

## Document image binarization: Evaluation of algorithms

Paul W. Palumbo, Puducode Swaminathan, and Sargur N. Srihari  
Department of Computer Science  
226, Bell Hall, State University of New York at Buffalo  
Amherst, NY 14260

### Abstract

The extraction of a binary image from a gray level image is a common image processing operation particularly for document image analysis and optical character recognition. Various methods for this task are described in the literature including global and adaptive binarization. This paper evaluates three adaptive binarization techniques viz., a contrast measure approach, a weighted running average approach and a second derivative approach, and compares them to global binarization methods. Experiments with noisy document (postal letter mail) images lead to the following conclusions: Image contrast binarization often yields nearly the same results as the edge operator, with considerably less computation and is less sensitive to parameter settings. In addition, the edge operator is more sensitive to image resolution than the contrast operator. The weighted running-average approach is highly sensitive to the parameters involved in the calculation of the average but produces a quick binarization.

*Keywords:* Document Image Analysis, Image Binarization, Edge Detection, Marr - Hildreth operator, Image Contrast.

### 1. Introduction

This paper concerns the task of document image binarization which is a preprocessing step in document image analysis and optical character recognition. The class of documents is assumed to be those that are printed with dark ink on white paper. However, some mention is made of images which deviate from this typical class of documents.

An optically scanned document image is usually an integer (gray-valued) array; it is obtained by a process of spatial sampling and simultaneous conversion of light photons to electric signals. A high resolution *bit-map* can be sufficient to capture the shades of gray in the eye of the human perceiver (such images are called *half-tone*). An optically scanned document image is converted into a bit-map by a *binarization technique*.

A document is generally synthesized (and meant to be perceived) as black object on a white background. The underlying assumption used to binarize most document images is that the original document page contains only two intensity levels but when the document was digitized, a gray level image of the page was produced. Given this gray level image of a binary scene, a binarization technique can be used to convert the input from a multivalued image to a more appropriate binary image. In effect, a bottom-up analysis of the image is performed where each pixel is examined and converted to either black or white based on its gray level intensity value. Combined with the bottom-up pixel classification of the image to its perceived form is the more ideal top-down segmentation of the image into a symbolic form where each object on the page is represented by a connected component. Using this top-down approach, each point in the image should be converted to black if and only if that point is a part of the foreground (or goal region) of the image. This top-down binarization technique would produce an image space where each black point in the binary image would represent the foreground (or goal region) in the corresponding gray level image. Seeing that the goal directed decision of whether one pixel is a part of the foreground or background is so difficult, the much easier bottom-up pixel intensity classification is usually used.

Without the idealized assumption that the document page contains only black and white information, the problem is more complicated. Some document images may contain not only black text on a white background but may also contain objects of various colors which when captured by a gray level digitizing device may be represented as various shades of gray. To binarize these colored images, the colored object of interest must be converted to black in the binary image as well as convert the originally black text in the image to the same black in the binary images. A similar problem occurs when binarizing an image which contains various shades of gray.

There are many ad-hoc techniques to do binarization. This paper describes and evaluates several binarization techniques. Section 2 describes global thresholding which is useful in simple situations but cannot be applied to images which have backgrounds of varying gray levels such is present in colorful and poorly illuminated image. Local adaptive thresholding and specific versions of it are described in Section 3. Section 4 contains details of evaluation of the adaptive binarization techniques. Section 5 concludes the paper with a discussion of the relative merits and demerits of the each locally adaptive technique.

### 2. Global thresholding

Many of the global thresholding techniques determine one threshold for converting the entire image from a range of gray-values to two values (black and white). If a pixel in the image is above that global threshold then it is converted to white (background) otherwise it is converted to black (foreground). The determination of this global threshold could be made either using histogram properties<sup>9,11</sup>, busyness properties<sup>1,10</sup>, global edge information<sup>4,5</sup> or determined a priori by the user. A good review of several global

binarization techniques is presented by J. Kittler, et. al.<sup>4</sup>.

Assuming the object is darker than the background, point (x,y) is part of the object if  $I(x,y) < T$ , where T is a threshold and  $I(x,y)$  is the gray level intensity of the image at point (x,y); otherwise it is part of the background. A similar rule can be used if the object is assumed to be brighter than the background. A good heuristic to pick the threshold T in the bimodal case is to search the histogram of gray levels and find an intensity value separating the two peaks; typically a high peak corresponds to the background (white) and an intermediate peak corresponds to the print (black). Finding the valley between the right peaks of a histogram can be difficult when the histogram does not contain a single pure valley between two peaks. When performing a global thresholding of the image, only one threshold parameter can be adjusted to improve the binarization. With typically only 256 values for this parameter, only limited improvements in the binarization can be achieved.

Global thresholding is useful in simple situations. A common problem with this technique occurs when the gray level transition between the object and the background is gradual and inconsistent across objects in the image. In the area of postal envelope images [Figures 1(a)-(b)], it is indeed the case that backgrounds are of varying gray levels. Many of them exhibit properties which are not ideal, including poor illumination and color fields, ruling out the adoption of global thresholding.

The problem of choosing a global threshold value is brought out in Figure 2. Figures 2(a)-(b) shows a row slice (profile) for the input images shown in Figures 1(a)-(b) respectively. This figure shows each pixel in a particular row plotted against its corresponding gray value. It can be noticed that a wide variation in gray value exists between neighboring pixels in the same row. Also, foreground and background regions of the image are comprised of similar gray values but are locally identifiable.

### 3. Local adaptive thresholding

All global binarization methods examine only the pixel to be converted when determining the subsequent binary value of that pixel. For locally adaptive thresholding of an image, not just the value of the pixel to be binarized is examined but a neighborhood of pixels around this pixel is examined. If the pixel is significantly darker than the neighboring pixels, it is converted to black; otherwise it is converted to white. Using this neighborhood information, the binary value for the pixel under consideration can be determined based upon local conditions and not just on a single threshold calculated from an image-wide standpoint. Using locally adaptive binarization techniques, a more top-down knowledge-based approach is utilized to provide a better binarization. Several local adaptive thresholding techniques have been developed based on a weighted running average<sup>12</sup>, local contrast measure<sup>2</sup> and a local edge based approach<sup>7</sup>.

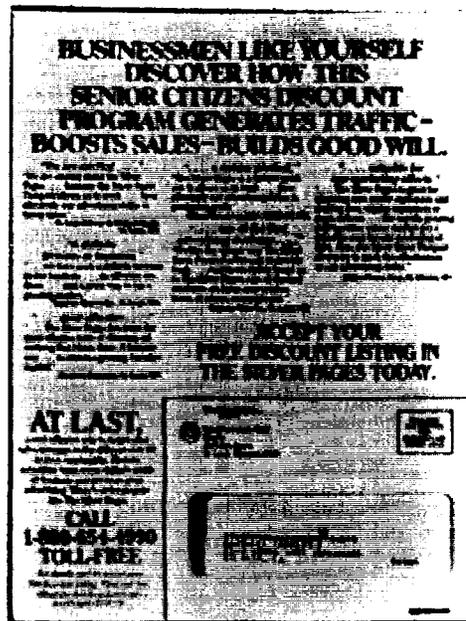
#### 3.1 Running average thresholding

White and Rohrer<sup>12</sup> describe a *dynamic thresholding algorithm* where the black/white decision is determined by a threshold which is continually changed as the scanned gray scale data stream changes. The basic threshold calculation may be represented (in each dimension) as a first order difference equation with a nonlinear coefficient determined by heuristics.

Given a scan line of a gray scale image represented by the array  $I(n)$ , the horizontal running average,  $y(n)$ , for each row up to column n can be calculated by:

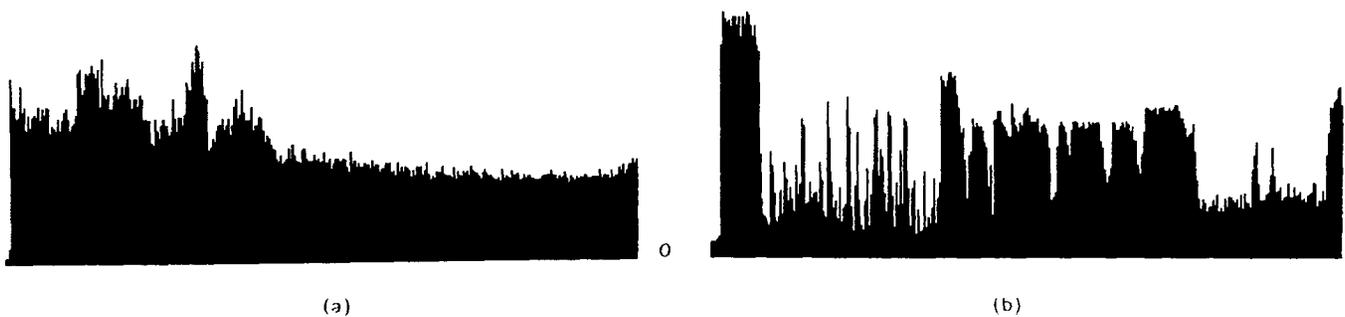
$$y(n) = y(n-1) + f[I(n) - y(n-1)],$$

where  $y(n-1)$  represents the weighted running average calculated for the current scan line up to and including the previous point in the row. The weighting function  $f[\delta]$  returns a value corresponding to the change in the average between successive points in a scan line. This weighting function should have a value of zero only when its argument has a value of zero, and otherwise it should



(a) (b)

Figure 1(a)-(b) Input mail images for the experimental evaluation.



Figures 2(a), and 2(b) - A profile of a single row from each of the images shown in Figures 1(a), and 1(b). The significant gray level variation between adjacent pixels makes it difficult to choose a single global threshold value.

have a value between zero and the value of its argument.

A similar running average can be calculated for the vertical direction. This vertical average utilizes the horizontal running average,  $y(n)$ , from above and is calculated using the following formula:

$$z(n) = z(n-L) + g[y(n) - z(n-L)],$$

where  $z(n-L)$  represents the vertical running average calculated using the previous scan line of image data. The updating function,  $g[\text{delta}]$ , is the vertical updating function constructed similarly to  $f[\text{delta}]$  described above.

The final judgement as to whether a gray level value  $l(n)$  in a particular scan line is to be converted to black or white is as follows:

$$\text{if } h[l(n-N)] > z(n) \text{ then BLACK else WHITE,}$$

where  $h[\text{pixel}]$  is a biasing function which adjusts the gray value of the pixel under consideration. The constant  $N$  causes a look ahead factor in the binarization calculation. This formula performs a comparison of a biased gray valued point against the a weighted running average of the image scanned so far as well as a few points still unclassified on the present scan line (look ahead). This technique is an efficient one-pass means for binarizing a gray level image.

### 3.2 Contrast measure thresholding

An adaptive thresholding scheme was developed at Elsag by Giuliano, Paitra and Stringa<sup>2</sup> for character recognition applications. The intensity  $I(x,y)$  of a point  $(x,y)$  in the input image is functionally compared to the intensities of pixels in a  $9 \times 9$  neighborhood centered at the pixel and a binary value  $Z(x,y)$  is output. The neighborhood is constructed as shown in Figure 3. Area  $A_1$  contains the pixels in the  $3 \times 3$  neighborhood about  $(x,y)$  and area  $A_2$  contains the four  $3 \times 3$  neighborhoods diagonally adjacent to  $A_1$ . This mask arrangement is presumably constructed so as to capture foreground pixels in area  $A_1$  while the background is captured in area  $A_2$  and vice versa. Pixels in  $A_1$  are most significant in the comparison to determine an output value, those in  $A_2$  are less significant and those in the remaining area of the input are not considered at all. A version of this scheme is the adaptive method considered in our experimentation.

The thresholding is done using the following steps. Predetermined static thresholds  $t_1$  and  $t_2$  are set so that intensities less than  $t_1$  are always black (similar to a global threshold). The subset  $A_{2t}$  of pixels in  $A_2$  whose intensities exceed  $t_2$  are determined. If the average intensity of pixels in  $A_{2t}$  exceeds the average intensity of pixels in  $A_1$  (weighted by parameters  $t_3$ ,  $t_4$  and  $t_5$ ) then the pixel is considered black; otherwise it is white. The algorithm is formally,

```
begin
  if ( $I(x,y) < t_1$ ) then  $Z(x,y) := \text{black}$ ;
  else
    begin
       $a_1 := \text{Average of 9 pixels in area } A_1$ ;
       $A_{2t} := \{(x,y) | (x,y) \text{ in } A_2 \text{ and } I(x,y) > t_2\}$ ;
       $a_2 := \text{Average of pixels in } A_{2t}$ ;
      if ( $(t_3 * a_2 + t_5) > (t_4 * a_1)$ ) then  $Z(x,y) := \text{black}$ ;
      else  $Z(x,y) := \text{white}$ ;
    end;
  end.
```

Thus it can be seen that the threshold level is dynamically varied depending upon the neighborhood values of the pixel under consideration. The algorithm also uses simplistic global thresholding to classify extreme cases and reduce the computational complexity.

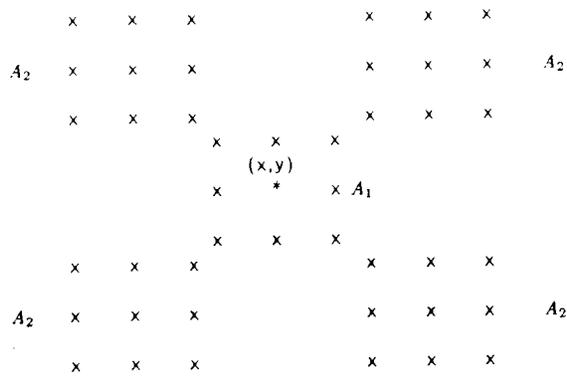


Figure 3 - Image point  $(x,y)$ , nearby  $3 \times 3$  neighborhood  $A_1$ , and four diagonally adjacent neighborhoods  $A_2$ .

### 3.3 Edge-based thresholding

#### Theory of Edge Detection

An edge or boundary in an image is the place where there is a more or less abrupt change in gray level. At the edge between two surfaces, brightness should change sharply from one value to another, but in practice the steplike change is generally corrupted. One way to detect such noisy edges involves four steps<sup>13</sup>. The first step is to make an average brightness array from the image. Local brightness averages reduce the influence of noise. The following one dimensional formula illustrates the computation that needs to be done. The local average brightness  $A_i$  around any point  $i$  is given by

$$A_i = (I_{i-1} + I_i + I_{i+1}) / 3,$$

where  $I_i$  is the image brightness at point  $i$ . Now, an average first difference array is made from the average brightness array. To do this, we average the right neighbor difference  $A_{i+1} - A_i$  and the left neighbor difference  $A_i - A_{i-1}$ . The resultant average  $F_i$  is given by

$$F_i = ((A_{i+1} - A_i) + (A_i - A_{i-1})) / 2 = (A_{i+1} - A_{i-1}) / 2.$$

Thus  $F_i$  is a finite difference approximation to the *first derivative*. Next, an average second difference array is made from the average first difference array. To do this, we average the average first difference array. This average second difference of average brightness  $S_i$  is given by

$$S_i = ((F_{i+1} - F_i) + (F_i - F_{i-1})) / 2 = (F_{i+1} - F_{i-1}) / 2.$$

Thus  $S_i$  is a finite difference approximation to the *second derivative*. Note that the first differencing procedure transforms smoothed steps into bumps. And the second differencing procedure transforms bumps into S shaped curves that cross zero steeply between positive and negative peaks. It can thus be observed that intensity changes in an input image correspond to maxima of the first derivative (or gradient) of the image surface or equivalently a place at which the second derivative crosses zero and changes sign. First and second derivative operators can be roughly realized by filters with receptive fields which are directionally selective. By contrast, the Laplacian operator, which is an approximation for  $\partial/\partial x + \partial/\partial y$ , is rotationally symmetrical and as such is non-directional. Most algorithms for implementing these operators have two stages: first the image is *convolved* with the operator (in sufficiently many directions if the operator is directional) and second, the output is analyzed to locate the orientation and extent of the intensity changes in the input. There have been many practical implementations of the operators described above, one of which is the Marr-Hildreth operator.

#### The Marr - Hildreth operator

Marr and Hildreth<sup>7,8</sup> point out that "a major difficulty with natural images is that changes can and do occur over a wide range of scales, so it follows that one should seek a way of dealing with the changes occurring at different scales". No single filter can be optimal simultaneously at all scales, hence one should seek a way of dealing separately with the changes occurring at different scales. There are two physical considerations that combine to determine the appropriate smoothing filter. The first is that the motivation for filtering the image is to reduce the range of scales over which the intensity changes take place. The filter spectrum should therefore be smooth and roughly bandlimited in the frequency domain. In other words, the filter should be localized in the frequency domain. The second consideration is the constraint of spatial localization. The things in the world that give rise to the intensity changes in the image can all be thought as spatially localized. Hence the filter we seek should be localized in the spatial domain too. Marr and Hildreth argue that a Gaussian filter optimizes localization in both domains simultaneously and so it is chosen as the band limiting filter. The Gaussian smoothing is followed by detection of zero crossings with a rotationally symmetric Laplacian operator. Being nondirectional, these are computationally more efficient than other directional operators.

The Marr-Hildreth approach to edge finding involves convolving the image with an operator which is the Laplacian of a Gaussian (or  $\nabla^2 G$  as it is known).  $\nabla^2 G$  is a circularly symmetric operator also called the *Mexican hat operator*<sup>13</sup> with its central peak and

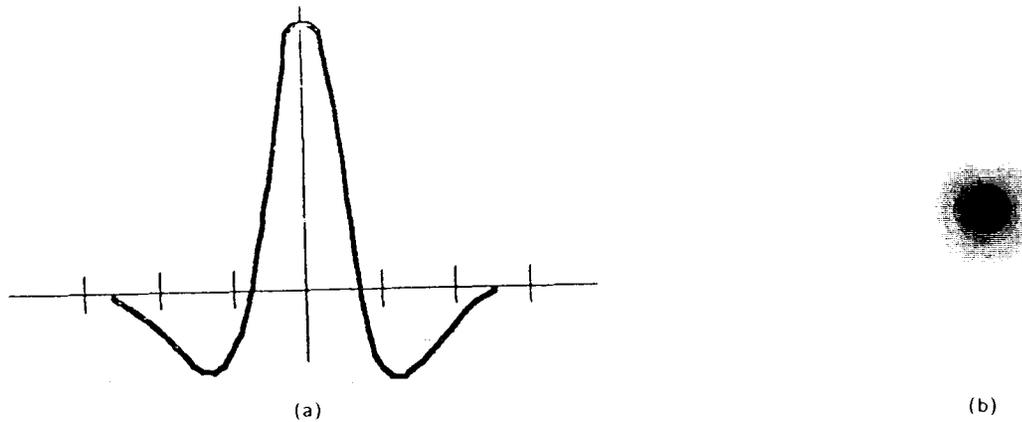


Figure 4(a) -  $\nabla^2G$  shown as a one dimensional function. 4(b)  $\nabla^2G$  shown as a two dimensional function.

negative lobes, whose distribution in two dimensions may be expressed in terms of the radial distance  $r$  from the origin by the formula

$$\nabla^2G(r) = \frac{-2}{\pi\sigma^4} \left(1 - \frac{r^2}{2\sigma^2}\right) e^{-\frac{r^2}{2\sigma^2}}$$

where  $\sigma$  is the space constant of the Gaussian. Figure 4 illustrates the one and two dimensional forms of this operator. The diameter of its positive central region,  $w$ , varies linearly with the space constant  $\sigma$  of the underlying Gaussian  $w = 2\sqrt{2}\sigma$ . The Gaussian part blurs the image effectively wiping out all structures at scales much smaller than the space constant of the Gaussian. This fact will have to be taken into consideration while choosing the space constant for actual implementation. Hence, in practice, the most satisfactory way of detecting intensity changes in an image is to first filter it with the operator  $\nabla^2G$  where the space constant of  $G$  is chosen to reflect the scale at which the intensity changes are to be detected and then to locate the zero crossings in the filtered image.

To determine the correlation between the Gaussian space constant  $\sigma$  and the effective part of the Gaussian several calculation must be made. That is, the operator size of the Marr-Hildreth operator should capture a significant portion, say 95%, of the absolute area under the  $\nabla^2G$  curve. After performing the mathematical integration, the area under both the positive and the negative portions of the  $\nabla^2G$  curve is:

$$\frac{4}{e\sigma^2}$$

The are under the negative portion of the curve up to radius  $r$  ( $r \leq \sqrt{2}\sigma$ ) is:

$$\frac{4}{e\sigma^2} - \frac{2r^2}{\sigma^4} e^{-\frac{r^2}{2\sigma^2}}$$

After performing the necessary calculations, the operator size to capture 95% of the effect of the Marr-Hildreth operator should be set to 6.2 times the Gaussian space constant  $\sigma$ .

### Binarization

Marr<sup>6</sup> suggested the following method of binarizing an image. Compute the convolved image using a second derivative operator such as the Mexican hat operator and Assign white to positive values and black to negative values. The critical fact about the operator  $\nabla^2G$  is that its zero-crossings mark the intensity changes, as seen at the Gaussian's particular scale. The numerical values in the  $\nabla^2G$ -filtered image are both positive and negative. Just to one side of the zero-crossing will lie a peak positive value and just to the other side, a peak negative value. The positive part represents that side of the edge where the intensity is high and the negative part corresponds to the side where the intensity is low. Hence, it is reasonable to perform thresholding based upon the above criterion.

The Marr - Hildreth operator accepts two parameters - (i) the  $\sigma$  value (space constant) for the Gaussian, and (ii) the size of the operator. A sample Marr - Hildreth operator with actual values is shown in Figure 5. Based on the  $\nabla^2G$  operator function, a mask of the required size is calculated and the input image is convolved with it. The resultant output has zero crossings precisely at those points in the image where there are sharp discontinuities in the intensity. This fact enables thresholding to be done based on the following criteria - if the pixel intensity is negative, it is made black (0); and if the pixel intensity is positive, it is made white (255) - thus performing the conversion into a black and white image.

The operator size will be set, as described above, to be 6.2 time the Gaussian space constant  $\sigma$ . As we decrease the  $\sigma$ , additional zero crossings reflecting noisy edges are introduced. The space constant  $\sigma$  for the Gaussian will have to be chosen taking into account the scale at which the intensity changes have to be detected. In other words, it determines the width of the thinnest edge that will be detected by the operator.

-0.004	-0.018	-0.050	-0.087	-0.104	-0.087	-0.050	-0.018	-0.004
-0.018	-0.073	-0.168	-0.242	-0.261	-0.242	-0.168	-0.073	-0.018
-0.050	-0.168	-0.270	-0.154	-0.015	-0.154	-0.270	-0.168	-0.050
-0.087	-0.242	-0.154	0.588	1.154	0.588	-0.154	-0.242	-0.087
-0.104	-0.261	-0.015	1.154	2.000	1.154	-0.015	-0.261	-0.104
-0.087	-0.242	-0.154	0.588	1.154	0.588	-0.154	-0.242	-0.087
-0.050	-0.168	-0.270	-0.154	-0.015	-0.154	-0.270	-0.168	-0.050
-0.018	-0.073	-0.168	-0.242	-0.261	-0.242	-0.168	-0.073	-0.018
-0.004	-0.018	-0.050	-0.087	-0.104	-0.087	-0.050	-0.018	-0.004

Figure 5 - Marr Hildreth operator mask of size 9x9 for a  $\sigma$  value of 1.4.

## 4. Experimental evaluation

An experimental evaluation of all three adaptive binarization techniques was done in the context of address block detection in images of postal letter mail. These images are inherently noisy and do not represent ideal documents. They also do not conform to any specific standard. As such it is not possible to use any a priori knowledge of the input layout. Hence, it is impractical to use any global binarization technique that has preset thresholding parameters which depend on the nature of the input image. What is required is a technique that can dynamically change after taking a look at a neighborhood of the image point to be classified. This experimental evaluation was done keeping this prerequisite in mind. Each algorithm was evaluated based on the results of over 10 images with a resolution of about 128 pixels per inch, some of which are shown in Figure 1.

### 4.1 Running average binarization

The performance of the running average binarization method on the images in Figure 1 can be seen in Figures 6(a)-(b). The algorithm is a fast one pass approach to binarization. The resulting binary images, however, are adequate but was unable to tolerate some of the noise present in the original image. This problem is primarily a result of binarizing a noisy point in the image which is darker than the current running average (even after the adjustment produced by the biasing function,  $h()$ ). Since this running average approach just compares the biased pixel value against the running average, this noisy point would be classified as black even when the surrounding points are all white. Also, once this running average binarization is tuned to a particular image resolution, the performance of the method deteriorates when images of other resolutions are used. This difficulty results from the vertical, and horizontal weighting functions  $f()$  and  $g()$  being resolution dependent. Once these adaptation functions are created for a given image resolution, they must be changed when the image resolution changes. Also, in a region which contains only background pixels, the overall running average will tend to converge on the average background gray value. This means that when comparing image points in these large background region against the nearly similar overall running average, noise will be introduced into the binary image. In general, this adaptive binarization algorithm is very fast but suffers under the addition of noise in the image.

### 4.2 Contrast measure binarization

The results of applying the adaptive thresholding to the input images shown in Figures 1(a)-(b) are given in Figures 7(a)-(b). A constant set of values for the five parameters  $t_1$ ,  $t_2$ ,  $t_3$ ,  $t_4$  and  $t_5$  was used on all the images. The values used were :  $t_1 = 20$ ,  $t_2 = 20$ ,  $t_3 = 0.85$ ,  $t_4 = 1.0$  and  $t_5 = 0$ . The choice of these parameters were fairly straightforward (with minimal experimentation).

As can be observed, the results are satisfactory in that the characters are clearly distinguished from the background. Computationally it is not too expensive. For a 512x512 image with the above parameter setting, it required approximately 2 minutes of CPU time, which in real time becomes 10 minutes under normal system load on VAX-11/780.

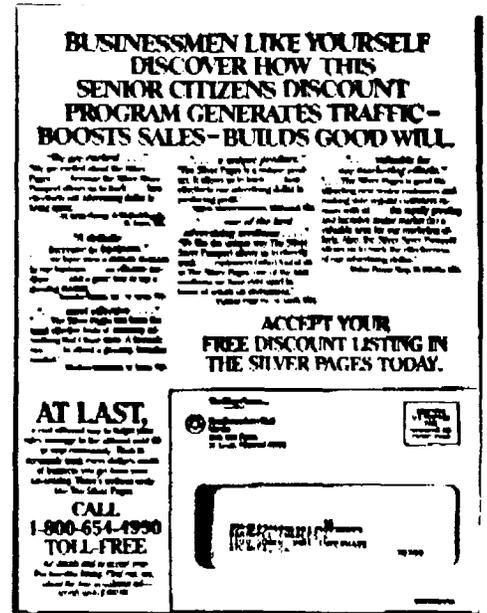
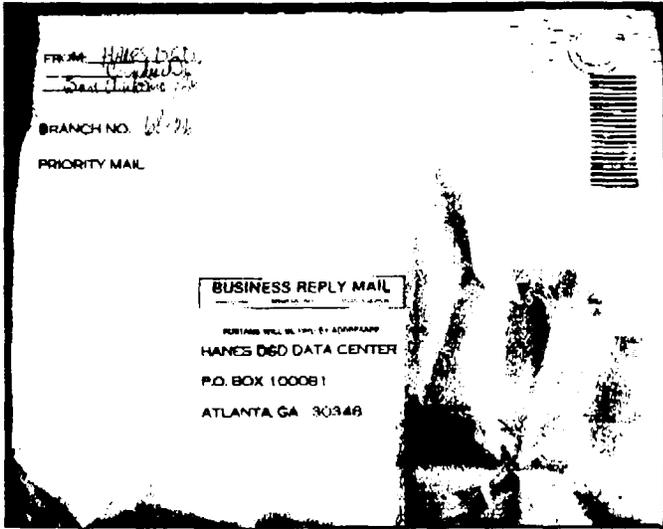
The only (minor) problem with the operator is at lower resolution horizontal (and at times vertical) smearing is visible in the binary output. This seems to take place due to the fact that we are always considering a 3x3 neighborhood for arriving at the threshold value; in a situation involving edges present inside the neighborhood the algorithm averages out the edge resulting in the extra smearing. It can be noted that the results obtained with medium resolution images do not suffer from this.

### 4.3 Second derivative binarization

The results of applying the Marr - Hildreth operator and subsequent binarization on the images of Figures 1(a)-(b) are shown in Figures 8(a)-(b). The operator used had a size of 19x19 and the  $\sigma$  value was 3.0. The results obtained indicate that the choice of the correct  $\sigma$  value for the operator given a fixed image resolution is crucial for obtaining good results. A small deviation of  $\sigma$  results in the introduction of a large number of noisy points in the output. The value for  $\sigma$  was chosen to capture intensity changes with a particular magnitude.

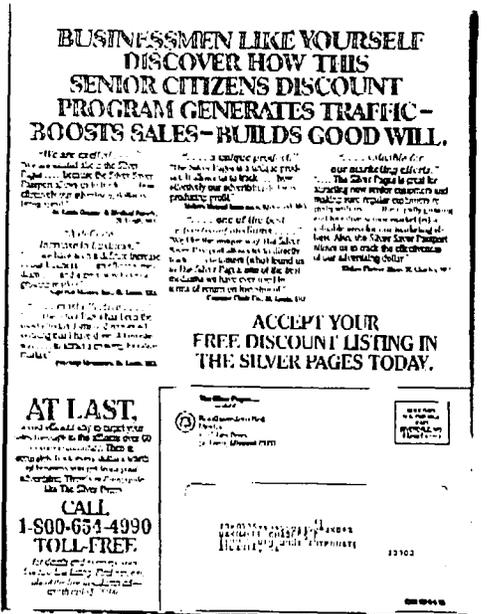
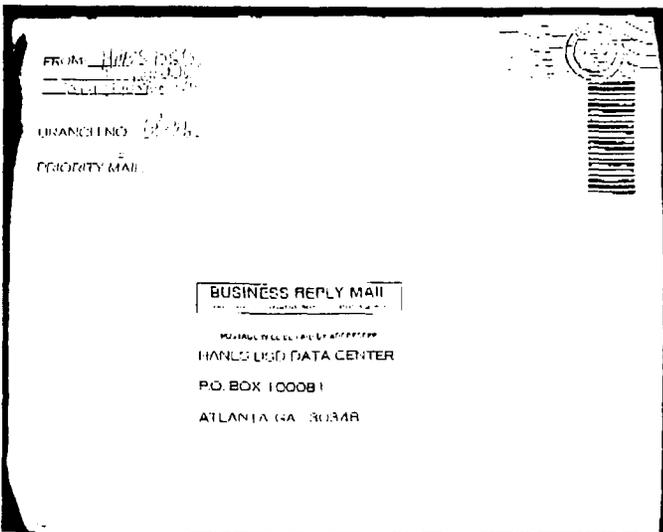
The Marr - Hildreth operator is computationally very expensive for large operator sizes. To get an idea of the number of arithmetic operations involved, consider a 31 by 31 operator running on a 512 by 512 image. Carrying out the convolutions will involve performing  $512*512*31*31$  multiplications and additions. A typical run with the above specifications requires CPU time of the order of 200 minutes (running on a VAX-11/780). In real time, this translates to 10-12 hours to obtain the binary image. Hildreth and Grimson<sup>3</sup> mention a hardware convolver constructed at MIT which can carry out the convolution in a matter of seconds. In the absence of such sophisticated hardware, the Marr - Hildreth operator is computationally very expensive for large values of  $\sigma$ . Some of this computational expense can be minimized by using an operator size equal to 6.2 times the Gaussian space constant  $\sigma$ . This operator size will reduce the needless calculation of insignificant portions of the Gaussian

A problem with the second derivative binarization is evident in Figures 8. Thick black regions in the input become hollow and sometimes get converted as thick black regions with small white holes inside. This has apparently come about when using small



(a) (b)

Figure 6(a)-(b) - Binary output of the running average binarization operator for the input images shown in Figures 1(a)-(b). Note the noise produced on the right of the images.



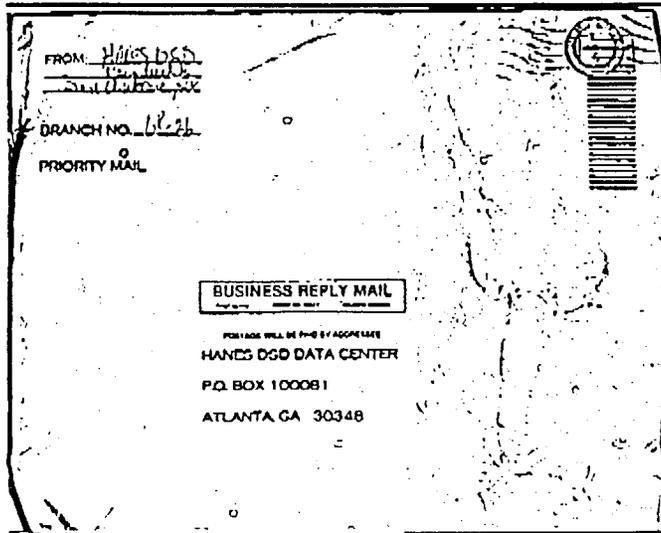
(a) (b)

Figure 7(a)-(b) - Binary output of the adaptive thresholding operator for the input images shown in Figures 1(a)-(b).

values for  $\sigma$  and results in the operator looking for intensity changes in a very small area surrounding the pixel. If the apparently black region in the input image is not uniformly black (in terms of intensity values) then these small discontinuities in intensity get detected as edges resulting in the appearance of holes in the output. That is, in the center of the black region the second derivative will approach zero from the negative (and could become positive if noise is present). For this black point in the gray level image, the corresponding point in the binary image could be either black or white depending on  $\sigma$  and the image resolution.

### 5. Conclusions

In general, global binarization techniques should only be applied in "ideal" situations where the transition between the foreground and the background is constant across the image and when a minimum noise factor is encountered. When these idealizing assumptions are not met, then a more adaptive approach must be used. Three of these adaptive binarization algorithms were described and evaluated on noisy postal letter mail document images. In terms of computational cost, running average thresholding is much better than the



(a) (b)

Figure 8(a)-(b) - Binary output of the Marr - Hildreth operator for the input images shown in Figures 1(a)-(b).

contrast and second derivative binarization. Selection of parameters is much simpler in the case of contrast thresholding compared to the other two approaches. The problem about having holes inside thick black regions in the output of Marr - Hildreth and sometimes the contrast operator was mentioned earlier. The running average scheme does not suffer from any such limitations but deteriorated under the presence of noise in the image. The contrast thresholding scheme introduces horizontal smearing at times, when the edges are too close to each other. But this seems to be much more tolerable compared to amount of noisy edges that appear by running a Marr - Hildreth operator with incorrect  $\sigma$  values.

To summarize, the running average thresholding scheme is a much better candidate for binarization of gray level images from the point of view of computational cost but performed poorly under the presence of noise in the image. The contrast thresholding scheme is moderately fast and had easy parameter adjustment capability. The second derivative approach was slow for large values of  $\sigma$  but was noise tolerant.

### Acknowledgements

The authors wish to thank Jonathan Hull who made many valuable suggestions for conducting the experimental evaluations. This work was supported in part by the United States Postal Service under contract 104230-85-M3349.

### References

- [1] F. Deravi and S. K. Pal, "Grey level thresholding using second-order statistics," *Pattern Recognition Letters* 1 (1983) pp 417-422.
- [2] E. Giuliano, O. Paitra and L. Stringa, "Electronic character reading system", *US Patent No 4,047,15* dated Sept 6, 1977.
- [3] E. Hildreth and W. Grimson, "Comments on 'Digital edges from zero crossings of second directional derivatives,'" *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-7, No. 1, pp 121-127, January 1985.
- [4] J. Kittler, J. Illingworth, J. Foglein, K. Paler, "Threshold selection based on a simple image statistic," *Computer Vision, Graphics, and Image Processing*, 30, 1985, pp 125-147.
- [5] J. Kittler, J. Illingworth, J. Foglein, K. Paler, "An automatic thresholding algorithm and its performance," *Proc. 7th International Conference on Pattern Recognition*, Montreal 1984, pp 287-289.
- [6] D. Marr, *Vision*, pp 54-61, W. H. Freeman, San Francisco, 1982.
- [7] D. Marr and E. Hildreth, "A theory of edge detection," *Proc. R. Soc. London B* 207 (1980), pp 187-217.
- [8] D. Marr and E. Hildreth, "The detection of intensity changes by computer and biological vision systems," *Computer Vision Graphics and Image Processing*, 22 (1983), pp 1-27.
- [9] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol 9, January 1979, pp 62-66.
- [10] J. S. Weszka and A. Rosenfeld, "Threshold evaluation techniques," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol 8, August 1978, pp 622-629.
- [11] J. S. Weszka and A. Rosenfeld, "Histogram modification for threshold selection," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol 9, August 1979, pp 38-52.
- [12] J. M. White and G. D. Rohrer, "Image thresholding for character image extraction and other applications requiring character image extraction," *IBM Journal of Research and development*, Vol 27, No. 4, pp 400-411, July, 1983.
- [13] P. Winston, *Artificial Intelligence*, Addison Wesley, Reading, Massachusetts, 1984.