of these pictures, shown in Fig. 6, contains two normally distributed gray level populations. The fraction of the total picture area occupied by the object was either 50, 25 or 10 percent. For each of the three object sizes pictures were created having object gray level mean 30, background mean 20, and standard deviations as follows:

Case	Standard Deviation of Object	Standard Deviation of background
1	3	3
2	3	5
3	5	3
4	5	5

Histograms of the pictures of Fig. 6 are shown in Fig. 7.

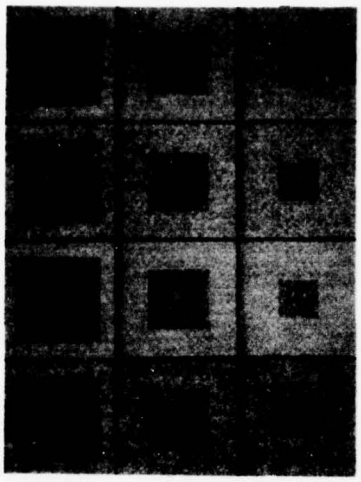

Fig. 6. Synthetic images.

For each of the above images the mean gray level of object points is 30 and the mean of background points 20. The standard deviations of object gray levels, σ_o, and of background gray levels, σ_b, are listed above. The object occupies 50% of the total picture area for pictures a, d, g and j; 25% for pictures b, e, h and k; 10% for c, f, i and l.

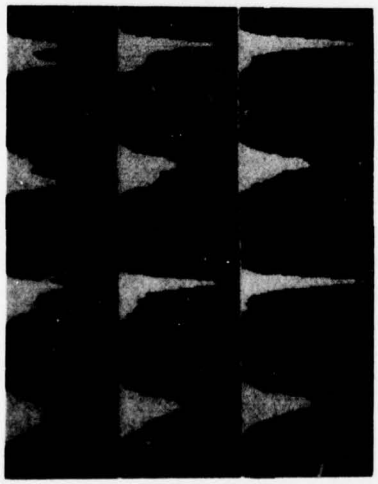

Fig. 7. Histograms of the pictures of Fig. 6.

Co-occurrence matrices computed on each of the twelve synthetic images and log-scaled for output are shown in Fig. 8. Note that since the correlations of pairs of gray levels in the object and in the background populations are zero, the component distributions that make up the co-occurrence matrices are circularly symmetric. The populations representing the co-occurrence of object gray levels with background gray levels are barely visible, if at all, in the matrices of Fig. 8. This is not surprising since the percentage of points on or adjacent to the border is a small fraction of the total area (4, 3, and 2% for the squares occupying 50, 25, and 10% of the picture area). When the distributions of object and background gray levels overlap to the extent of Fig. 8j-l, the component populations are not distinguishable in the co-occurrence matrix.

Graphs of busyness for a range of thresholds are shown in Fig. 9 for the pictures of Fig. 6. Note that each of these curves resembles the shape of its corresponding gray level histogram in Fig. 7. In addition, the overall shape of the busyness curve is smoother and less noisy than that of its corresponding histogram.

The busyness curves of Fig. 9a, b, c, d, g, h, i and j are bimodal while those of Fig. 9e, f, k and l are unimodal. For the bimodal curves busyness thresholds are chosen at the valley bottoms. The unimodal curves of Fig. 9e and f show a slight slope change about midway between the minimum and maximum. Busyness thresholds were chosen at the points on the curve corresponding to the slope

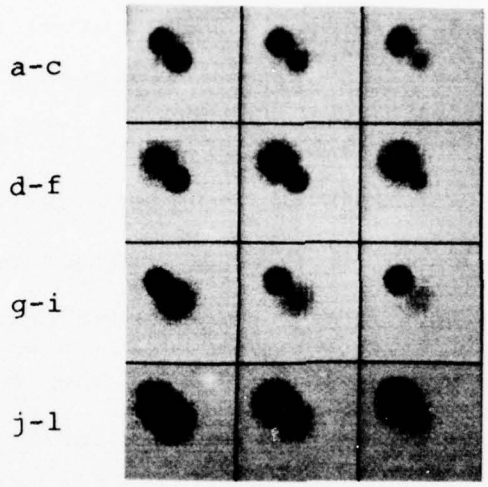

Fig. 8. Co-occurrence matrices (log-scaled) for the pictures of Fig. 6.

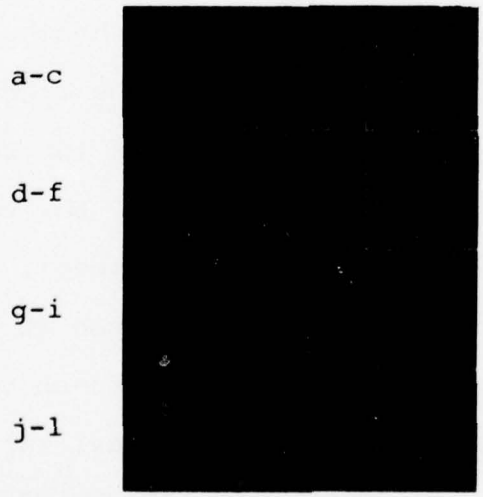

Fig. 9. Graphs of busyness (y-axis) vs. threshold (x-axis) for the pictures of Fig. 6.

change and their positions are indicated by vertical lines in Fig. 9. (This slope change can be detected automatically by the method described in [7].) The busyness thresholds chosen are listed in Table 1. The curves of Fig. 9k and l show an almost linear rate of falloff of busyness and so the busyness thresholds were selected at the object mean.

For comparison with the busyness curves, graphs of discrepancy are shown in Fig. 10. Discrepancy thresholds are also listed in Table 1. Note that only for Fig. 6ℓ is the discrepancy threshold chosen at the object mean.

For nine out of the tweleve pictures the busyness and discrepancy thresholds differ by only one gray level. In six cases the thresholds are identical (Fig. 6a, b, c, g, j and l). For the pictures in which the thresholds differ by 2, 3 or 4 (Fig. 6e, k and f, respectively) the busyness curves are unimodal. The busyness thresholds of Fig. 6e and f were chosen at the point to the right of the peak where the slope changes; for Fig. 6k the threshold was chosen at the object mean. For Fig. 6ℓ, the only other picture whose busyness curve is unimodal, both the busyness and discrepancy thresholds were chosen at the object mean. In summary, the discrepancy and busyness thresholds differ by at most 1 in 9 out of the 12 cases, including all the cases where the busyness curves are bimodal.

The results of thresholding the pictures of Fig. 6 at every threshold in the range t = 17 to t = 32 are shown in Fig. 11. A threshold at gray level t maps all gray levels greater than t into the object and all other levels into

Part of Fig. 6	Std. Dev of Object	Std. Dev. of Background	% of Total Area Occupied by Object	Busyness Threshold	Discrepancy Threshold
a	3	3	50	24	24
b	3	3	25	25	25
c	3	3	10	26	26
d	3	5	50	24	25
e	3	5	25	25	27
f	3	5	10	25	29
g	5	3	50	24	24
h	5	3	25	26	25
i	5	3	10	28	27
j	5	5	50	24	24
k	5	5	25	30	27
l	5	5	10	30	30

Table 1. Busyness and discrepancy thresholds for the pictures of Fig. 6.

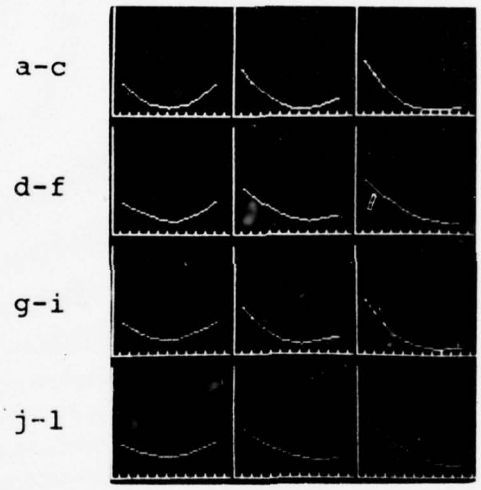

Fig. 10. Graphs of discrepancy (y-axis) vs. threshold (x-axis) for the pictures of Fig. 6.

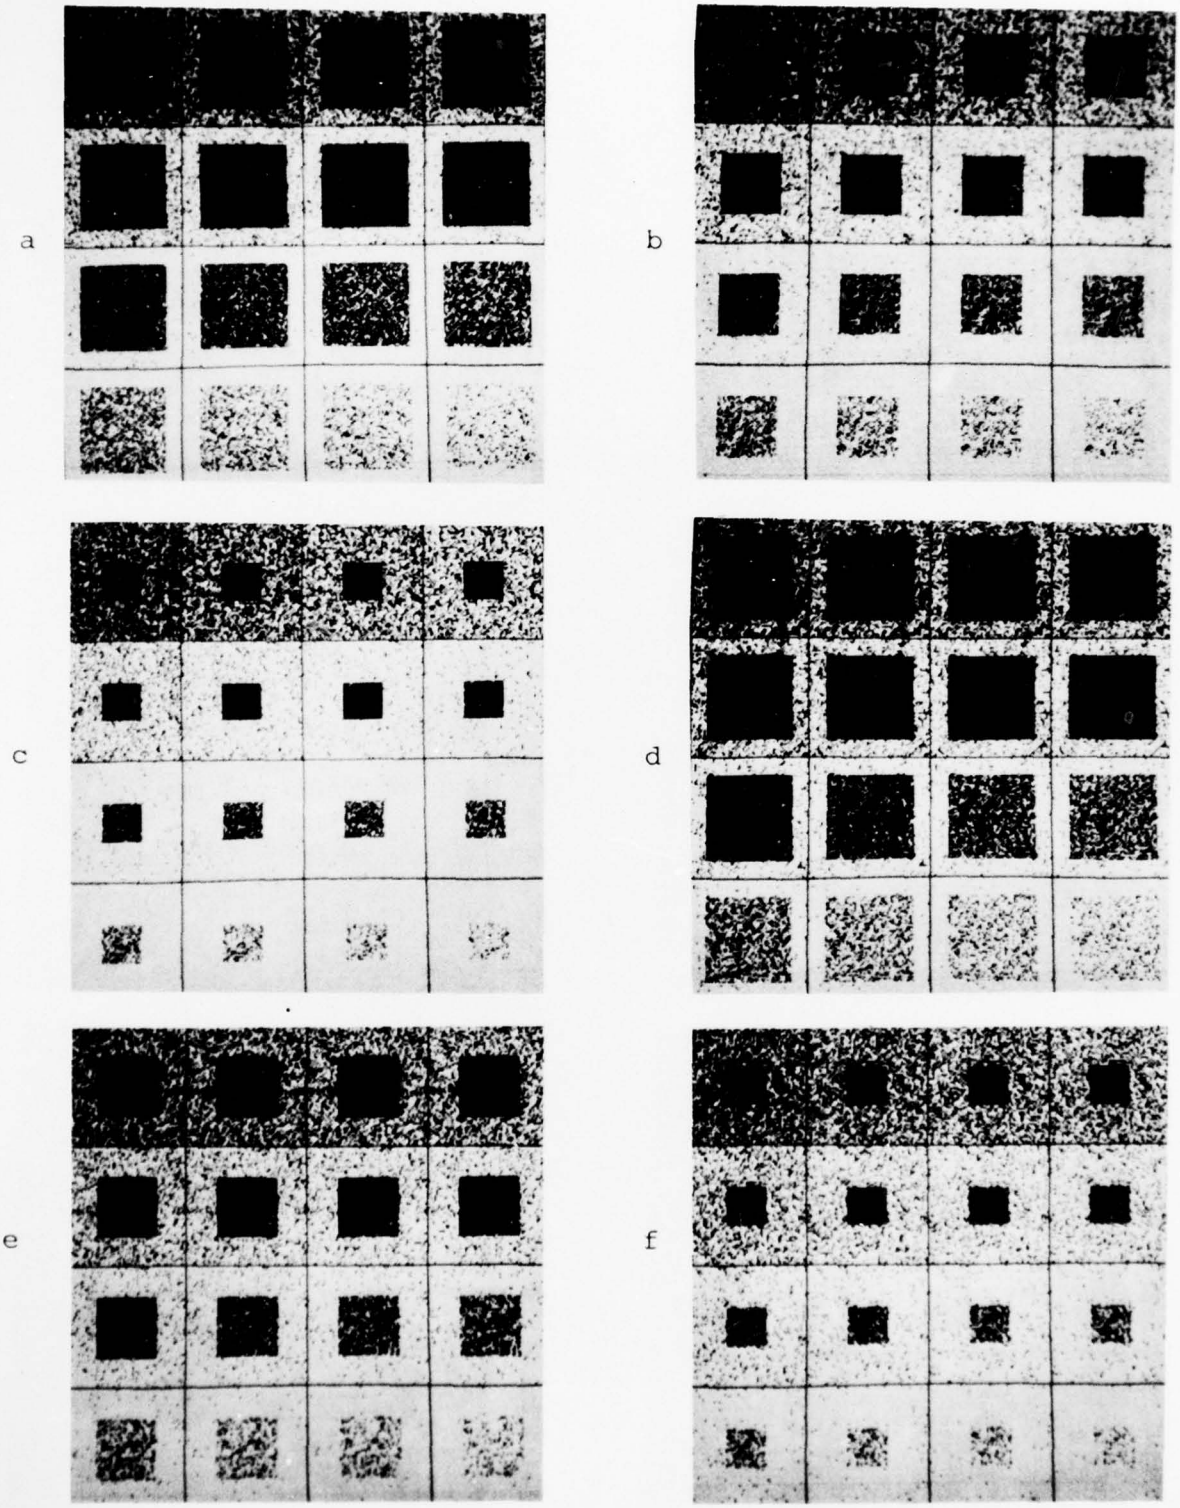

Fig. 11a-f). Results of thresholding the pictures of Fig. 6a-f at every threshold in the range t = 17 to t = 32.

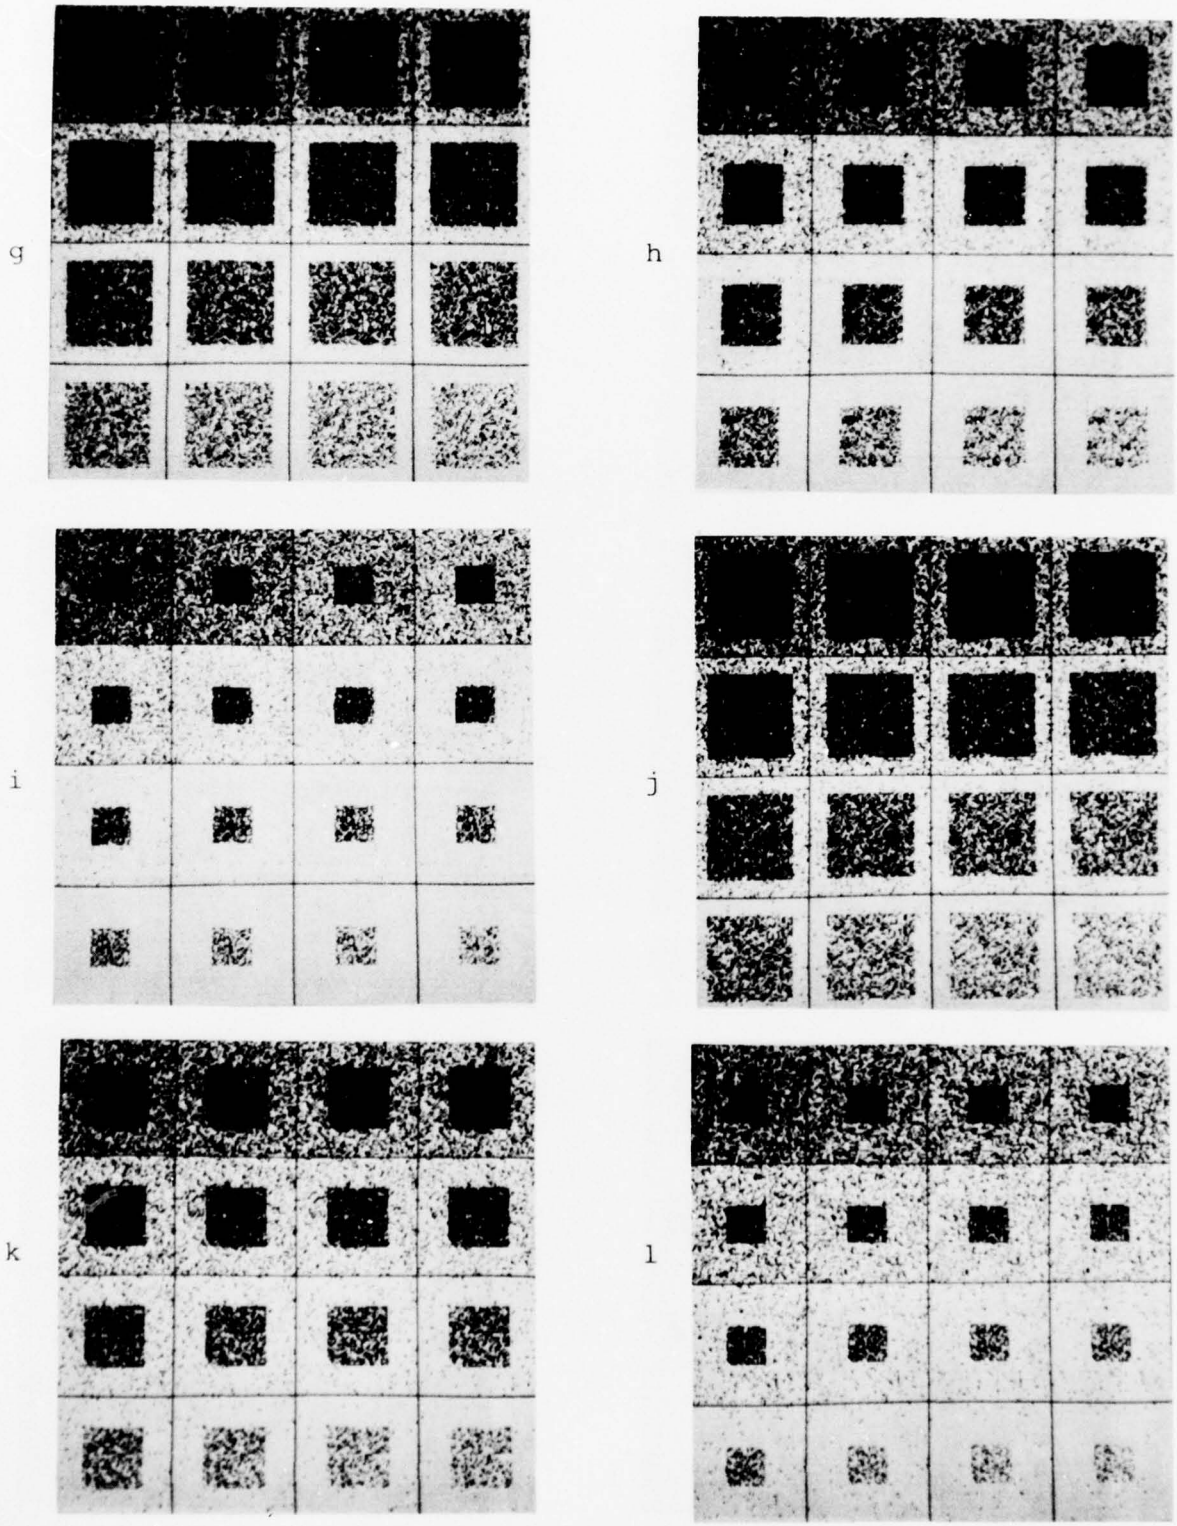

Fig. 11g-1). Results of thresholding the pictures of Fig. 6g-1 at every threshold in the range t = 17 to t = 32.

the background. From Fig. 11 we see that in most cases the thresholds listed in Table 1 are reasonable thresholds.

3.2. Images with Correlated Gray Levels

In our model up to now we have made an unrealistic assumption about the correlations of each of the three populations in the co-occurrence matrix. We have assumed that pairs of gray levels in the object (or background) are independent, i.e., have correlation zero. This implies that the density functions in the co-occurrence matrix representing these points are circularly symmetric.

A more realistic model would assume a non-zero correlation between neighboring points. A correlation near 1 would tend to elongate the density function along the diagonal, while a negative correlation would elongate the density along a line perpendicular to the diagonal. In real images the correlations of points in the object and background populations should be positive, since adjacent points which are not near the border would tend to have the same or similar gray levels. The equations for the bivariate density functions assuming correlated data are

$$p_1(x_1,x_2) = \frac{1}{(2\pi)|\Sigma_1|^{1/2}} \exp[-\tfrac{1}{2}(X-M)' \Sigma_1^{-1}(X-M)]$$

$$\text{where } \Sigma_1 = \begin{bmatrix} \sigma_{11}^2 & \sigma_{12}^2 \\ \sigma_{21}^2 & \sigma_{22}^2 \end{bmatrix} = \begin{bmatrix} \sigma_o^2 & \sigma_{oo}^2 \\ \sigma_{oo}^2 & \sigma_o^2 \end{bmatrix} ;$$

$$X = (x_1,x_2);$$

$$M = (\mu_o,\mu_o);$$

and μ_o, σ_o^2, are the mean and variance of gray levels

in the object, and σ^2_{oo} is the covariance of adjacent gray levels in the object.

$$p_2(y_1,y_2) = \frac{1}{(2\pi)|\Sigma_2|^{1/2}} \exp[-\tfrac{1}{2}(X-M)' \Sigma_2^{-1}(X-M)]$$

where $\Sigma_2 = \begin{bmatrix} \sigma^2_{11} & \sigma^2_{12} \\ \sigma^2_{21} & \sigma^2_{22} \end{bmatrix} = \begin{bmatrix} \sigma^2_b & \sigma^2_{bb} \\ \sigma^2_{bb} & \sigma^2_b \end{bmatrix}$;

$X = (y_1,y_2)$;

$M = (\mu_b,\mu_b)$;

and μ_b, σ^2_b are the mean and variance of gray levels in the background, and σ^2_{bb} is the covariance of adjacent gray levels in the background.

$$p_3(x,y) = \frac{1}{(2\pi)|\Sigma_3|^{1/2}} \exp[-\tfrac{1}{2}(X-M)' \Sigma_3^{-1}(X-M)]$$

where $\Sigma_3 = \begin{bmatrix} \sigma^2_{11} & \sigma^2_{12} \\ \sigma^2_{21} & \sigma^2_{22} \end{bmatrix} = \begin{bmatrix} \sigma^2_o & \sigma^2_{ob} \\ \sigma^2_{bo} & \sigma^2_b \end{bmatrix}$

$X = (x,y)$:

$M = (\mu_o,\mu_b)$;

and σ^2_{ob}, σ^2_{bo} are the covariances of object gray levels with background gray levels and vice versa.

Another assumption we have made which does not correspond to what occurs in real images is that the edges between objects and background are sharp. This assumption does not apply to most classes of real pictures since the edges are blurred.

We will study a simple class of blurred synthetic images created by unweighted averaging of the synthetic images of Section 3.1 over square neighborhoods. Although a linear blurring process does not exactly model blurring in real images, it does provide a reasonable first approximation to many real blurring processes. These images have nonzero correlations between adjacent gray levels in the object (or in the background), and also have blurred edges.

Three sizes of blurring neighborhoods were applied to the pictures of Fig. 6: 2x2, 4x4, and 8x8 (see Figs. 12, 13 and 14, respectively). Histograms of these pictures are shown in Figs. 15, 16, and 17. As can be seen from the histograms, the blurring process resulted in a sharpening of the peaks corresponding to the object and background (since the standard deviations of the object and background distributions were lowered). Hence, the task of threshold selection for these pictures is much easier than for those of Fig. 6. Thresholds could be chosen at the valley bottoms on the histograms of Fig. 15 and 16; for Fig. 17 the valley is broader and relatively flat so we need a method for deciding which threshold in the valley to choose. The threshold selection schemes we have described are especially useful when the histogram is not strongly bimodal; in these cases conventional threshold selection techniques are not readily applicable. However, we will examine the results of using these threshold selection methods on the blurred images of Fig. 12-14 in order to attempt to understand how our methods are affected by the blurring process.

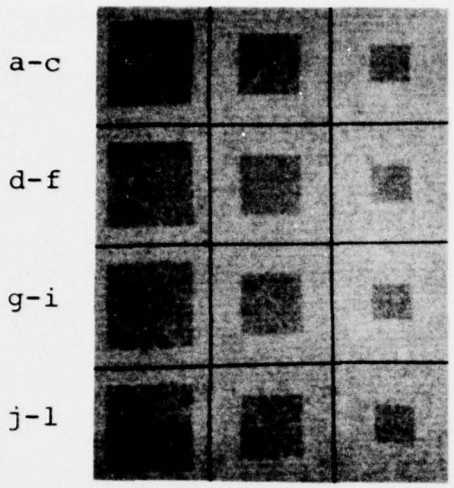

Fig. 12.

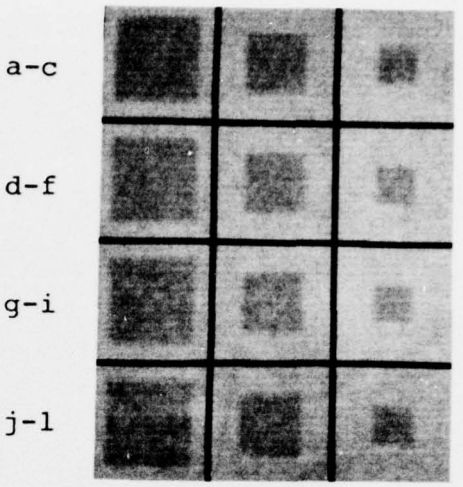

Fig. 13.

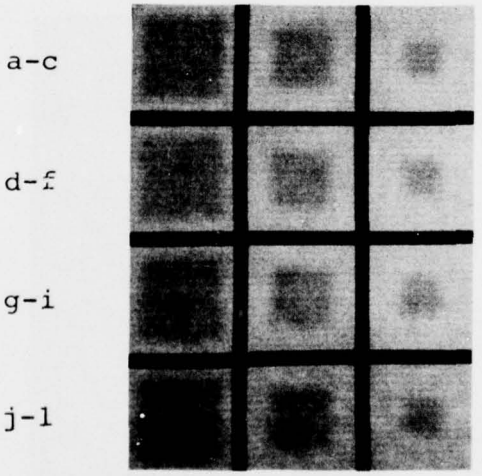

Fig. 14.

Figs. 12, 13, 14. Results of averaging the pictures of Fig. 6 over 2x2, 4x4 and 8x8 neighborhoods, respectively.

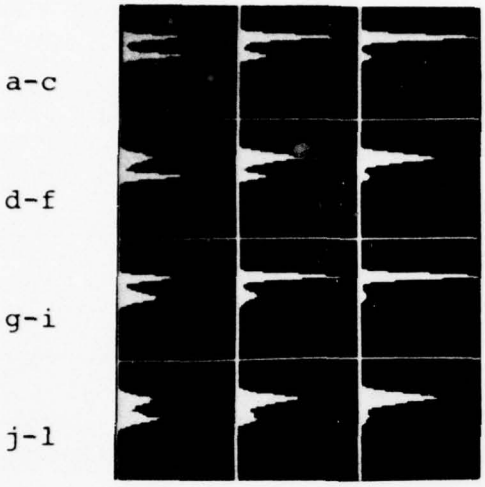

Fig. 15.

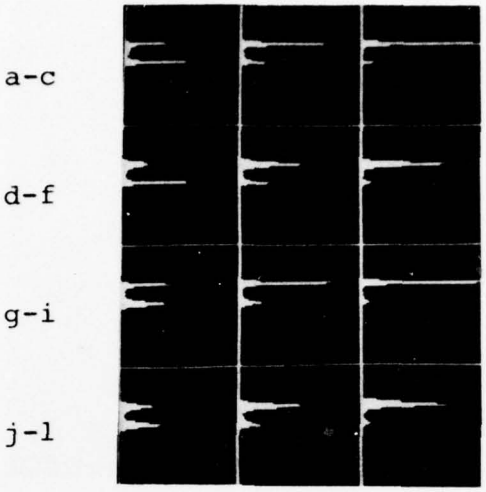

Fig. 16.

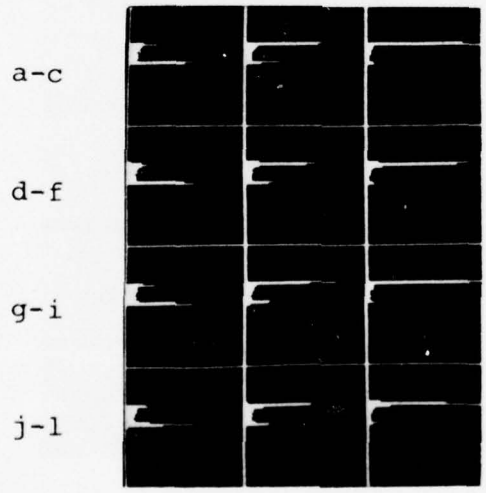

Fig. 17.

Figs. 15, 16, 17. Histograms of the pictures of Figs. 12, 13 and 14, respectively.

We can compute the resulting correlations of adjacent gray levels in the object and in the background for each blurred picture. The correlations of a pair of horizontally or vertically adjacent gray levels in the objects or in the backgrounds of the blurred pictures of Figs. 12-14 are .5, .75, and .875, respectively, since these gray levels arise from pairs of blurring neighborhoods that overlap by 50%, 75%, and $85\frac{1}{2}$%. The mean gray levels of points in the objects and in the backgrounds of the blurred pictures are the same as those for the pictures containing uncorrelated gray levels (means 20 and 30, respectively). The standard deviations of the object and background distributions in the blurred picture are given hy $\sigma/\sqrt{n}$ where n is the number of points in the blur neighborhood and σ is the standard deviation of gray levels before the blurring operation is applied. Table 2 contains a list of the standard deviations of the object and background populations of Figs. 6 and 12-14.

Co-occurrence matrices computed on the blurred images of Figs. 12-14 are shown in Figs. 18-20. The following observations can be made by comparing these matrices to those of Fig. 8, which were computed on the sharp images of Fig. 6:

(1) The population representing the co-occurrence of object gray levels and that representing the co-occurrence of background gray levels are elongated along the main diagonal. The degrees to which these populations are

Figure No.	Figure Part	Object Standard Deviation	Background Standard Deviation
6	a, b, c	3.0	3.0
	d, e, f	3.0	5.0
	g, h, i	5.0	3.0
	j, k, l	5.0	5.0
12	a, b, c	1.5	1.5
	d, e, f	1.5	2.5
	g, h, i	2.5	1.5
	j, k, l	2.5	2.5
13	a, b, c	.75	.75
	d, e, f	.75	1.25
	g, h, i	1.25	.75
	j, k, l	1.25	1.25
14	a, b, c	.375	.375
	d, e, f	.375	.625
	g, h, i	.625	.375
	j, k, l	.625	.625

Table 2. Standard deviations of object and background gray level populations for the synthetic pictures of Figs. 6, 12, 13, and 14.

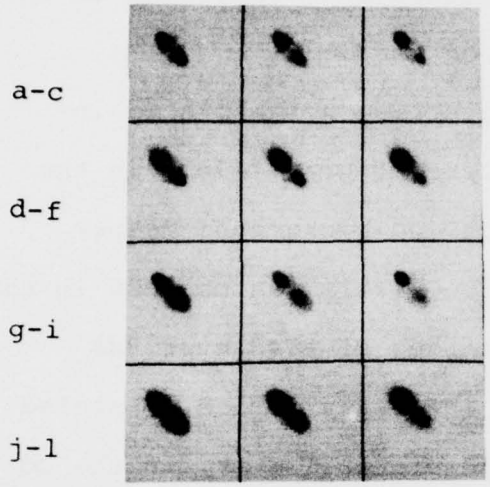

Fig. 18.

Fig. 19.

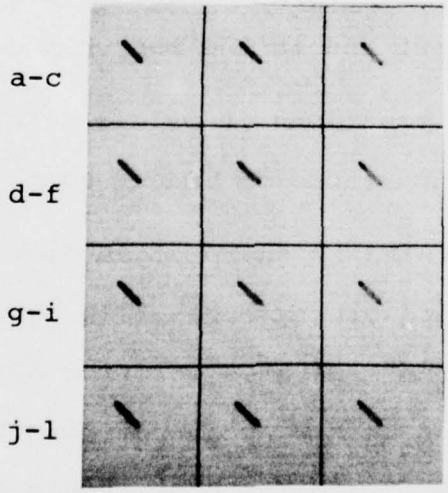

Fig. 20.

Figs. 18, 19, 20. Co-occurrence matrices (log-scaled) computed on the pictures of Figs. 12, 13 and 14, respectively.

elongated and pulled in towards the main diagonal are dependent on their correlations, which vary according to the size of the blurring neighborhood used.

(2) As the size of the blurring neighborhood increases, the distribution of pairs of gray levels in the middle range (between the object and background means) approaches a uniform distribution. This can be seen in the co-occurrence matrices of the images blurred over 8x8 neighborhoods. The gray level intensity (which is scaled logarithmically according to the probabilities in the co-occurrence matrix) is relatively constant on or near the diagonal between the two dark spots, which correspond to hills in the three-dimensional plot of the co-occurrence matrix, and which represent adjacencies of pairs of gray levels in the object and in the background.

From the above observations it is clear that our model for the co-occurrence matrices of images containing objects on a background separated by sharp edges does not apply to the case of objects with blurry edges. In blurred images, since the blurry edge zone separates interior object points from interior background points, pairs of points on the blurry edge have the same or similar gray levels and contribute to the co-occurrence matrix along or near the main diagonal. Since the gray levels on the edge lie between the object mean and the background mean, the population of pairs of points on the edge connects the populations of co-occurring object points and co-occurring background points. Thus, it is reasonable that a model for the co-occurrence matrix of

a blurred image would consist of three populations, all of which are centered on (and symmetric about) the main diagonal.

Several possible models could be used to quantitatively describe the co-occurrence matrix of a blurred image. We could once again assume that both the object and background gray level populations can be modeled by Guassian distributions. The population of edge points, however, could more naturally be modelled by a uniform distribution. An alternative would be to model the edge points by another Gaussian distribution whose mean lies between that of the object and background and whose standard deviation is such that the population is relatively flat and straddles the object and background populations. Neither of these models will be examined in detail in this paper.

Busyness measures based on co-occurrence matrices were computed for thresholds in the range t = 17 to t = 32 for each of the blurred images of Figs. 12-14. Busyness thresholds are listed in Table 3. Graphs of busyness values are shown in Figs. 21-23. On these curves, local minima are always obtained between the means. Note that for Figs. 13 and 14 there is a range of thresholds between the means over which busyness changes very slowly.

The results of thresholding the pictures of Figs. 12-14 at every threshold in the range t = 17 to t = 32 are shown in Figs. 24-26. The busyness thresholds for the pictures blurred over 2x2 neighborhoods seem to minimize the total amount of noise in the object and background of the resulting

Part of Fig.	Minimum-Busyness Thresholds			
	No Blur (Fig. 6)	2x2 Blur (Fig. 12)	4x4 Blur (Fig. 13)	8x8 Blur (Fig. 14)
a	24	24	25	27
b	25	24	26	28
c	26	25	27	28
d	24	25	27	28
e	25	25	25	27
f	25	26	27	28
g	24	23	25	27
h	26	24	25	27
i	28	24	25	28
j	24	24	25	27
k	30	25	24	27
l	30	26	25	27

Table 3. Busyness thresholds for the pictures of Fig. 6 and the blurred pictures of Figs. 12, 13 and 14.

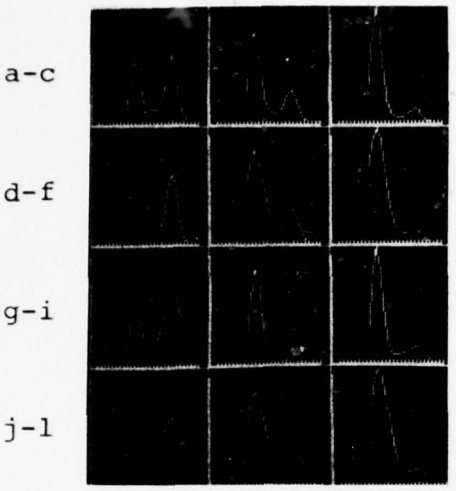

Fig. 21.

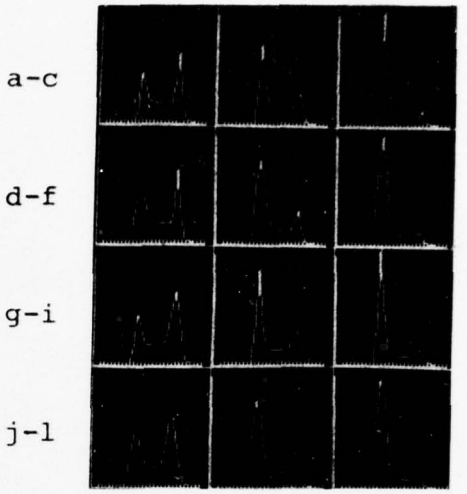

Fig. 22.

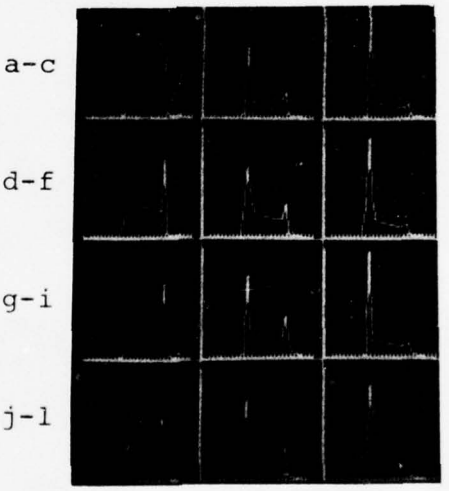

Fig. 23.

Figs. 21, 22, 23. Graphs of busyness (y-axis) vs. threshold (x-axis) for the pictures of Figs. 12, 13, and 14, respectively.

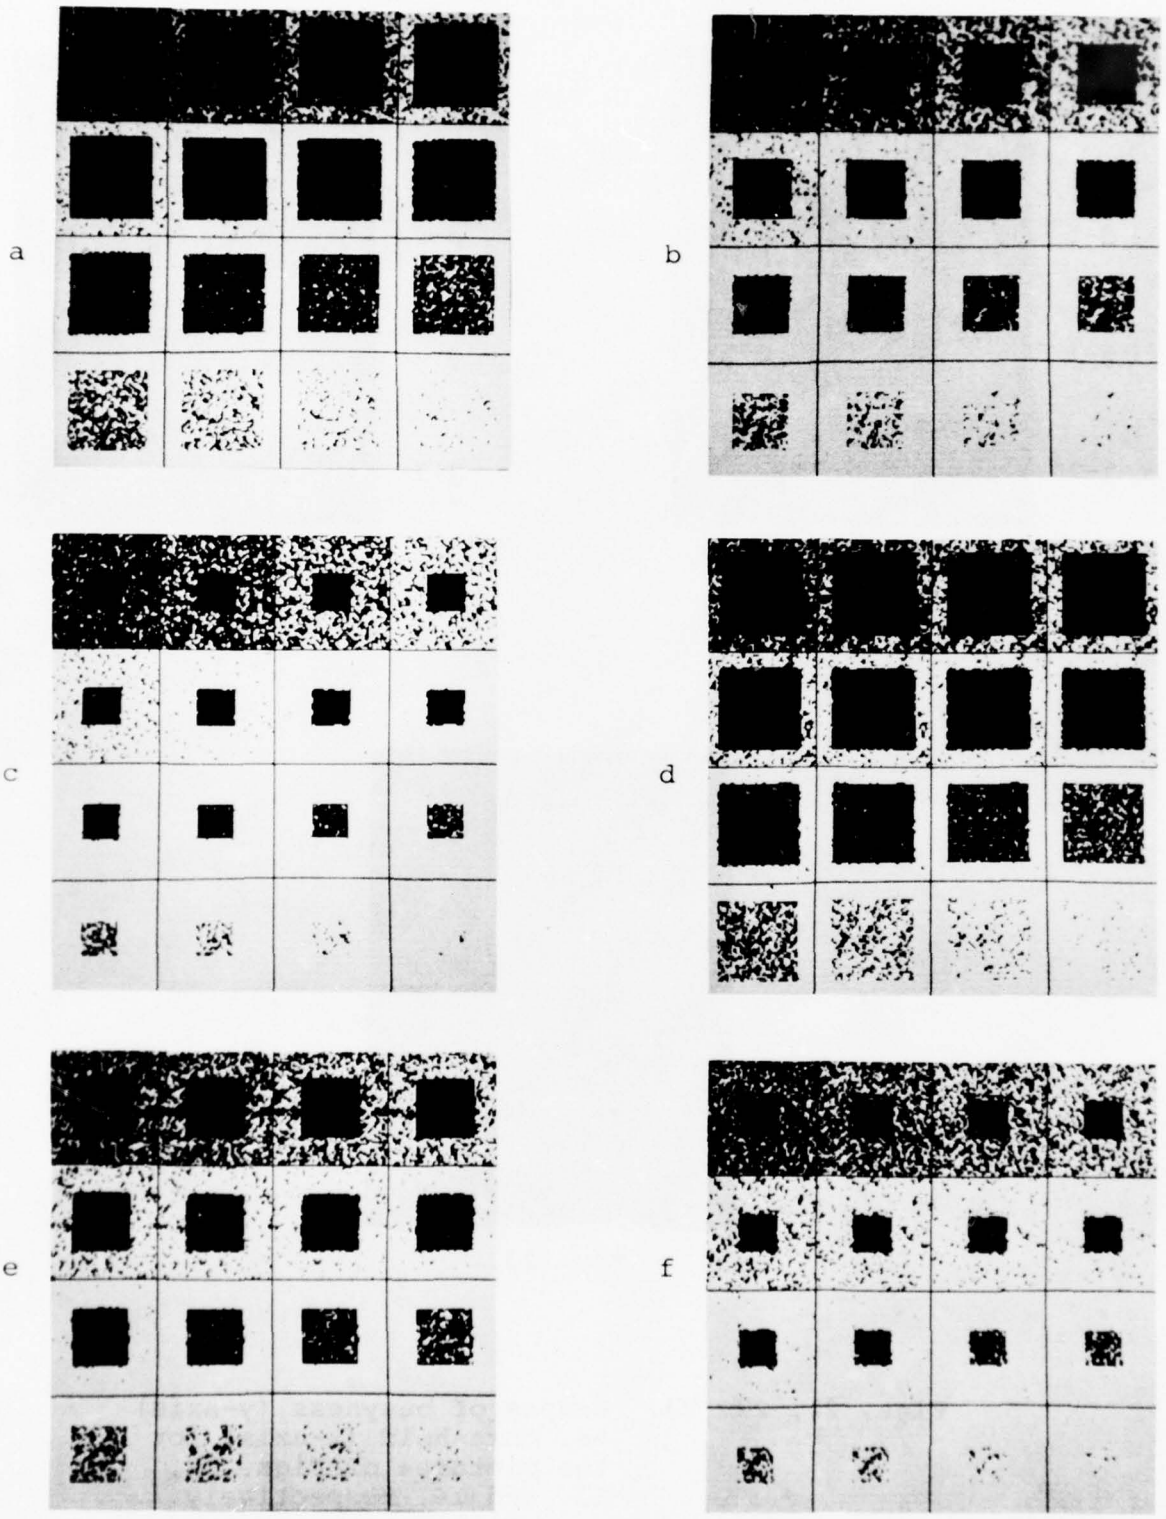

Fig. 24a-f). Results of thresholding the pictures of Fig. 12a-f at every threshold in the range t = 17 to t = 32.

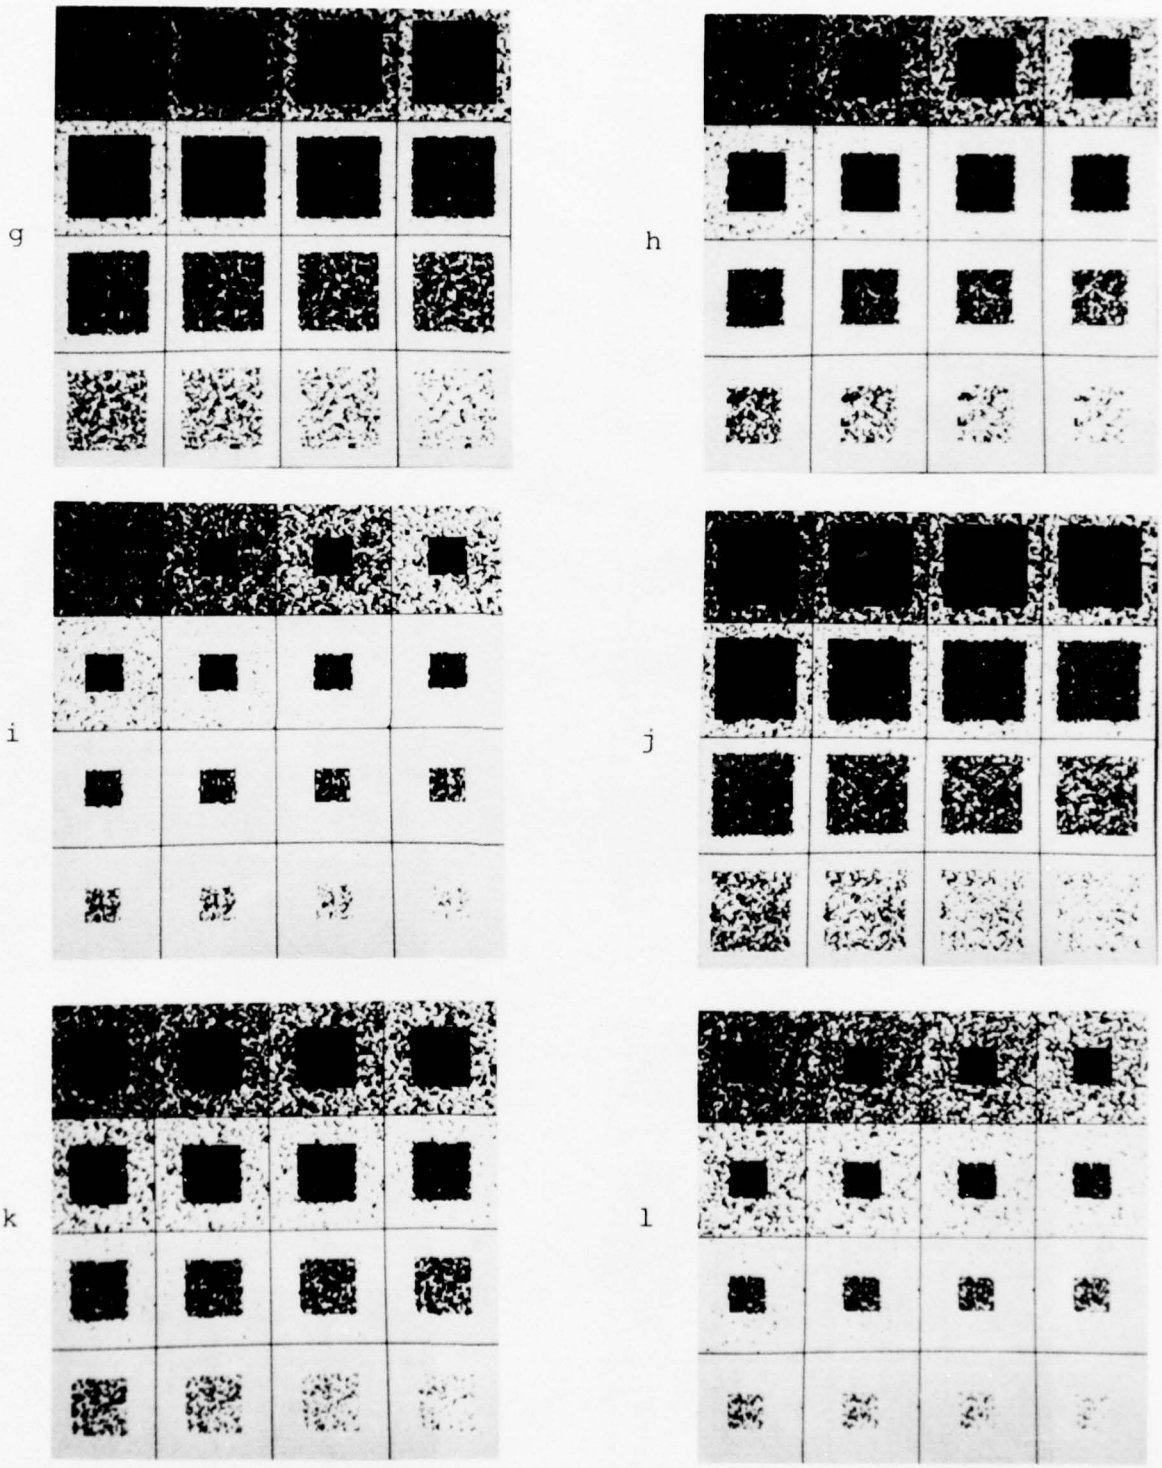

Fig. 24g-1). Results of thresholding the pictures of Fig. 12g-1 at every threshold in the range t = 17 tc t = 32.

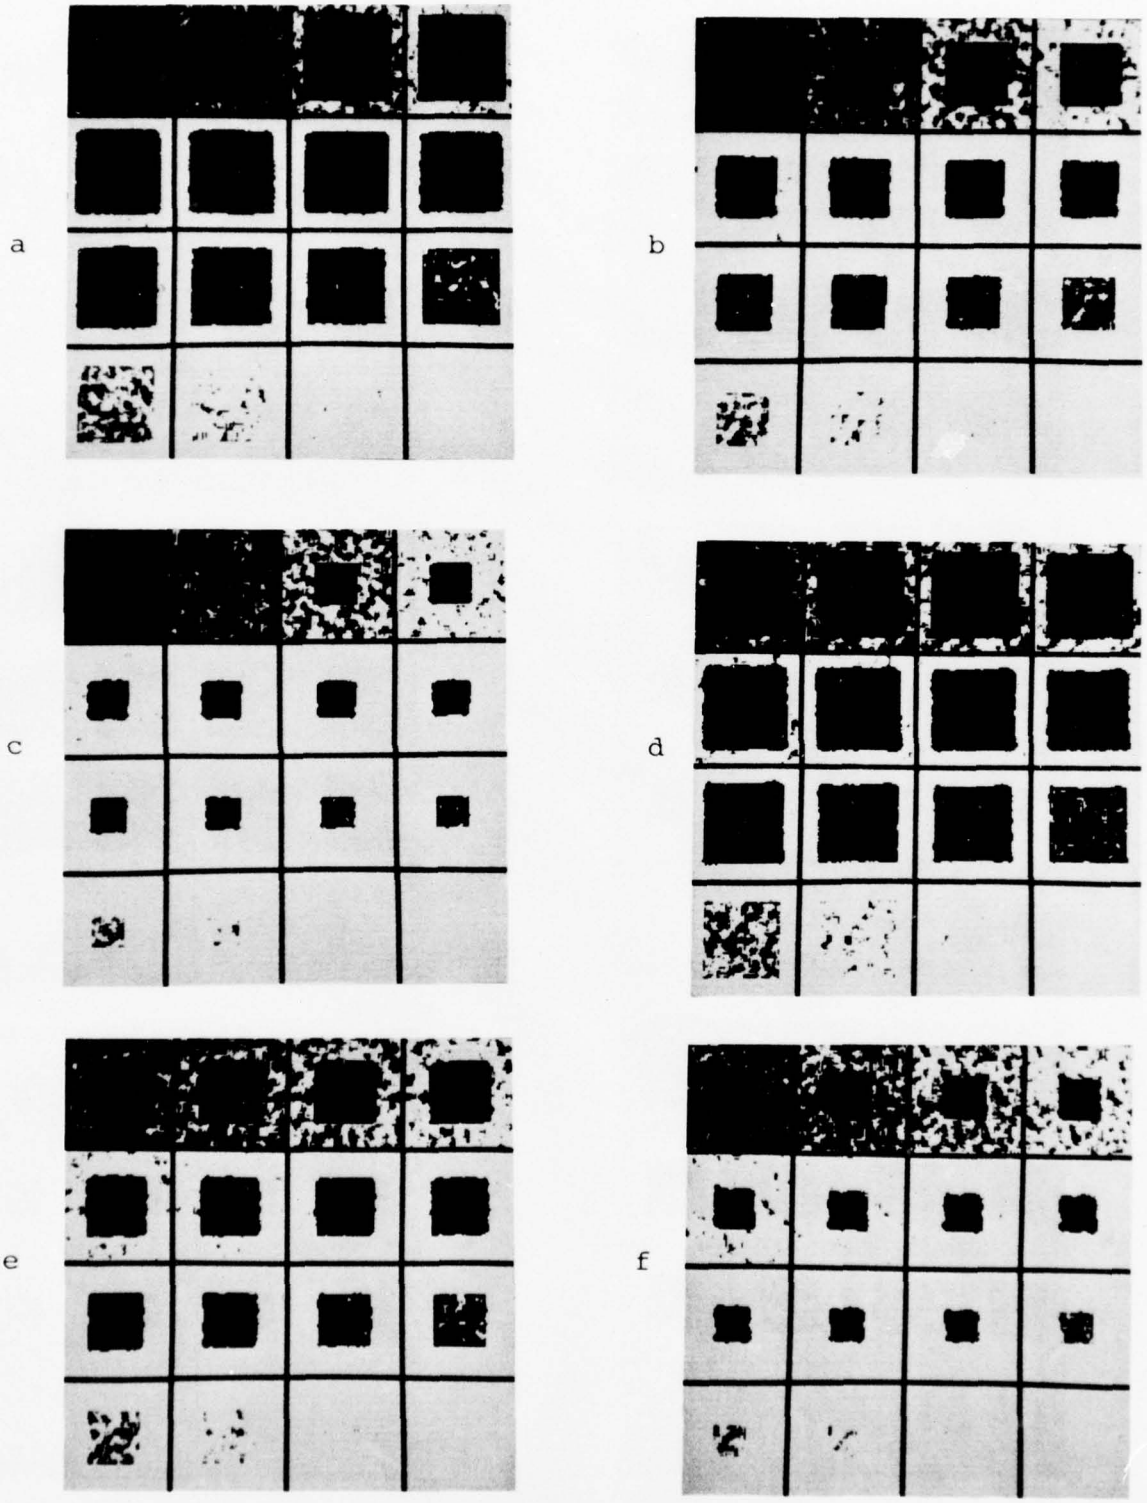

Fig. 25a-f). Results of thresholding the pictures of Fig. 13a-f) at every threshold in the range t = 17 to t = 32.

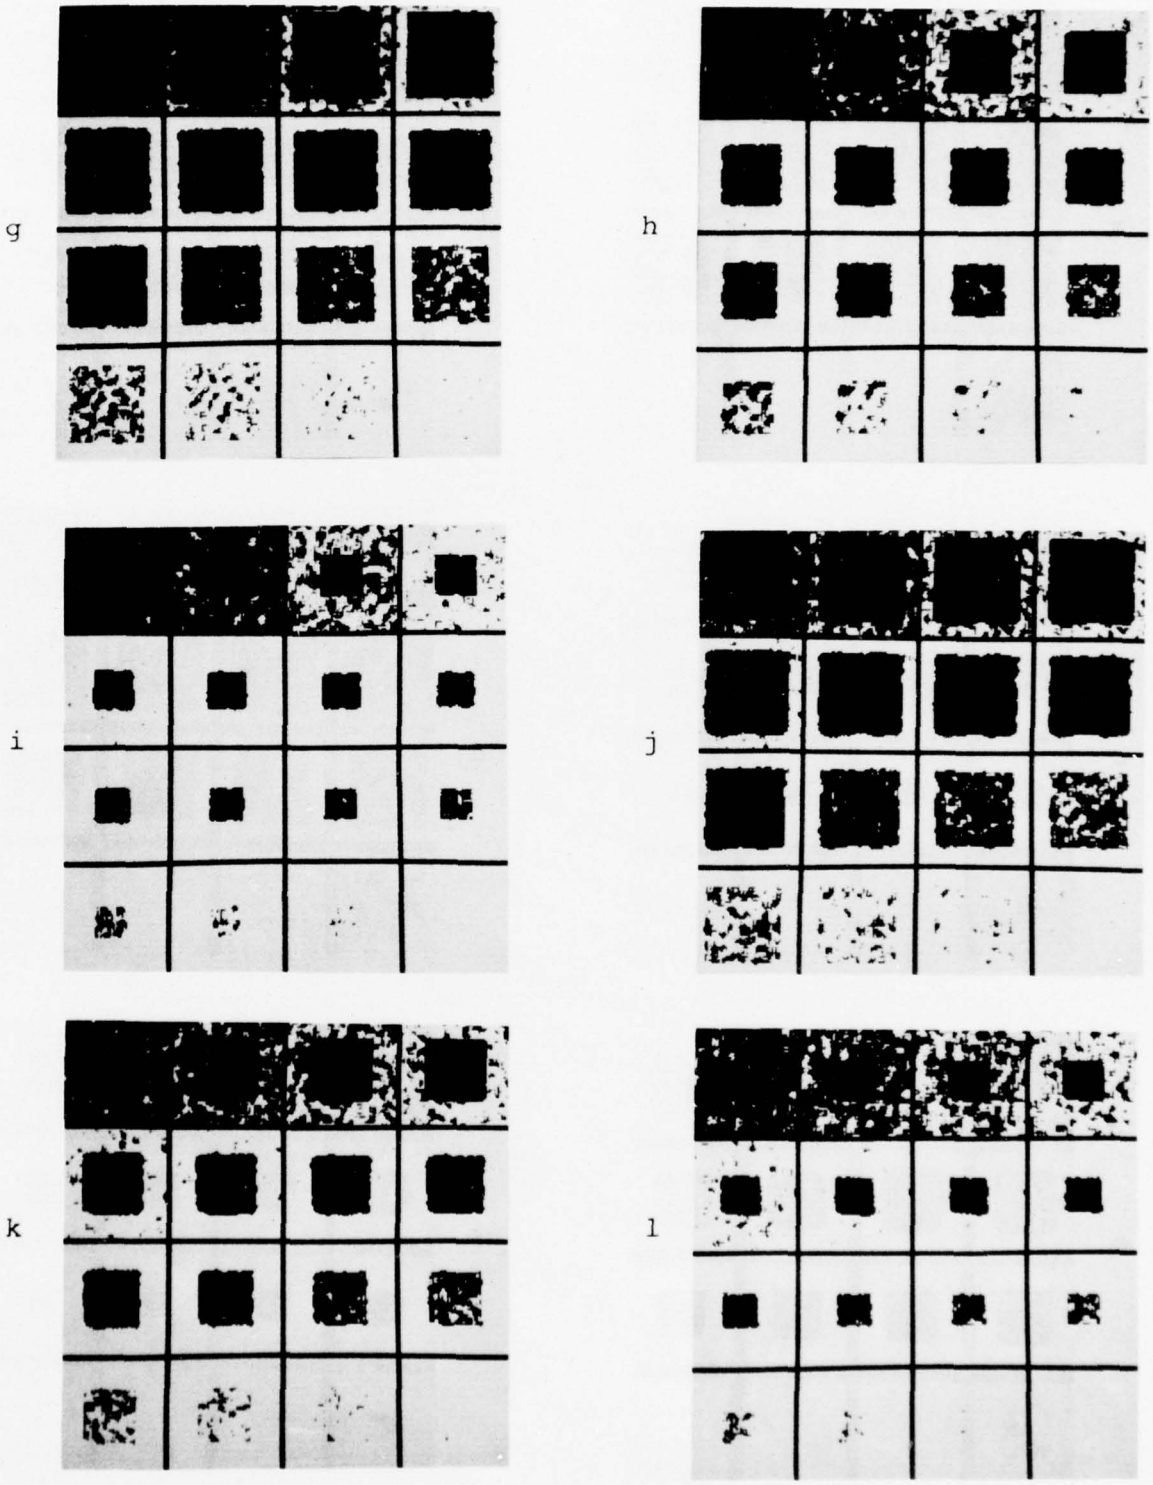

Fig. 25g-1). Results of thresholding the pictures of Fig. 13g-1 at every threshold in the range t = 17 to t = 32.

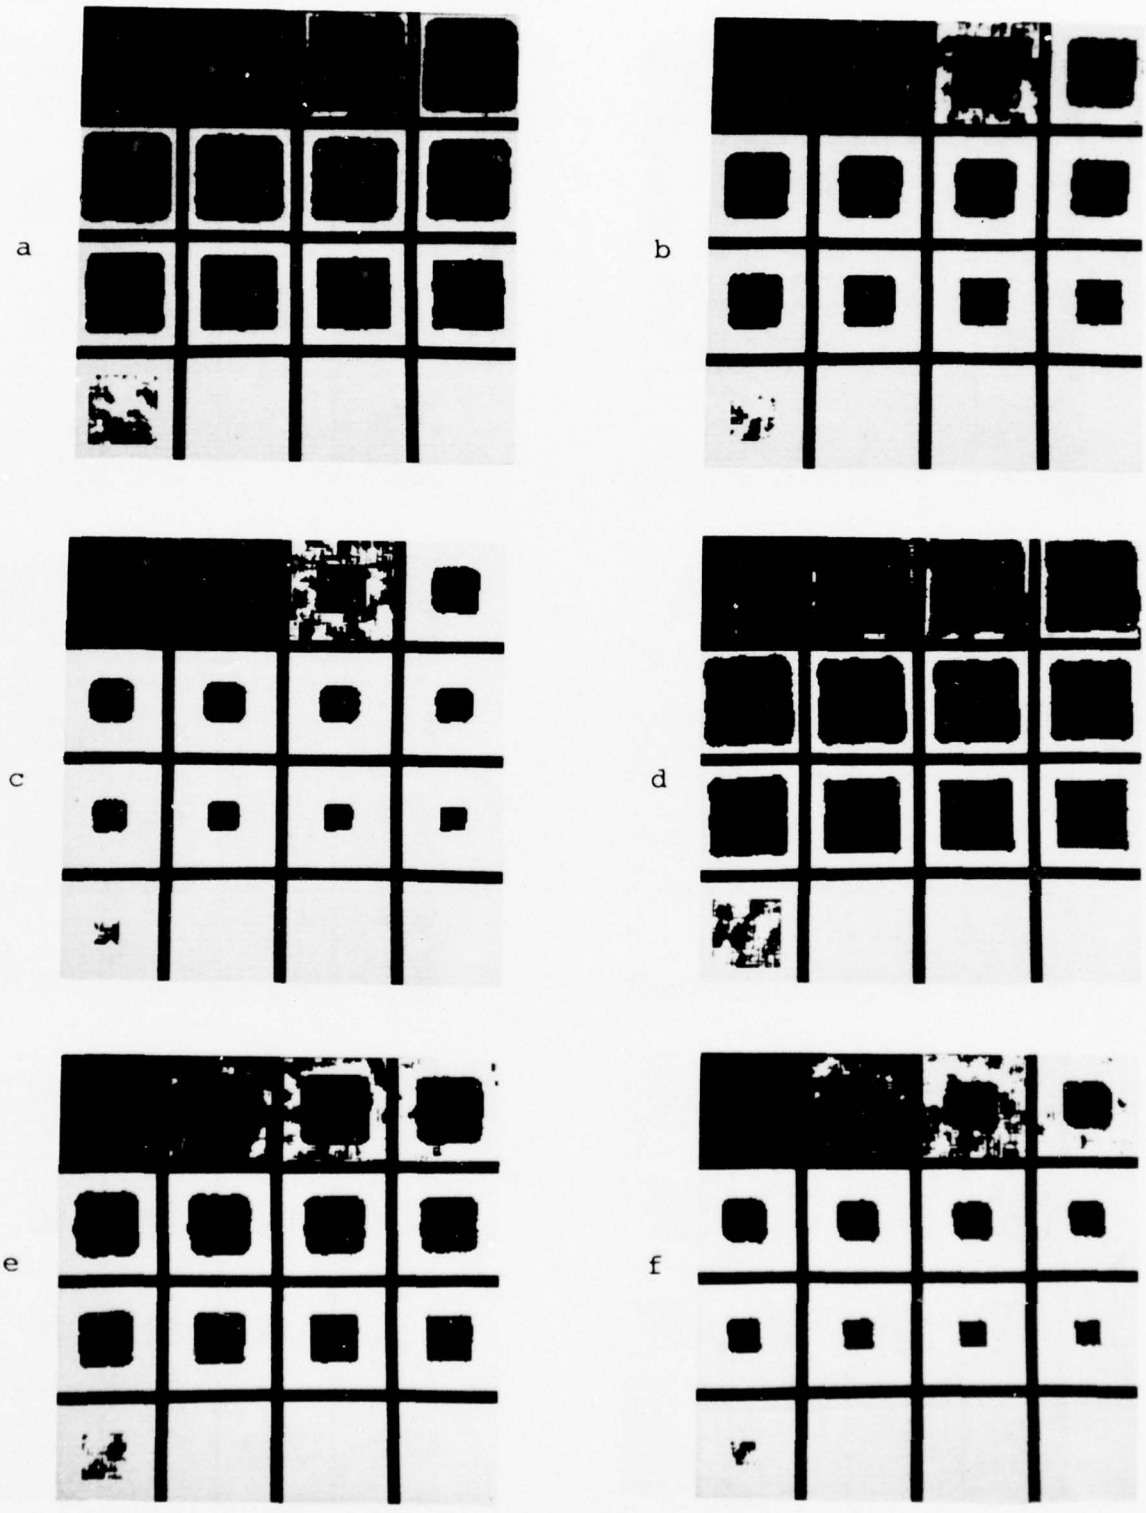

Fig. 26a-f). Results of thresholding the pictures of Fig. 14a-f at every threshold in the range t = 17 to t = 32.

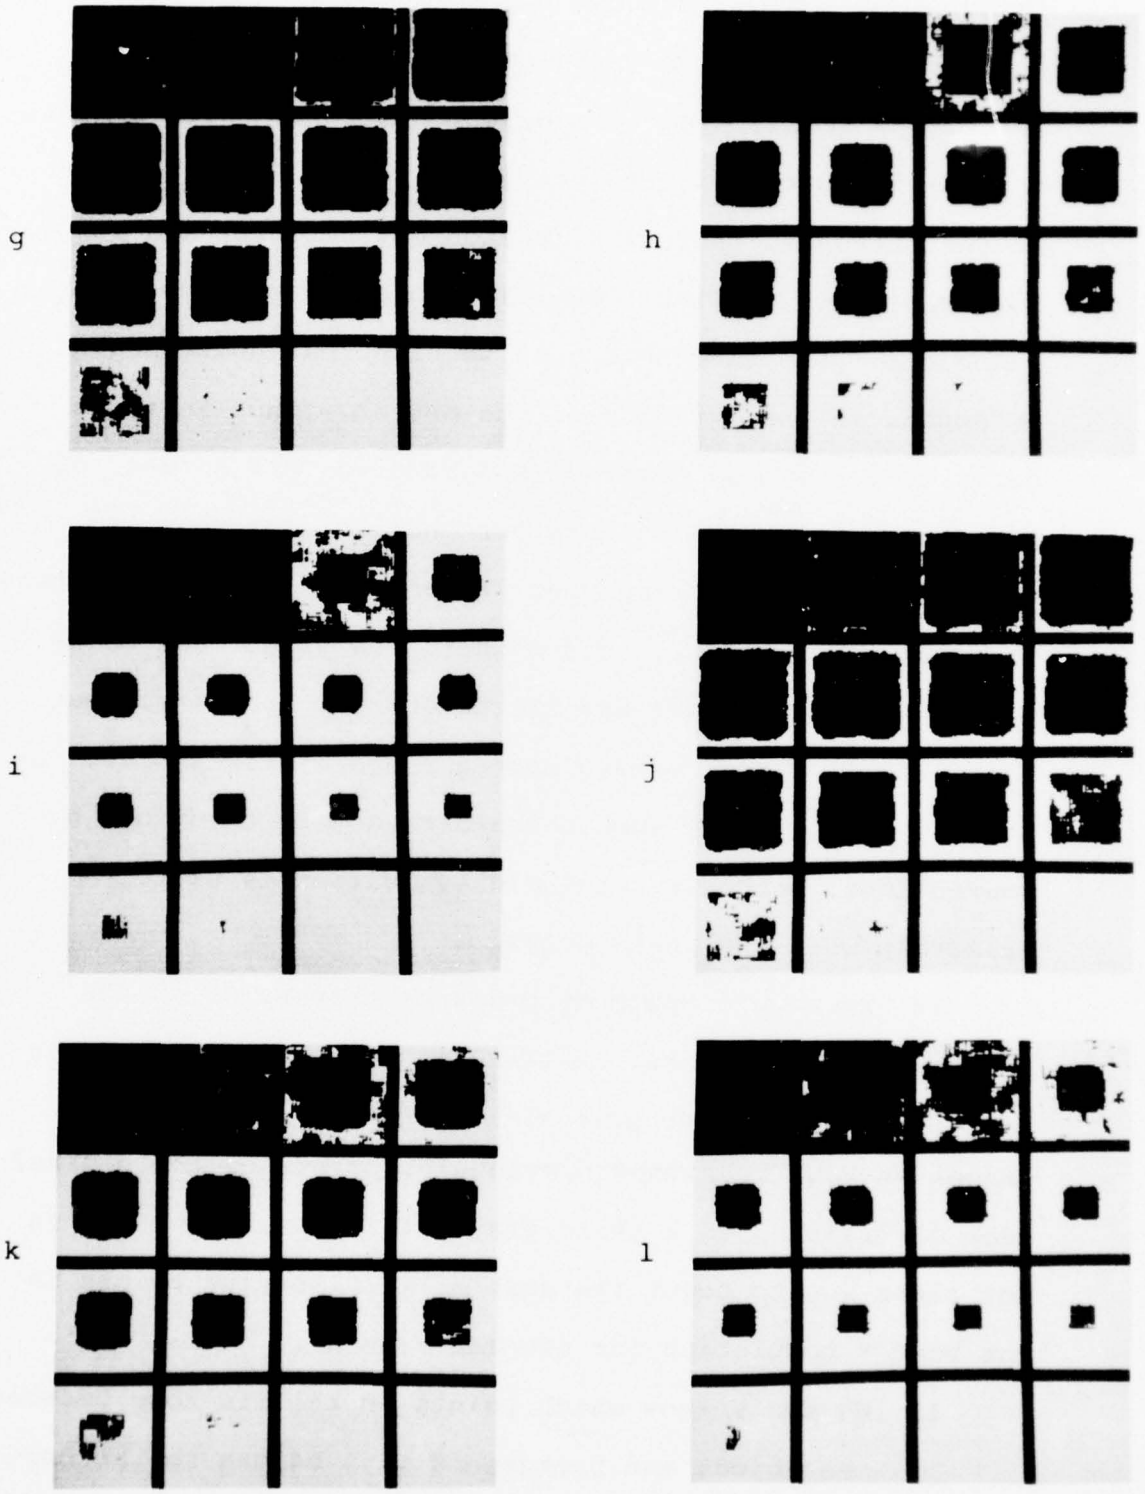

Fig. 26g-1). Results of thresholding the pictures of Fig. 14g-1 at every threshold in the range t = 17 to t = 32.

thresholded image. For the higher degrees of averaging, 4x4 and 8x8, the noise in the object and background interiors is nearly zero over a range of thresholds, so that the chief contribution to busyness is made by the border. As the threshold increases, the object shrinks, and its border becomes smaller, resulting in a lower busyness value until the threshold gets so high that noise begins to appear in the interior of the object. Thus in these cases, the busyness measure is dominated over a range of thresholds by object size rather than by noisiness, so that it is not measuring what we wanted it to. Fortunately, this can only happen in cases where the object and background histogram peaks are narrow and far apart, so that noisiness is close to zero over a range of thresholds. As pointed out earlier, such cases are easy to threshold by conventional methods; threshold evaluation is useful primarily when the peaks overlap.

In order to compute discrepancy thresholds on the blurred images of Figs. 12-14 we must assign some cost to misclassifying border points. In the model for blurred images we have mentioned previously, points on the blurred edge contributed to a third gray level population. It is not clear how to count the cost of misclassying points in the border population for several reasons:

1) We don't know which points in the mid-zone between the object and background mean belong to the border population.

2) Even if we could isolate the border population, since thresholding requires a two-level output, we would be faced with the task of deciding which border points belong to the object population and which belong to the background population.

One solution to the problem of assigning costs to misclassified border points is to assume that the misclassification cost is zero for these points. This corresponds to using the parameters of the object and background distributions and their a priori probabilities to determine the probability of error as a function of t, as in Section 2.1. The weights of the object, background and border populations are listed in Table 4 for each of the pictures of Fig. 12, 13 and 14. These percentages were computed using the actual numbers of points in each of the three areas of the images.

Assuming that the misclassification cost for border points is zero, minimum error thresholds were computed for the blurred pictures of Fig. 12-14 using the probabilities of the object and background distributions listed in Table 4. These thresholds, which are listed in Table 5, always occur strictly between the object and background means.

Graphs of error for consecutive thresholds of Fig. 13 and 14 show that there is a range of thresholds between the means where the error is relatively constant. This would be expected since the histograms of Fig. 13 and 14 have valleys which are also relatively flat.

Figure No.	Parts	Picture Area	Percentage Occupied by Object	Background	Border
6	a, d, g, j	4096	.505	.495	.0
	b, e, h, k	4096	.250	.750	.0
	c, f, i, l	4096	.103	.897	.0
12	a, d, g, j	3969	.499	.455	.046
	b, e, h, k	3969	.242	.726	.032
	c, f, i, l	3969	.096	.884	.021
13	a, d, g, j	3721	.485	.368	.147
	b, e, h, k	3721	.226	.671	.103
	c, f, i, l	3721	.082	.852	.066
14	a, d, g, j	3249	.456	.152	.392
	b, e, h, k	3249	.192	.532	.276
	c, f, i, l	3249	.056	.767	.177

Table 4. Fraction of picture area occupied by object, background and border for synthetic images.

Figure Part	No Blur (Fig. 6)	2x2 Blur (Fig. 12)	4x4 Blur (Fig. 13)	8x8 Blur (Fig. 14)
a	24	24	24	24
b	25	25	25	25
c	26	25	25	25
d	25	26	26	26
e	27	26	26	26
f	29	26	26	26
g	24	23	23	23
h	25	24	23	23
i	27	24	24	23
j	24	24	24	24
k	27	25	25	25
l	30	26	25	25

Table 5. Discrepancy thresholds for synthetic pictures assuming zero cost for misclassifying border points.

Assigning zero cost to misclassifying border points is unsatisfactory since we would like to choose a threshold which results in the object having a smooth border. If we assume a cutoff point k such that border points greater than k should be classified as object points and all others as background points then we can count the cost of misclassifying border points if we know the distribution of these points. For simplicity, we will assume that the gray levels in the border are uniformly distributed between the object and background means. We will choose k as the midpoint between the object and background means since we are assuming that the border is ramp-like, and border points whose gray levels are closer to the object mean should be classified as object points and those closer to the background mean as background points. We can now express P'(error|t) as

$$p_1\left\{\frac{1}{\sqrt{2\pi}\,\sigma_1}\int_{-\infty}^{t}\exp\left[-\frac{1}{2}\left(\frac{x-\mu_1}{\sigma_1}\right)^2\right]dx\right\} + p_2\left\{\frac{1}{\sqrt{2\pi}\,\sigma_2}\int_{t}^{\infty}\exp\left[-\frac{1}{2}\left(\frac{x-\mu_2}{\sigma_2}\right)^2\right]dx\right\}$$

$$+ \; p_3\left|\frac{t-k}{\mu_1-\mu_2}\right|$$

where p_1, p_2 and p_3 are the a priori probabilities of the object, background, and border distributions, respectively. The third term in the above sum is the border misclassification cost.

Using the percentages listed in Table 4, P'(error|t) was computed for a range of threshold for each picture. Graphs of these values are shown in Figs. 27-29. The discrepancy

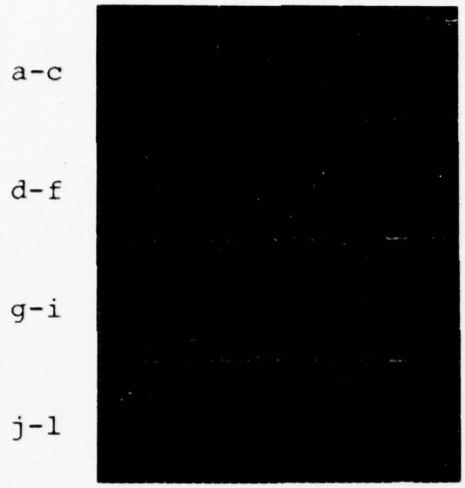

Fig. 27.

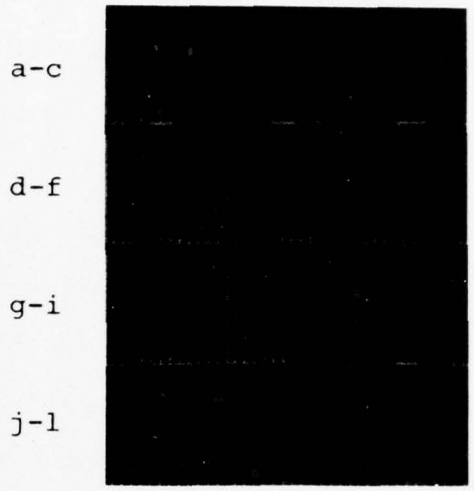

Fig. 28.

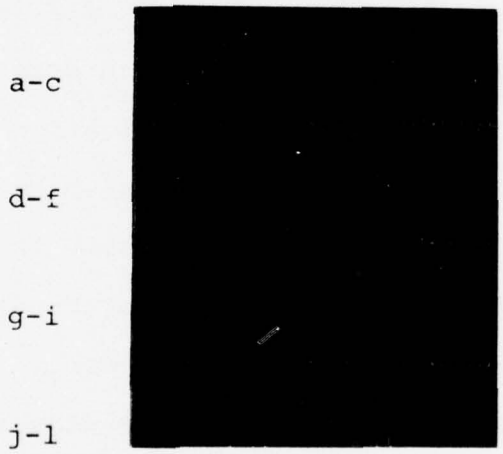

Fig. 29.

Figs. 27, 28, 29. Graphs of discrepancy (y-axis), assuming nonzero cost for misclassifying border points, vs. threshold (x-axis) for the pictures of Figs. 12, 13 and 14.

thresholds listed in Table 6 were chosen at the minima of these curves. Note that for higher degress of averaging (4x4 and 8x8) the graphs of P'(error|t) always achieve minima at 24 or 25. This is because the probability of error assuming zero cost for misclassifying border points is relatively constant over a range of thresholds including both the minimum and the midpoint between the means. Since the border misclassification cost is greater than zero everywhere except at the midpoint (25), P' is minimum at or near 25. Thus, the integer threshold occurs at either 24 (if the minimum is less than 25) or 25 (if the minimum is greater). Comparing the discrepancy listed in Tables 5 and 6 we can see that the thresholds differ by at most one gray level. Thus, the discrepancy thresholds chosen assuming that the border misclassification cost is nonzero differ little from those computed assuming that the border misclassification is zero. This result is of interest since in real blurred images it is often difficult to estimate the area of the blurred border.

If we compare the busyness thresholds in Table 3 for the 2x2 blurred images to the discrepancy threshold of Table 5 or Table 6 we see that in all cases they differ by at most one gray level. For the unblurred pictures of Fig. 6 we also found that busyness and discrepancy thresholds agreed (to within one gray level) in the majority of cases.

For the more severe cases of blur (Fig. 13 and 14) the busyness thresholds of Table 4 are always greater than

Figure Part	No Blur (Fig. 6)	2x2 Blur (Fig. 12)	4x4 Blur (Fig. 13)	8x8 Blur (Fig. 14)
a	24	24	24	24
b	25	25	25	25
c	26	25	25	25
d	25	25	25	25
e	27	26	25	25
f	29	26	25	25
g	24	24	24	24
h	25	24	24	24
i	27	24	24	24
j	24	24	24	24
k	27	25	25	25
l	30	26	25	25

Table 6. Discrepancy thresholds for synthetic pictures assuming nonzero cost for misclassifying border points.

or equal to the corresponding discrepancy thresholds of Tables 5 and 6, with differences as large as four gray levels. As already pointed out, this is because the busyness measure in these cases is influenced by the object size, so that a higher threshold yields lower busyness because it makes the object smaller.

CHAPTER IV

EXPERIMENTS WITH FLIR IMAGES

The results obtained in Chapter III for synthetic images are encouraging, but it is difficult to predict performance on real images from such results, since the normal distribution assumptions may not be satisfied by such images. In this chapter we present some results obtained on a class of real images obtained from a FLIR (Forward Looking InfraRed) sensor.

The image data base used was obtained from the U. S. Army Night Vision Laboratory, Fort Belvoir, Virginia. These images, shown in Fig. 30, each contain a dark target on a light background with a considerable amount of noise in both populations. In addition, the edges of the objects are not sharp but blurred. It can be seen by examining the histograms shown in Fig. 31 that a few of them exhibit a bimodal structure (Fig. 31b, e, f, g and ℓ), but most of them are unimodal. It is not obvious how to select thresholds for these latter cases. In fact, because of the amount of noise in the pictures, the degree of busyness in the thresholded image is very sensitive to the choice of a threshold.

Co-occurrence matrices computed on the pictures of

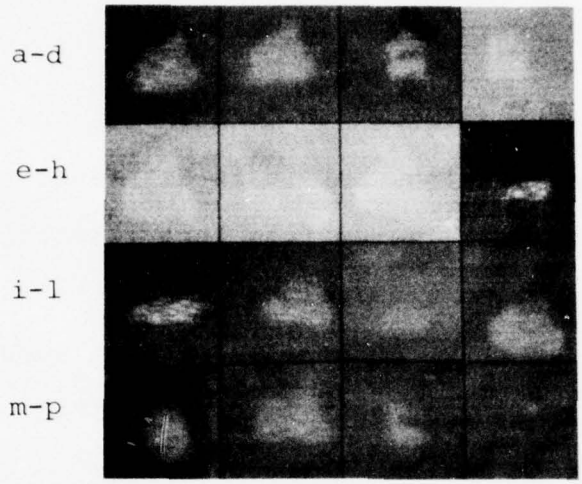

Fig. 30. FLIR images containing targets.

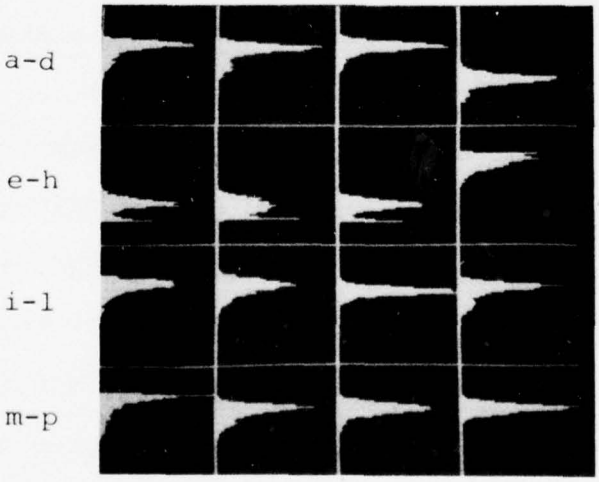

Fig. 31. Histograms of the pictures of Fig. 30.

Fig. 30 are displayed in Fig. 32 using log scaling. It should be noted that the structure of these matrices resembles that of the matrices computed for the blurred synthetic images. The black circular population on the main diagonal in each matrix represents the co-occurrence of gray levels of points in the background and accounts for the largest number of adjacencies in the image. Pairs of gray levels in the object are represented by a fuzzy population which is also centered on the main diagonal but whose boundaries are not always easily defined (see for example Fig. 32c). The smeared population between those of the object and background is due to the blur. In some of the matrices, it is possible to visually select a threshold which would partition the matrix in such a way as to minimize the resulting busyness measure. For example, in Fig. 32e there appears to be a light region along the main diagonal between the object and background populations. Selecting a threshold in this area would seem to result in the fewest number of object-background gray level adjacencies. However, when the distribution along the diagonal is cigar-shaped, as in Fig. 32m, and appears to be of nearly constant gray level, it is not clear that busyness attains a minimum other than by choosing t at one of the extreme ends of the gray scale.

Plots of busyness versus threshold are shown in Fig. 33 for the target pictures of Fig. 30. We see that for each of these pictures the shape of the busyness curve again resembles the shape of its gray level histogram.

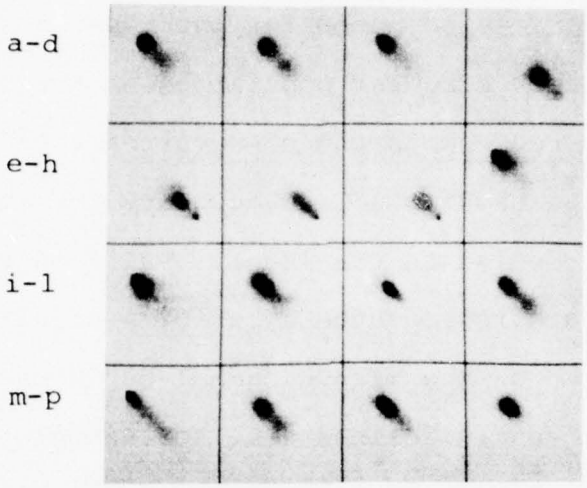

Fig. 32. Co-occurrence matrices (log-scaled) for the pictures of Fig. 30.

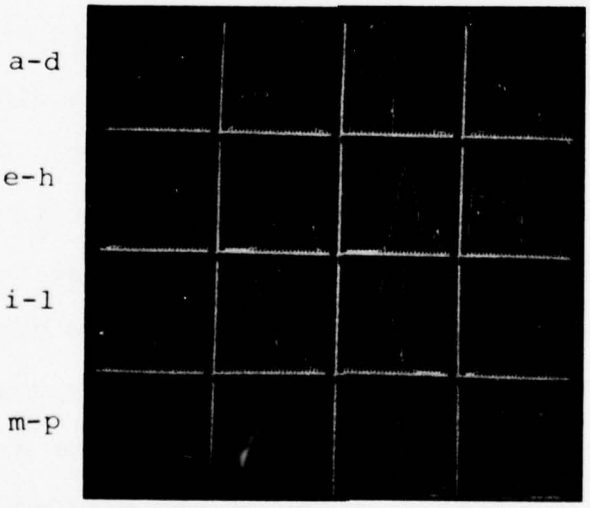

Fig. 33. Graphs of busyness (y-axis) vs. threshold (x-axis) for the pictures of Fig. 30.

The histograms and busyness curves for the pictures of Fig. 30b, e, f, g, and l are bimodal (compare Fig. 31 and Fig. 33). For these pictures we want to choose the threshold corresponding to the valley bottom on the busyness curve in order to minimize the amount of noise in the thresholded image. Each of these thresholds, which are listed in Table 7, is identical to the threshold obtained by selecting the minimum in the valley on the gray level histogram. The busyness curves in Fig. 33a, m and n contain relative minima but are not bimodal in shape. The thresholds corresponding to the relative minima for these curves are also listed in Table 7. Both the gray level histograms and the busyness curves for the remaining pictures (Fig. 30c, d, h, i, j, k, o and p) consist of a single mode and a shoulder to the right of the mode. The shoulders on the busyness curves are better defined and less noisy than those on the histograms. For these curves the left edge of the shoulder on the busyness curve is listed as the busyness threshold in Table 7 and its position is indicated by a vertical line on the graphs of Fig. 33.

For comparison with busyness thresholds, discrepancy thresholds were also computed for the pictures of Fig. 30. This involved fitting pairs of Gaussians to the histograms of the images using an algorithm described by Hartigan in [3]. The parameters of the distributions resulting from using Hartigan's programs are listed in Table 8. To evaluate visually the goodness of these fits to the histo-

Part of Fig. 30	Busyness Threshold	Discrepancy Threshold
a	27	24
b	28	26
c	26	25
d	43	45
e	46	47
f	47	47
g	47	47
h	27	26
i	27	27
j	27	28
k	26	26
ℓ	28	26
m	22	21
n	28	26
o	27	28
p	25	25

Table 7. Busyness and discrepancy thresholds for the pictures of Fig. 30.

Part of Fig. 31	μ_o	σ_o	θ	μ_b	σ_b	$1-\theta$
a	29.4	5.5	.26	20.3	2.1	.74
b	31.4	3.7	.22	21.9	2.2	.78
c	29.2	4.6	.13	20.6	2.2	.87
d	47.4	1.8	.07	38.1	2.9	.93
e	49.4	.8	.16	40.6	3.5	.84
f	48.7	.5	.16	41.8	3.5	.84
g	49.5	.7	.17	41.3	2.7	.83
h	29.0	10.0	.10	18.0	3.3	.90
i	30.2	10.0	.16	20.2	2.9	.84
j	30.1	4.1	.12	20.9	3.4	.88
k	26.2	2.6	.17	23.6	2.3	.83
l	30.9	3.7	.23	21.6	2.2	.77
m	27.9	5.1	.27	16.1	2.4	.73
n	29.8	3.7	.22	21.4	2.5	.78
o	30.7	2.8	.06	21.8	2.7	.94
p	24.5	2.1	.18	21.2	2.2	.82

Table 8. Parameters and weights of Gaussian distributions used to "fit" the histograms of Fig. 31. μ_o, σ_o, μ_b, σ_b are the mean and standard deviation of the object and background distributions, respectively; θ and $1-\theta$ are the a priori probabilities of the object and background populations, respectively.

grams, Fig. 34 shows the histograms superimposed on the normal curves. The curves fit the histograms quite well. Using the parameters listed in Table 8, the discrepancy thresholds listed in Table 7 were computed. No attempt was made to estimate the area occupied by the border populations.

Comparing the busyness and discrepancy threshold listed in Table 7 shows that the thresholds are usually close, with few exceptions. For all but one picture (Fig. 30a) the difference between the busyness and discrepancy thresholds is at most two gray levels. As can be seen from Fig. 35, the thresholds chosen all result in reasonable segmentations of the images.

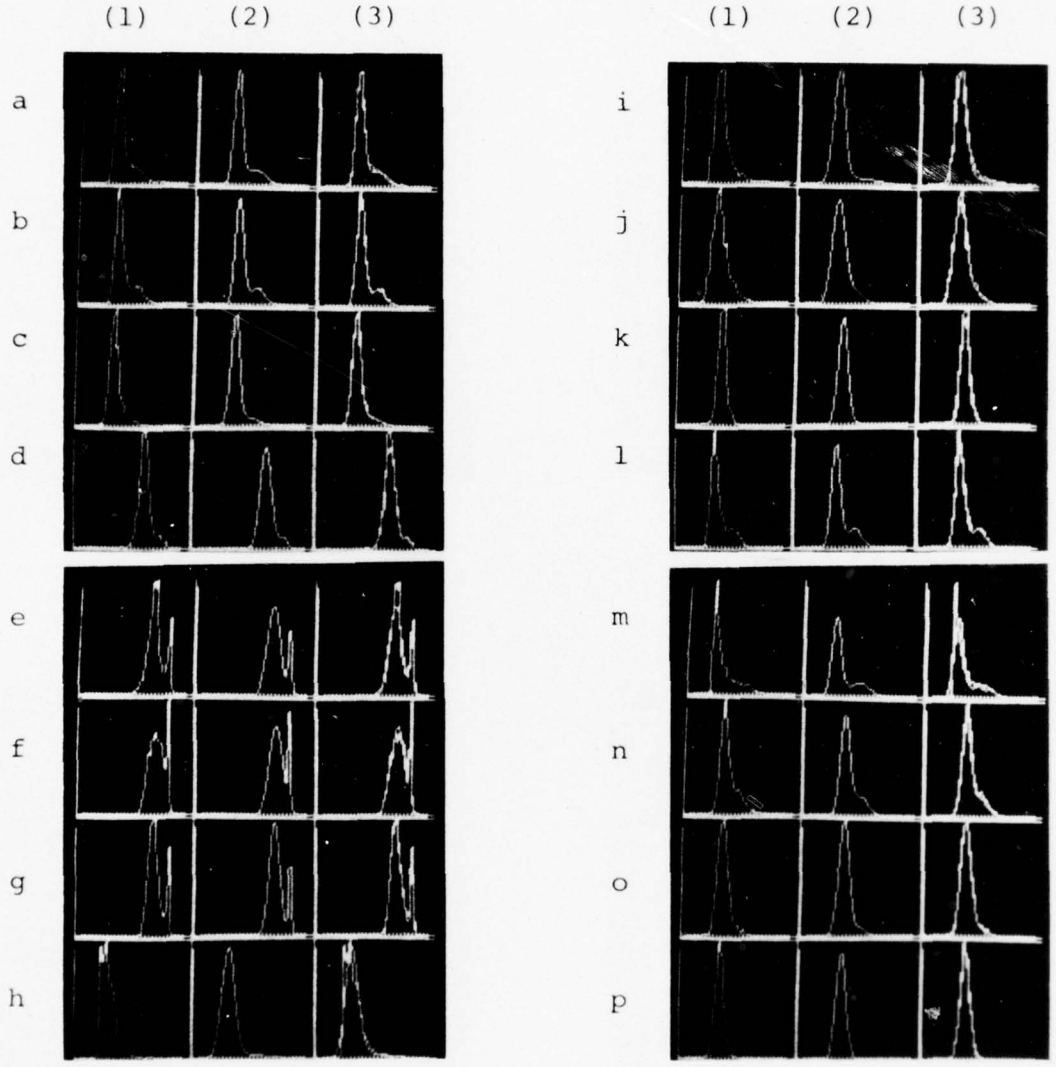

Fig. 34. Histograms, curves defined by the parameters listed in Table 8, and superimposed curves.

a-p) For the pictures of Fig. 30a-p are shown (1) gray level histograms (2) Gaussian curves using "fitted" parameters, and (3) superimposed histograms on the normal curves.

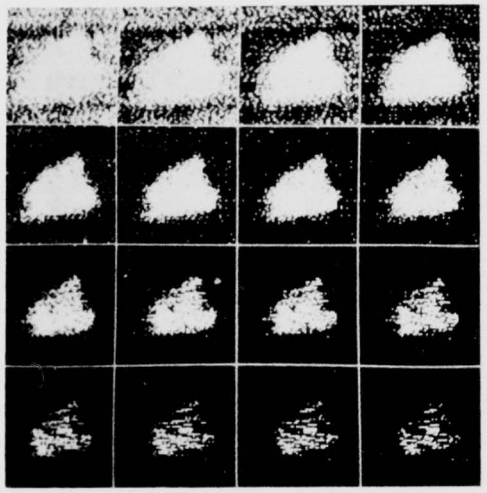

a

t = 19 to t = 34

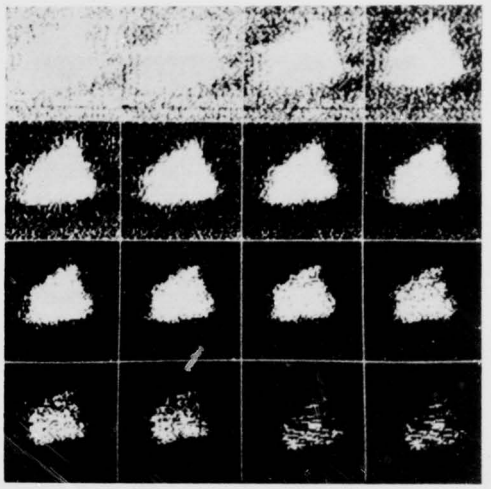

b

t = 19 to t = 34

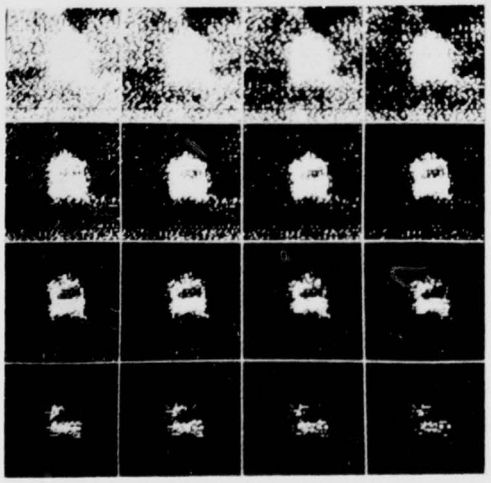

c

t = 19 to t = 34

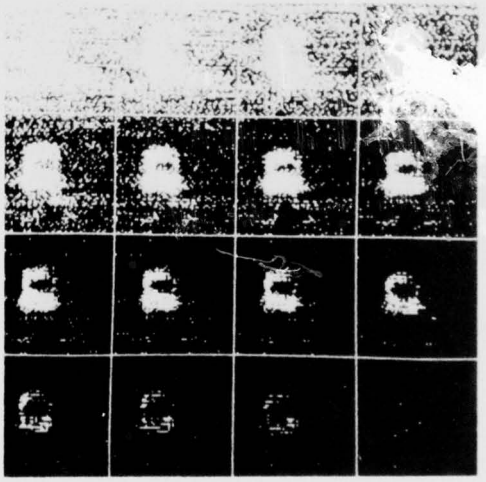

d

t = 35 to t = 50

Fig. 35a-d). Results of thresholding the pictures of Fig. 30a-d at every threshold in the ranges specified.

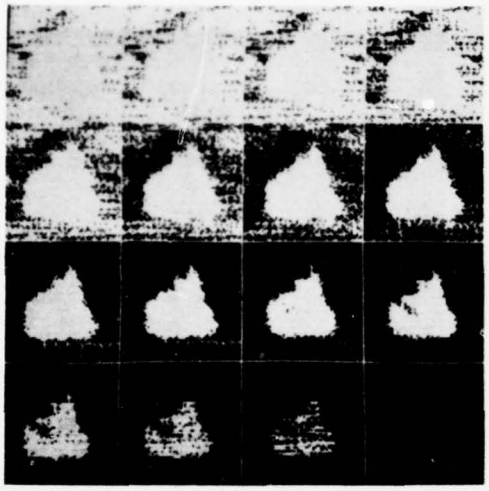

e

t = 35 to t = 50

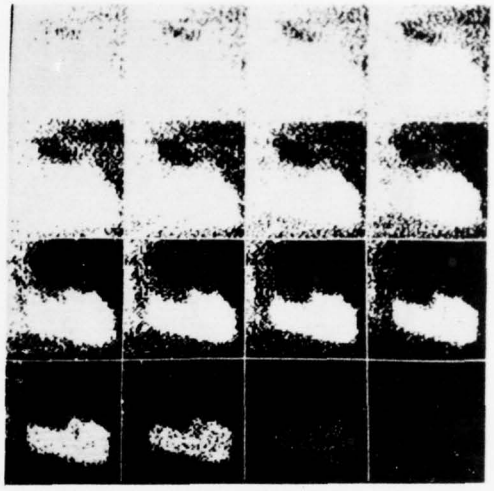

f

t = 35 to t = 50

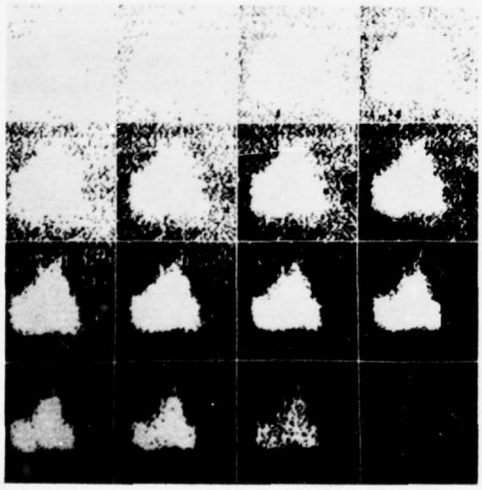

g .

t = 35 to t = 50

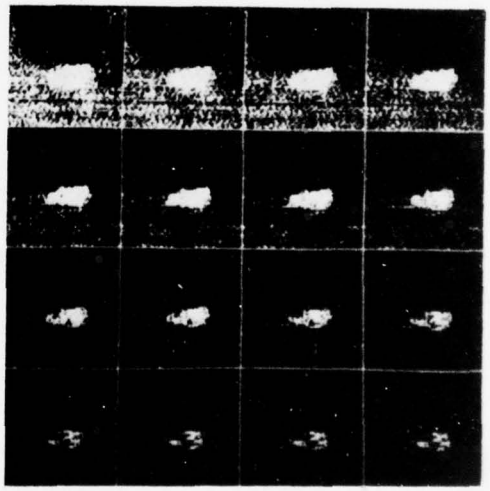

h

t = 19 to t = 34

Fig. 35 e-h). Results of thresholding the pictures of Fig. 30e-h at every threshold in the ranges specified.

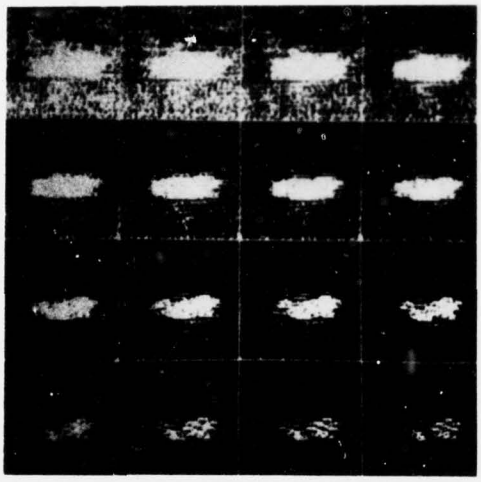

i

t = 19 to t = 34

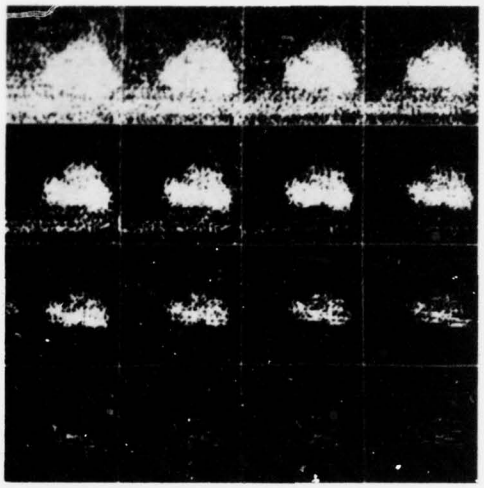

j

t = 19 to t = 34

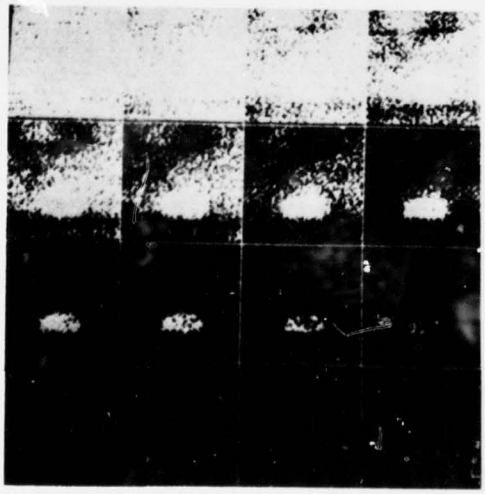

k

t = 19 to t = 34

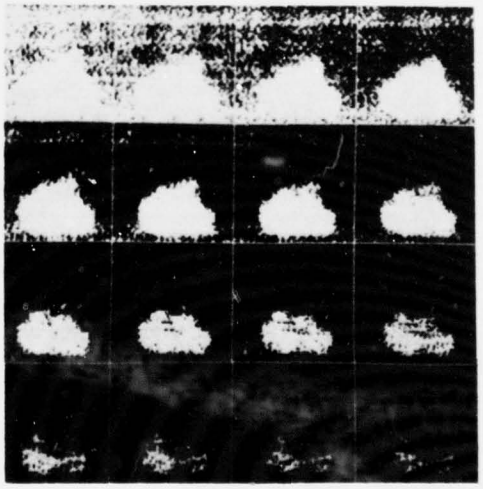

l

t = 19 to t = 34

Fig. 35 i-1). Results of thresholding the pictures of Fig. 30i-1 at every threshold in the range specified.

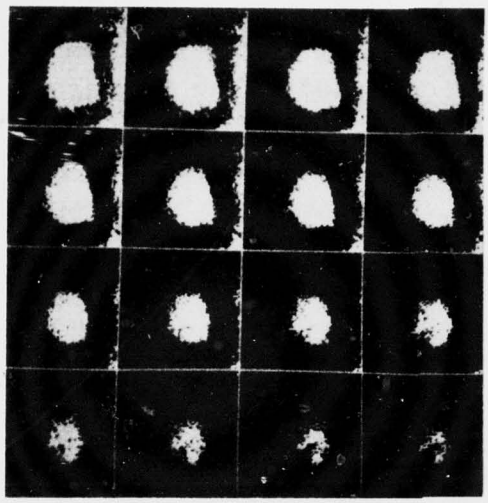

m

t = 19 to t = 34

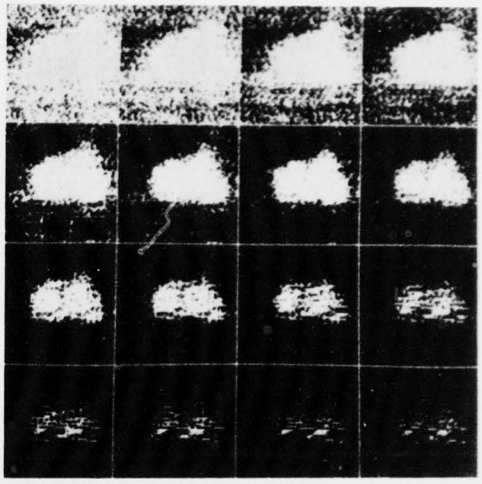

n

t = 19 to t = 34

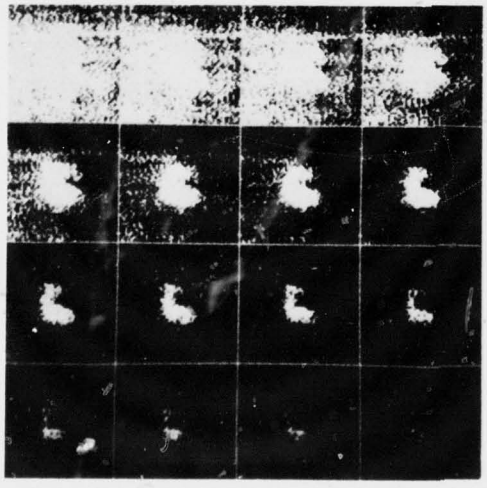

o

t = 19 to t = 34

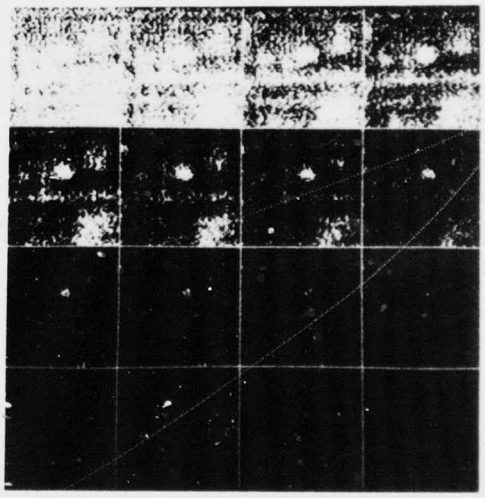

p

t = 19 to t = 34

Fig. 35 m-p). Results of thresholding the pictures of Fig. 30m-p at every threshold in the ranges specified.

CHAPTER V

CONCLUSIONS

We have compared thresholds obtained using busyness and discrepancy methods of threshold selection on both synthetic and real data. These methods yielded similar good results on unblurred synthetic images, slightly blurred synthetic images (2x2 averaging neighborhood) and real FLIR images of targets.

The similarity between thresholds obtained using these two methods is of interest since busyness thresholds are less costly to compute than discrepancy thresholds. As discussed in Section 2.1, the process of fitting normal curves to the histogram in order to obtain the parameters to be used in computing discrepancy thresholds is an expensive, iterative procedure. After the parameters have been obtained, however, the discrepancy threshold can be computed in only a few operations. The cost of computing the gray level histogram which is used as input to the fitting procedure is proportional to the picture area. Busyness measures involve calculation of co-occurrence matrices; this also requires a number of operations proportional to the picture area. Sums of entries in the co-occurrence matrix can be efficiently computed to obtain busyness measures for a range of thresholds (see Section 2.2). From these busyness measures the busyness threshold can be obtained by 1) selecting the minimum in the valley or 2) using a scheme for shoulder detection. This latter procedure is complex,

but can be done efficiently for small numbers of thresholds. Even with the additional processing required for examining busyness measures for a range of thresholds, the selection of busyness thresholds would seem to be a computationally less expensive procedure than selecting discrepancy thresholds.

The busyness and discrepancy threshold selection methods can be extended to deal with images whose histograms are multimodal. An example is provided by images of white blood cells, in which there are three types of regions: nucleus, cytoplasm, and background; this results in a tri-modal histogram. In these cases the thresholded picture should have an output level corresponding to each histogram mode. The discrepancy thresholds can be selected by fitting Gaussians to the histogram and finding intersections of the fitted curves. Busyness can be computed from the co-occurrence matrix by summing those entries corresponding to the types of busyness which are of interest. For example, if there are three output levels, then the two thresholds used partition the co-occurrence matrix into nine rectangular areas. The regions along the main diagonal represent the co-occurrences of pairs of gray levels in the objects (e.g., nuclei and cell bodies) or in the background. The six off-diagonal areas measure the contributions of three types of busyness: object 1 - background, object 2 - background, and object 1- object 2 gray level adjacencies. Appropriate thresholds can be chosen to minimize the sum of these three busyness contributions.

In summary, busyness thresholds are cheaper to compute than discrepancy thresholds and result in good image segmentations. The performance of the busyness threshold selection method makes it a viable tool for image analysis.

APPENDIX A

A SURVEY OF THRESHOLD SELECTION TECHNIQUES

This appendix presents a survey of threshold selection techniques. Most of the methods discussed deal with images comprised of two types of regions, object and background. However, several of these methods are also applicable to images containing three or more homogeneous populations.

The purpose of the methods described is to segment the image into regions which can subsequently be analyzed based on their shapes, sizes, relative positions, and other characteristics. The thresholded image also provides a representation for the image which requires less storage than the original.

In its simplest form, threshold selection involves choosing a gray level t such that all gray levels greater than t are mapped into the "object" label (denoted, say, by gray level 1), and all other gray levels are mapped into the "background" label (gray level 0).

In a more general form, a threshold operator can be viewed as a test involving a function T of the form

$$T(x,y, N(x,y), g(x,y))$$

where $g(x,y)$ is the gray level of the point (x,y) and $N(x,y)$ denotes some local property of the point (x,y), e.g., the average gray level over some neighborhood. For each point (x,y) in the image, if $g(x,y) > T(x,y,N(x,y),g(x,y))$ then (x,y) is labelled as an object point; otherwise, (x,y) is labelled as a background point.

When T depends only on g(x,y), the threshold will be called global. Simple thresholding at a specified gray level t is an example of a global threshold. If T depends on both g(x,y) and N(x,y), then the threshold chosen by T will be called a local threshold. If T depends on the coordinate values x, y as well as on g(x,y) and N(x,y), then the thresholding scheme will be called dynamic. One might, for example, use a different thresholding criterion in areas of the image which are known to be blurred. When using either local or dynamic thresholding techniques, attention must be given to the problem of artificial edges which are introduced when the threshold at one point of the image differs significantly from the threshold at an adjacent point.

We will discuss schemes for global, local and dynamic thresholding in the sections which follow.

A.1 Global Threshold Selection Based on Gray Level Histograms

Some of the earliest techniques for automatic threshold selection were global methods based on analysis of the image's gray level histogram. If the objects in an image are darker (i.e., have higher gray levels) than the background, and, in addition, occupy a fixed percentage of the picture area, then we can use a simple "p-tile" method which was suggested by Doyle [8]. This scheme chooses a threshold at the gray level on the histogram which most closely corresponds to mapping at least q% of the gray levels into the object. If, for example, dark objects occupy 20% of the picture area, then we should threshold at the 80th percentile or, more precisely, at the largest gray level allowing at least 20% of the picture points to be mapped into the object. This method is impractical if the object area is unknown or varies from picture to picture.

For the purpose of segmenting images of white blood cells, Prewitt and Mendelsohn [9] chose thresholds at the valleys (or antimodes) on the histogram. This technique is called the mode method. For four out of five of the white blood cell types the histograms were trimodal, with modes corresponding to three distinguishable areas in the images: background, nucleus, and cytoplasm. In these cases, two thresholds were chosen, at the two valley bottoms. The fifth blood cell type, basophil, produced a bimodal histogram, so a single threshold was selected in this case. The automatic selection scheme, which involved some smoothing of the histogram data, searched for modes and placed thresholds at

the minima between them.

Prewitt and Mendelsohn's method relied heavily on the bimodal and trimodal structure of the gray level histogram, which contained peaks and valleys corresponding to aspects of the image model. Object and background regions (represented by histogram peaks) were assumed to be of fairly constant gray level, and to differ in average gray level. Edges (represented by valleys) were composed of intermediate gray levels and were less heavily populated than either object or background interiors. Relationships between object regions depended on whether the object was modeled by one or more modes. Other work by Wall et al. [10] has also been done on modeling objects and their edges in an attempt to derive the structure of the corresponding gray level histogram.

A.2 Global Threshold Selection Based on Local Properties

Local property statistics have been used to aid in the selection of global thresholds in two ways. Values of local properties have been used to improve the shape of the gray level histogram by, for example, making it more strongly bimodal. This facilitates the use of global methods such as the mode method (described in Section A.1) which can then easily be applied. Methods involving statistics of local properties have also been suggested to directly compute a global threshold. The local operators employed in the methods described in the following subsections include both gradient and Laplacian operators.

A.2.1 Methods that "Improve" Histograms

Techniques described in this subsection make use of local properties to improve the shape of the gray level histogram.

Mason et al. [11] have proposed a method for making histogram valleys deeper, in order to facilitate the use of the mode method. They computed a histogram in which not all points were counted equally. The lower the value of a difference operator at a point, the more weight was given to that point. Thus, points interior to the object or background, i.e., those points whose difference values were low, were counted heavily, while those points on the border, i.e., those with high differences, were counted less heavily. The overall effect of the weighting process should be to make the histogram peaks sharper and higher, and the valley deeper so that the valley bottom can more easily be selected as a threshold. This method has been shown to derive reasonable thresholds for images, though the effect of the valley deepening is not substantial in many cases [12].

For histograms in which the valley is broad and the peaks are very unequal in size, it is often difficult to choose a threshold using the methods described in Section A.1. In these cases a technique based on a digital "Laplacian" operator can be used to produce a strongly bimodal histogram for which it is easier to select a threshold. This method was suggested by Weszka et al. [13].

The "Laplacian" operator is computed by taking absolute differences between the gray level at a point and

the average gray level in some neighborhood centered at the point. If we restrict consideration to only those points having Laplacian values in, say, the 90th percentile then we would be ignoring all points except those which lie on or near the borders of objects. For a gray level z which survives this filtering operation we would expect the probability that z lies in the object to be about equal to the probability that z lies in the background. Thus, a gray level histogram of only those points having high Laplacian values should be more symmetrical than the histogram of the entire picture. In addition, since the Laplacian is a second-difference operator which produces low values in the middle of ramps of ideal blurred edges, the center of the filtered histogram valley should be sparsely populated. This results in a deepened and sharpened valley bottom separating the two peaks. Experimental results show that the location of the valley bottom is relatively insensitive over a range of percentiles used to filter the Laplacian [14].

A.2.2 Methods that Compute Thresholds

Techniques described in this subsection make use of values of local properties to directly compute a threshold for an image.

Techniques analogous to the Laplacian based technique described in Section A.2.1 have also been developed using gradient operators. These schemes involve applying a digital gradient to an image (rather than the Laplacian) and then histogramming only those points having high gradient values. We would again expect these points to lie on or near the borders of objects, and to have a bimodal gray level histogram. Thus, the valley bottom should correspond to the gray level at which the edges of the objects are steepest. However, as we use higher and higher p-tile ranges, the histogram peaks get closer and closer together (since the points being histogrammed lie closer and closer to the edges of objects) until finally a single mode is produced by the high gradient points. Thus, if a high p-tile range is used, then whether the filtered histogram is bimodal or unimodal, the average gray level of the filtered points should be a reasonable place to threshold the image in order to separate objects from background. Katz used this method in [15]. Thresholding at the gray level corresponding to the mode on the filtered histogram (if it is unimodal) should yield results similar to those obtained by thresholding at the mean. Weszka and Rosenfeld suggested this in [16].

Watanabe [17] developed a method also based on a differencing operation p'(x,y) applied at every point p(x,y)

in the image. He computed for each gray level z,

$$d_z = \sum_{(x,y) \in s_z} p'(x,y)$$

where s_z is the set of points having gray level z. The threshold was then chosen at the level z for which d_z was highest. Since this level has a high proportion of high-difference points, it should occur just at the borders between objects and background. Watanabe achieved good segmentations with this method on a data base of cell images, but poor results were obtained by Weszka et al. [18] on images of chromosomes, handwriting, and cloud cover.

A.3 Local Threshold Selection

Automatic threshold selection techniques for use in optical character recognition systems have received much attention. Here, the thresholded output is used as input to the recognition logic, the performance of which can be used to evaluate the thresholding scheme. Since OCR systems must deal with a wide range of print quality distortions over a single document (or even over a single character), a combination of threshold operations is often used, with each operator designed to be sensitive to a different type of video distortion.

Bartz [19] described a system making use of four linear threshold operators. They were then combined to form a single threshold which was adapted to the different types of print quality distortions that occurred in the input. A threshold operator T_1 sensitive to contrast, for example, can be expressed as

$$T_1 = k_1 \bar{v} + c_1$$

where $\bar{v}$ is the average contrast over previously scanned characters, and k_1 and c_1 are optimizing parameters. The final threshold output, combining T_1 with other similarly generated thresholds, can then be applied locally or to several characters.

To handle large variations in shading which can occur even within a single character, Wolfe [20] developed a scheme for thresholding analog gray levels. In each neighborhood of a point, two operators were applied. The first operator

AD-A043 410 MARYLAND UNIV COLLEGE PARK COMPUTER SCIENCE CENTER F/G 9/2
THRESHOLD EVALUATION TECHNIQUES.(U)
APR 77 J S WESZKA
UNCLASSIFIED TR-531 AFOSR-TR-77-0980 F44620-72-C-0062 NL

2 OF 2
AD A043410

END
DATE FILMED
9-77
DDC

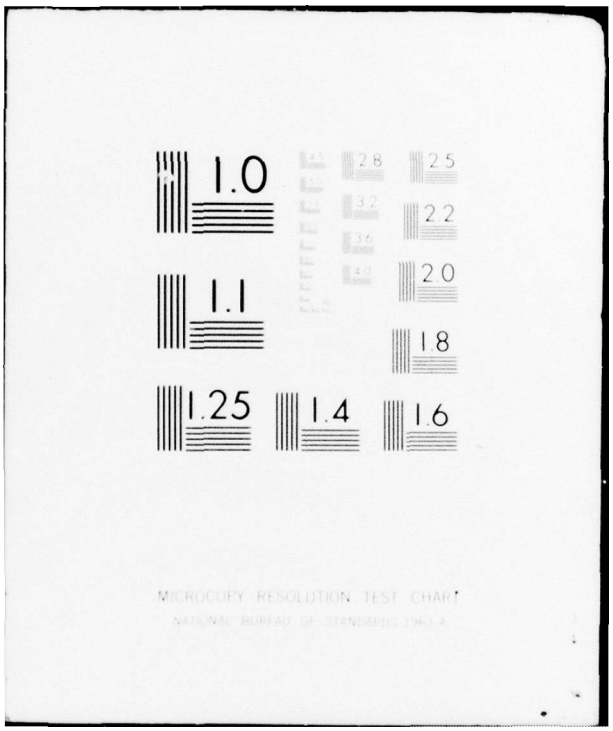
1.0
1.1
1.25
1.4
1.6
2.8
2.5
3.2
2.2
3.6
2.0
4.0
1.8
MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963 A

averaged gray levels in the image over 4x4 neighborhoods and then compared the mean gray level at a point p to those of its neighbors at distance 4 from p. If the gray level of p is darker than the gray levels of the two points whose orientations differ by 180° and which are distance 8 apart, then p is labelled as belonging to the object. The second operator is similar, but uses a larger neighborhood so that wider areas of characters could also be detected.

Another technique for selecting thresholds based on the gray levels in a neighborhood of a point was used by Ullmann [21]. Only the following points labelled n in a 5x5 neighborhood of a point p contributed to the threshold decision at p:

```
o n n n o
n n o n n
n o p o n
n n o n n
o n n n o
```

Based on experimental results, two rules were used to select a threshold at a point. The rule chosen depended on the values of the highest (whitest) gray level in the neighborhood of p. This value was denoted by $\hat{n}_p$. Rule (a) which follows was applied when $\hat{n}_p \leq 40$ and rule (b) was applied when $\hat{n}_p > 40$. The two rules are:

(a) Label p as an object point if for some point n in the neighborhood we have

$$p-n < \tau$$

where τ is some predefined threshold; otherwise, label p as a background point.

(b) Label p as an object point if for at least one of its neighbors n,

$$p < \frac{n}{\mu}$$

where μ is some predefined constant; otherwise, label p as a background point.

Notice that rule (a) uses differences between a point p and its neighbors to determine a threshold.

Two-dimensional plots of gray level versus gradient value have also been used to ease the task of threshold selection. In an early paper, Morrin [22] used these plots to convert grayscale images into black-white sketches whose resolution and contrast are superior to halftone representations. Recent work by Panda [23] involves methods in which a threshold applied at a point depends on both the gray level and edge value of the point. Several segmentation procedures were proposed based on image models suggested for FLIR images containing targets [24, 25].

The proposed models all resulted in trimodal distributions in the plot of frequency as a function of gray level and edge value. The three modes in this joint histogram correspond to image points interior to the object, interior to the background, and on the borders between the object and the background.

Segmentation procedures based on joint histogram analysis are applied to FLIR images in [23] and their results are compared with thresholded images obtained using a global thresholding scheme (the mode method). One method which was also used

by Katz selects a threshold at the mean gray level of those points having high edge values. A second scheme involves valley finding on a gray level histogram of those points having low edge values. Other more complex procedures require finding curves which follow valleys in the joint histogram. The best segmentations were obtained using a hybrid scheme that made use of the valley in the histogram of low edge value points and the mean of the high edge value points.

A.4 Dynamic Threshold Selection

Chow and Kaneko [26] used a dynamic threshold selection scheme for detecting boundaries in radiographic images. Their method was designed to perform well on low quality radiographic images in which a single global threshold is inadequate due to contrast differences throughout the image. As in the local threshold selection techniques described in Section A.2, local statistics were measured in overlapping windows of the image.(Window size 7x7 was used in the experiments.) These statistics consisted of the histogram in each window and its variance. In each window the histogram was modeled by one or two normal distributions, depending on whether the histogram was unimodal or bimodal in shape. For each histogram whose variance exceeded some predefined threshold, parameters of the component distributions were estimated. For those windows whose mixture distributions satisfied a bimodality test, minimum error thresholds were selected. Since these windows consist of mixtures of object and background populations, they should be located on the boundaries of objects. Using the thresholds selected for such windows, interpolation was employed to obtain thresholds for every point in the image. Notice that in this last step of the algorithm the thresholds were selected dynamically; that is, the value of the threshold at a point depended on its proximity to boundary points (or their neighbors) whose thresholds were determined using the local histogram analysis described above. The interpolation step was also applied to the thresholds computed for the bimodal

windows in order to ensure a gradual transition. Interpolating on a pointwise basis eliminates artificial edges at the boundaries of the windows.

APPENDIX B

ON MEASURING THE DISCREPANCY BETWEEN A THRESHOLDED IMAGE AND ITS ORIGINAL

As pointed out in Chapter II, it is not obvious how to measure the discrepancy between a thresholded image and its original, since this discrepancy depends on the output gray levels used in the thresholded image. In this appendix we investigate an approach to measuring this discrepancy.

Let μ be the mean and σ the standard deviation of g's gray level. We define the <u>α-thresholding</u> of g at t to be the image g_t given by

$$g_t(x,y) = \mu + \alpha\sigma \quad \text{if } g(x,y) \geq t$$

$$g_t(x,y) = \mu - \alpha\sigma \quad \text{if } g(x,y) < t$$

The reason for using this concept of α-thresholding, where the output gray levels depend on the grayscale of g, is that otherwise we could not fairly evaluate the discrepancy between g and g_t. Suppose, for example, that we used fixed output gray levels, say 0 and 100, and suppose that g's gray levels were all between 10 and 20. In this case, $(g-g_t)^2$ would be minimized by choosing $t \geq 20$, i.e., $g_t \equiv 0$; we could never find an interesting thresholding of g in this way.

How should the factor α be chosen? Let us consider several examples. Suppose first that the gray levels of g are uniformly distributed with mean μ and range 2r (i.e., uniformly-distributed in the interval $[\mu-r, \mu+r]$). In this case we have $\sigma = r/\sqrt{3}$. If we threshold g at μ (which, by

symmetry, should be "optimum"), and use output levels $\mu \pm \alpha\sigma$, we evidently have minimum discrepancy when $\alpha\sigma = r/2$, i.e., when $\alpha = \sqrt{3}/2$. On the other hand, suppose that g contains only the gray levels $\mu-r$ and $\mu+r$; here we have $\sigma = r$, so that if we threshold at μ (or anywhere strictly between $\mu-r$ and $\mu+r$), and use output levels $\mu \pm \sigma$, we obtain zero discrepancy.

For any probability density, if we use output levels A and B, then the minimum-discrepancy threshold is always at the mean of the output levels. To see this, let the density be f; then the mean squared error for threshold t is

$$\int_{-\infty}^{t} (z-A)^2 f dz + \int_{t}^{\infty} (z-B)^2 f dz$$

To find the t that minimizes this error, we differentiate with respect to t and equate the result to 0, obtaining

$$(t-A)^2 f - (t-B)^2 f = 0$$

or $(t-A)^2 = (t-B)^2$, so that $(t-A) = \pm(t-B)$. Here the positive root gives the absurd result A=B, whereas the negative root gives $2t = A+B$, or $t = (A+B)/2$. Thus if we use output levels $\mu \pm \alpha\sigma$, the minimum-discrepancy threshold is always at the mean.

Let us next consider the case where g's histogram is a mixture of two normal distributions with means and standard

deviations μ, ν and σ, τ, respectively, and which occupy fractions θ and $1-\theta$ of the picture area. It is not hard to verify that the mean and variance of this mixture are given by

$$m = \theta\mu + (1-\theta)\nu$$

$$s^2 = \theta\sigma^2 + (1-\theta)\tau^2 + \theta(1-\theta)(\mu-\nu)^2$$

To find the α that minimizes the mean squared error, we must differentiate

$$\int_{-\infty}^{m} (z-(m-\alpha s))^2 \, f dz + \int_{m}^{\infty} (z-(m+\alpha s))^2 f dz$$

with respect to α and equate the result to zero, where f is the mixture distribution. Since these are definite integrals that depend on the parameter α, there is no elementary approach to solving this equation. Thus as soon as we make a relatively realistic assumption about g's histogram, the problem of choosing α becomes non-elementary. For this reason, this approach to measuring thresholding discrepancy will not be pursued further.

References

1. A. Martelli and U. Montanari, Optimal smoothing in picture processing: an application to fingerprints, Proc. IFIP Congress 71, North-Holland, Amsterdam, The Netherlands, 1972, Booklet TA-2, pp. 86-90.

2. A. Rosenfeld and A. C. Kak, Digital Picture Processing, Academic Press, New York, 1976, pp. 269-273.

3. J. A. Hartigan, Clustering Algorithms, John Wiley and Sons, New York, 1975, pp. 113-129.

4. R. M. Haralick, K. Shanmugam and I. Dinstein, Textural features for image classification, IEEE Trans. Systems, Man, and Cybernetics SMC-3, 1976, 610-622.

5. R. O. Duda and P. E. Hart, Pattern Classification and Scene Analysis, John Wiley and Sons, New York, 1973, pp. 10-22.

6. Ibid., p. 40, problem 2.

7. J. S. Weszka and A. Rosenfeld, An improved method of angle detection on digital curves, IEEE Trans. Computers 24, 1975, 940-941.

8. W. Doyle, Operations useful for similarity-invariant pattern recognition, J.ACM 9, 1962, 259-267.

9. J.M.S. Prewitt and M. L. Mendelsohn, The analysis of cell images, Annals N.Y. Acad. Sci. 128, 1966, 1035-1053.

10. R. J. Wall, A. Klinger and K. R. Castleman, Analysis of image histograms, in Second Internat. Joint Conference on Pattern Recognition (IEEE Publ. 74CH-0885-4C), Copenhagen, August 1974, pp. 341-344.

11. D. Mason, I. Lauder, D. Rutovitz and G. Spowart, Measurement of C-bands in human chromosomes (preprint).

12. J. S. Weszka and A. Rosenfeld, Threshold Selection, 4, University of Maryland Computer Science Center TR-336, October 1974.

13. J. S. Weszka, R. N. Nagel and A. Rosenfeld, A threshold selection technique, IEEE Trans. Computers 23, 1974, 1322-1326.

14. J. S. Weszka, R. N. Nagel and A. Rosenfeld, A technique for facilitating threshold selection for object extraction from digital pictures, University of Maryland Computer Science Center TR-243, May 1973.

15. Y. H. Katz, Pattern recognition of meteorological satellite cloud photography, Proc. Third Symposium on Remote Sensing of Environment, Institute of Science and Technology, University of Michigan, February 1965, pp. 173-214.

16. J. S. Weszka and A. Rosenfeld, Threshold Selection Techniques, 5, University of Maryland Computer Science Center TR-349, February 1975.

17. S. Watanabe and the CYBEST group, An automated apparatus for cancer prescreening: CYBEST, Computer Graphics and Image Processing 3, 1974, 350-358.

18. J. S. Weszka, J. A. Verson and A. Rosenfeld, Threshold Selection Techniques, 2, University of Maryland Computer Science Center TR-260, August 1973.

19. M. R. Bartz, Optimizing a video processor for OCR, Proc. of the Internat. Joint Conference on Artificial Intelligence, Mitre Corp., Bedford, Mass., 1969, pp. 79-90.

20. R. N. Wolfe, A dynamic thresholding scheme for quantization of scanned images, Proc. of Automatic Pattern Recognition, National Security Industrial Association, Washington, D. C., May 1969, pp. 143-162.

21. J. R. Ullmann, Binarization using associative addressing, Pattern Recognition 6, 1974, 127-135.

22. T. H. Morrin, A black-white representation of a gray-scale picture, IEEE Trans. Computers 23, 1974, 184-186.

23. D. P. Panda, Segmentation of FLIR images by pixel classification, University of Maryland Computer Science Center TR-508, February 1977.

24. Algorithms and hardware technology for image recognition, Quarterly Report, 1 May - 31 July 1976, Contract DAAG53-76C-0138, Computer Science Center, University of Maryland.

25. Algorithms and hardware technology for image recognition, Semi-Annual Report, 1 May - 31 October 1976, Contract DAAG53-76C-0138, Computer Science Center, University of Maryland.

26. C. K. Chow and T. Kaneko, Boundary detection of radiographic images by a threshold method, Proc. IFIP Congress 71, North-Holland, Amsterdam, The Netherlands, 1972, Booklet TA-7, pp. 130-134.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER **AFOSR-TR- 77-0980**	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* THRESHOLD EVALAUTION TECHNIQUES		5. TYPE OF REPORT & PERIOD COVERED: Technical
		6. PERFORMING ORG. REPORT NUMBER: TR-531
7. AUTHOR(s) Joan S. Weszka		8. CONTRACT OR GRANT NUMBER(s): F44620-72C-0062
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Science Ctr. Univ. of Maryland College Park, MD 20742		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Math.& Info. Sciences, AFOSR/NM Bolling AFB Wash., DC 20332		12. REPORT DATE: April 1977 13. NUMBER OF PAGES: 102
14. MONITORING AGENCY NAME & ADDRESS*(if different from Controlling Office)*		15. SECURITY CLASS. *(of this report)*: UNCLASSIFIED 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT *(of this Report)*

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

18. SUPPLEMENTARY NOTES

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

Image processing
Pattern recognition
Segmentation
Thresholding

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

Threshold selection techniques have been used as a basic tool in image segmentation, but little work has been done on the problem of evaluating a threshold of an image. This paper addresses the problem of threshold evaluation and proposes two methods for measuring the "goodness" of a thresholded image, one based on a busyness criterion and the other based on a discrepancy or error criterion. These evaluation techniques are applied to both synthetic and real images and are shown to be useful in facilitating threshold selection.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

ABSTRACT: (continued)

In fact, both methods usually result in similar or identical thresholds which yield good segmentations of the images.