CVGIP: GRAPHICAL MODELS AND IMAGE PROCESSING
Vol. 54, No. 3, May, pp. 259–267, 1992

NOTE

Adaptive-Neighborhood Histogram Equalization for Image Enhancement

RAMAN B. PARANJAPE, WILLIAM M. MORROW, AND RANGARAJ M. RANGAYYAN

Department of Electrical and Computer Engineering, The University of Calgary, Calgary, Alberta, Canada T2N 1N4

Received March 18, 1991; accepted December 16, 1991

By modifying the histogram of an image, a dramatic improvement in the perceptibility of details can often be achieved. However, the two commonly used methods of full-frame histogram equalization and local-area histogram equalization often fail to produce adequate enhancement when the image contains relatively small but variable-sized regions in which there are objects or features of interest with low visual contrast. A new method of adaptive-neighborhood histogram equalization that is effective in enhancing these types of images is proposed in this paper. In this method, an adaptive neighborhood is developed for each pixel in the image. The adaptive neighborhood is a compound region made up of a foreground that contains 8-connected pixels close in gray level to that of the seed pixel, and a background of neighboring pixels molded around the foreground. The histogram of this adaptive neighborhood is equalized to provide the transformation that is applied to the seed pixel. Major advantages of this method are the avoidance of block edge artifacts that are encountered in local-area histogram equalization, and improved perceptibility of image detail. Examples of images transformed using the three methods of histogram modification are presented along with a discussion of the merits of the adaptive-neighborhood method.

1. INTRODUCTION

By modifying the histogram of an image, a dramatic increase in the perceptibility of image detail can often be obtained. This was first demonstrated by Hall [1] when he attempted to equalize the gray-level density of pixels in an image. The basic assumption used was that the information conveyed by an image is related to the probability of occurrence of each gray level in the image. By uniformly redistributing the probability of occurrence of gray levels in the image, it often becomes easier to perceive the information content of the image.

Frei [2], while accepting the concept of uniformly distributing information, suggested that since the human visual system has a logarithmic response to stimuli, the image information should be redistributed to be hyperbolic over the gray scale. Both Hall's and Frei's methods of histogram modification are referred to as global or full-frame histogram methods because they are based on a transformation using the histogram of the complete image. Both methods are in common usage today because they are relatively easy to implement, require no user interaction, and often provide a substantial increase in the perceptibility of image detail. These methods are, however, less than optimal when images containing small, relatively uniform regions, in which there are objects or other details of interest, are enhanced. Typically, the gray levels corresponding to these objects are shifted into the same gray level as that of their background after full-frame histogram equalization (FFHE) and are therefore no longer visible. The basic problem is that these details do not show a sufficiently high level of occurrence in the full-frame histogram, and thus are considered to have low information content from an information theoretic viewpoint. This difficulty was first addressed by Ketchum [3] when he suggested using local-area histogram equalization (LAHE). In LAHE, rather than redistributing the gray levels of pixels on the basis of the histogram of the whole image, the equalization is based on the histogram of the portion of the image under a two-dimensional sliding window that is centered over the pixel being processed. Only the gray level of the pixel in the center of the sliding window is modified by the equalization procedures. Pizer [4], who independently suggested this method for medical images, refers to the region surrounding the center pixel as its "contextual region." In Pizer *et al.* [5] LAHE is described as "an excellent contrast enhancement method" for both natural and medical images. They also suggest that the most severe problem with LAHE is that the method is computationally slow, and propose an effective approach for reducing the computation time of the LAHE method. In this approach, the entire image is divided into a small number of rectangular, nonoverlapping regions. Local-area histograms are calculated over each region, and his-

1049-9652/92 $5.00

togram-equalizing transforms are applied to the center pixels of the respective rectangular regions. For the pixels that are not center pixels, bilinear interpolations of the four neighboring center-pixel transformations are used to approximate the local-area histogram transformation. Leszczynski and Shalev [6, 7] recently published papers indicating that some undesirable artifacts associated with bilinear interpolation may be avoided at no significant increase in computational burden if interpolations are carried out only in the horizontal direction, and the regions over which histograms are calculated are slid down the image row by row.

While LAHE and its modifications are clearly very useful, we suggest in this paper that a significant improvement over these methods may be obtained by application of histogram equalization to a more appropriate local area than a rectangle. Rather than using an arbitrarily defined rectangular region centered over the pixel being processed (as suggested by Ketchum [3] and Pizer [4]), we suggest using an adaptive region grown around the pixel. By carefully defining the growth tolerance, a more appropriate "contextual" region than a rectangular area may be determined. We find that by equalizing the histogram of this adaptive neighborhood, the perceptibility of objects and details in an image can be greater than that achieved by using LAHE. In addition, we observe that the same adaptive neighborhood will be grown by all pixels within an adaptive neighborhood that have the same gray-level value as the first pixel used to grow the region on hand. By removing the redundancy of calculating regions and histograms for these pixels, the adaptive-neighborhood histogram equalization (ANHE) algorithm can be sped up to rates similar to those of the sample and interpolate version of LAHE [5–7].

2. METHODS

In principle, FFHE defines a transformation for distributing the gray levels of pixels in an image uniformly over the available range of gray levels. However, because gray levels are quantized in a digital image, the requirement of equal numbers of pixels at each gray level can be fulfilled only approximately [1, 8]. The FFHE transformation is obtained by first finding the histogram of the whole image. By integrating the histogram, a cumulative distribution function is obtained. The cumulative distribution is normalized to the range of gray levels desired in the final image to yield the FFHE transformation [8].

The application of LAHE to an image involves using a new and unique gray-level transformation at each pixel in the image. The gray-scale transformation is obtained by equalizing the histogram of the portion of the image enclosed by a rectangular window surrounding the pixel being processed. The steps involved in LAHE are first to find the histogram of the rectangular region surrounding the pixel being processed. From this local histogram, a local cumulative distribution is then determined. The local cumulative distribution is normalized to the full dynamic range of the output image. The normalized cumulative distribution is the transformation that determines the gray-level value of the center pixel. The rectangular window is then slid over to the next pixel and the above procedure is repeated.

ANHE is similar to LAHE except that the region over which the histogram is calculated and subsequently equalized is determined contextually from the image. That is, the region's shape and size are dependent upon the characteristics of the pixels that surround the pixel being processed (called the seed pixel in ANHE). The region (called the adaptive neighborhood) is composed of two different layers. The first layer is defined as the set of 8-connected pixels that are within a certain gray-level deviation of the seed pixel. This layer can be algebraically defined as the 8-connected points, $p(k, l)$, which have the property

$$|p(k, l) - p(i, j)| \leq t, \tag{1}$$

where $p(i, j)$ is the gray level of the seed pixel and t is the maximum allowed deviation in gray level from that of seed pixel.

A second layer of pixels, of width s pixels, is grown molded to the outline of the 8-connected first layer of the adaptive neighborhood. Thus the adaptive neighborhood is a compound region made up of two kinds of pixels: those that are 8-connected and have gray levels close to that of the seed pixel (called foreground) and those with gray levels that are different from that of the seed pixel (called background).

The foreground typically identifies pixels that are similar to the seed pixel in terms of both gray level and proximity (connectivity). These pixels are usually contextually related and are likely to belong to the same object or region. For example, the region developed around a seed pixel might identify the portion of an image that is generally darker because of a large shadow in the image, or, for example, the region might identify a relatively uniform object on which there are patterns or features of interest. The idea is to equalize the histograms of regions whose shape and size are based on the information content of the image rather than of arbitrarily formed rectangular regions.

The region-growing technique used to define the foreground in this work is a simple, graphical, seed-fill algorithm known as pixel aggregation [8–10]. A flow chart for the algorithm is provided in Fig. 1. The algorithm starts with a pixel identified to be the pixel to be processed or

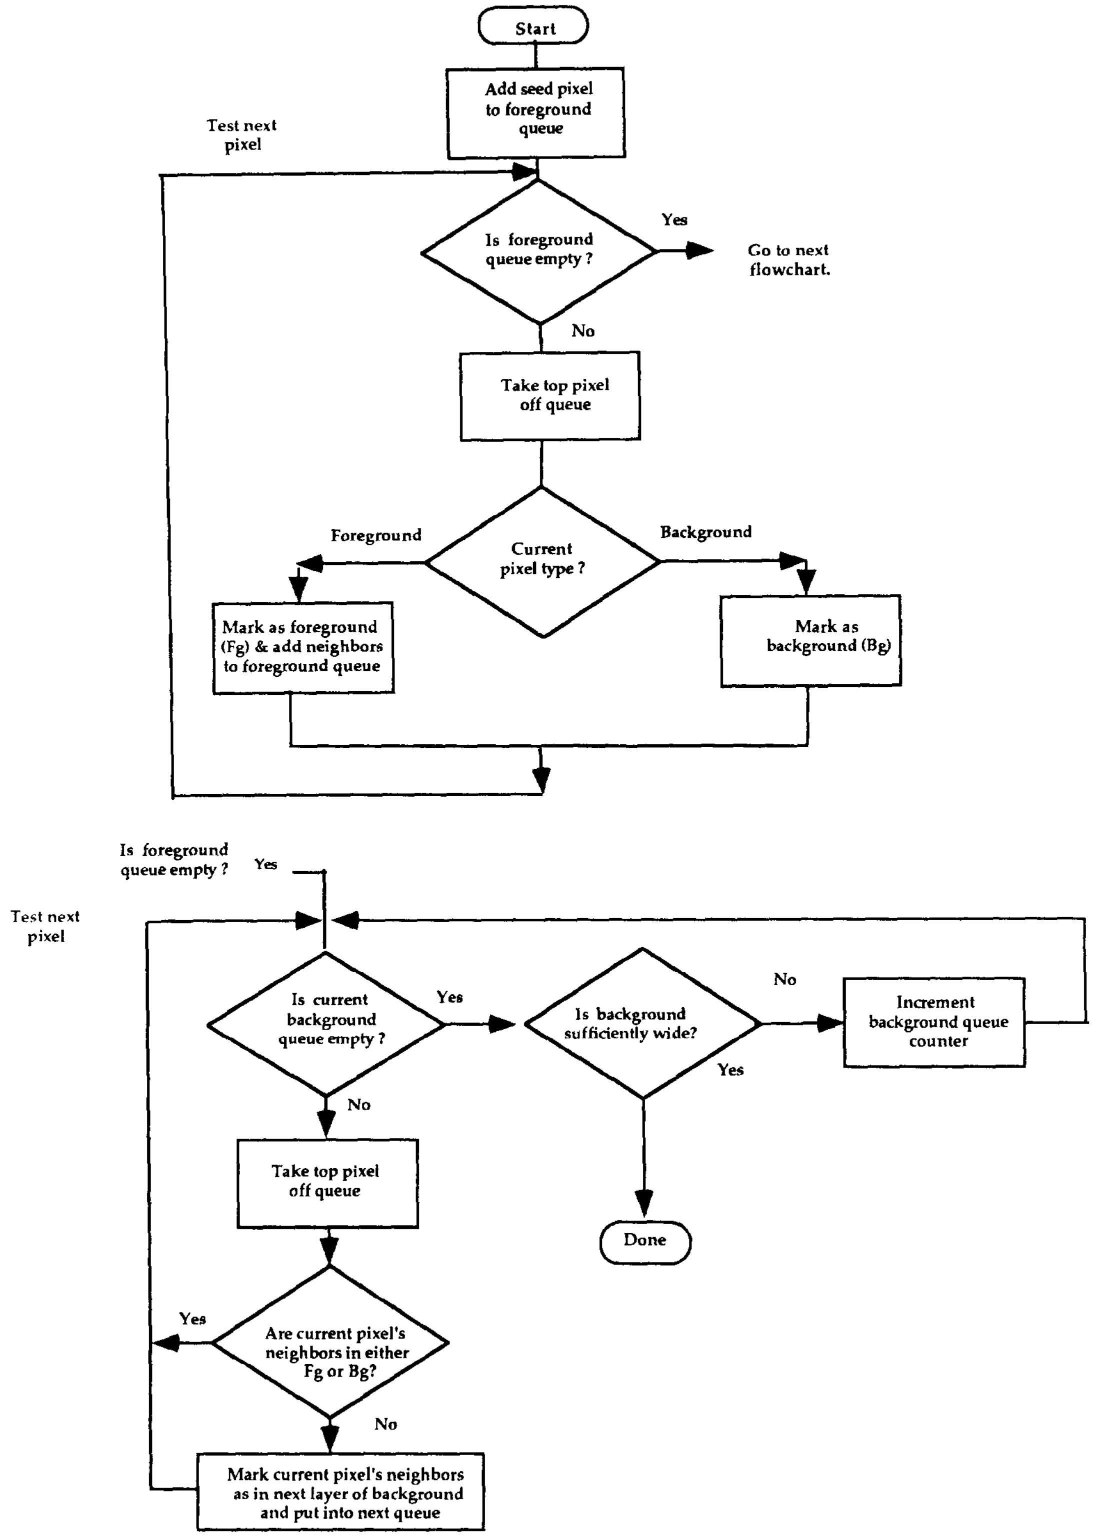

FIG. 1. A flow chart of the algorithm used to determine the adaptive neighborhood. The method used involves forming a series of queues of pixels and testing for inclusion in the foreground and in the various layers of the background.

the seed pixel. This seed pixel is placed in a queue that holds the pixels to be processed. The algorithm is executed until the queue is made empty. The mechanism for removing a pixel from the queue is to test if it is within $\pm t$ of the seed pixel value. If the pixel is within $\pm t$ of the seed pixel it is removed from the queue, but its 8-connected neighbors, which have not been processed, are added to the queue. If the pixel being tested is not within $\pm t$ of the seed pixel, it is marked as a background pixel and added to a background queue. This process is repeated until all the pixels in the first queue have been processed. When the queue is empty, a foreground will have been identified as the set of pixels which are 8-connected and within $\pm t$ of the seed pixel's gray level, and the single layer of pixels marking the start of the background will also have been identified.

The final step is then to extend the background to a width of s pixels. This is done by starting with any background pixel and testing its 8-connected neighbors. If these pixels are either part of the foreground or of the first layer of the background, they are disregarded. Otherwise, they are marked as the second layer of the background. This process continues with the third, fourth, and subsequent layers until a layer of s pixel width is identified as the background. Figure 2 presents a series of images that shows an adaptive neighborhood being formed.

After the adaptive neighborhood has been defined, its histogram is calculated, including both the foreground and the background pixels, and integrated in a manner similar to that used in FFHE or LAHE. The seed pixel is then modified using the normalized cumulative distribution of the adaptive neighborhood. The values assigned to t and s are very important in determining the type of features that will be enhanced by the ANHE method. Processing an image requires that each pixel in the image be considered as a seed pixel for region growing and subsequent histogram equalization. Thus the above process is applied to every pixel in turn.

It is of interest to observe that identical adaptive neighborhoods are formed for all pixels in an adaptive neighborhood with the same gray-level value as that of the

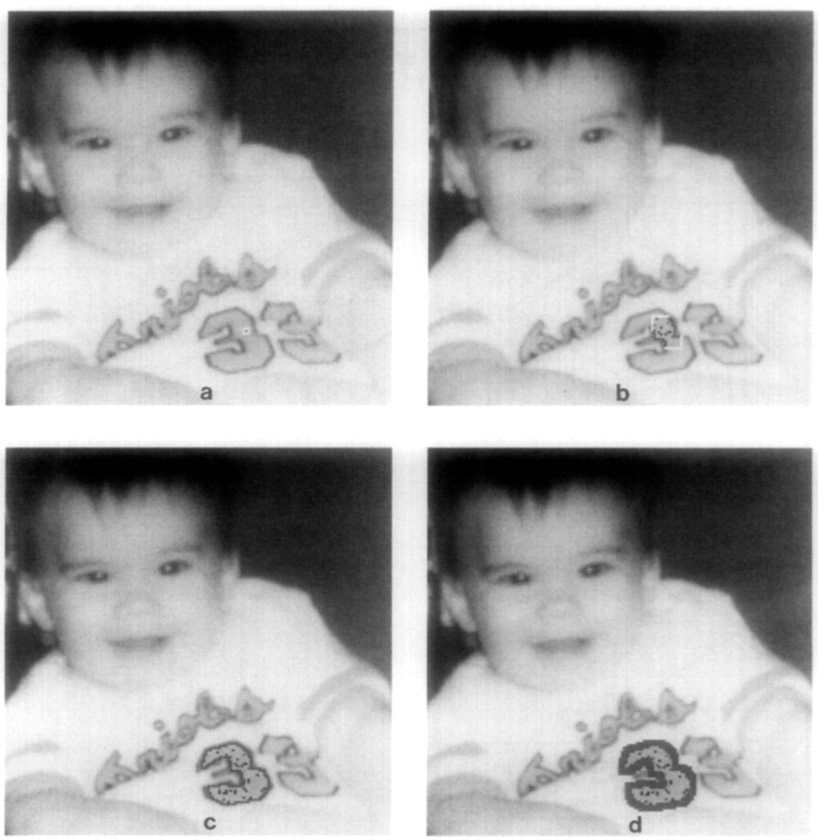

FIG. 2. A demonstration of region growing. The white pixels indicate pixels in the foreground queue, the light gray pixels indicate pixels in the foreground, the dark gray pixels indicate pixels in the background, and the black pixels indicate pixels in the foreground that have the same gray level as that of the seed pixel. (a) The seed pixel (black) and its 8-connected neighbors in the foreground queue (white), and (b) further development of the foreground. (c) The completed foreground and the first layer of the background and (d) the completed adaptive neighborhood with a 3-pixel-wide background layer. The image, with 176 × 176 pixels and 256 gray levels, was obtained from a publicly-accessible electronic bulletin board at vax.eedsp.gatech.edu in the directory /database/images.

seed pixel. Thus the transformation applied to the seed pixel is also applicable to all the pixels in the adaptive neighborhood that have the same gray-level value as that of the seed pixel. These pixels are called redundant seed pixels in this discussion as no further calculations are required to determine their output values. Since the seed pixel gray-level value may occur many times in the adaptive neighborhood, a very substantial reduction in computation is obtained by updating the gray-level value of all of these pixels using the same ANHE transformation.

3. RESULTS

Figure 3 demonstrates the characteristics and the limitations of FFHE. Two different scenes (Fig. 3a, girl in snow cave (GSC), and Fig. 3c, girl on beach (GOB)) are presented. Each image contains some relatively uniform regions inside of which there are features of interest. FFHE enhances the GSC image, making it possible to see some of the details of the girl (Fig. 3b). However, it continues to be difficult to see many details such as the girl's face and the wrinkles and folds in her coat. In the GOB image, rather than increasing the perceptibility of the features of interest, FFHE clearly degrades them further (Fig. 3d).

In Fig. 4, the LAHE techniques is applied to the GSC image. Four versions of this image are presented with sizes of the sliding window (local area) of 11 × 11, 21 × 21, 41 × 41, and 101 × 101 pixels. The size of the window determines the size of the detail that is most enhanced. The best overall image appears to be the version enhanced using the 101 × 101 window (Fig. 4d). In Fig. 5, the LAHE technique is applied to the GOB image. Four versions of the images are again presented. The LAHE method significantly enhances many of the details of interest in this image.

Results of ANHE enhancement of the GSC image are presented in Fig. 6. Figure 6a corresponds to $t = 0$ and $s = 16$; thus, the adaptive neighborhood is formed with a foreground that contains only those pixels that have the same gray-level value as that of the seed pixel and a background that is 16 pixels wide and surrounds the foreground. The figure shows that very small details, like the characteristics of the girl's face, are enhanced significantly; however, other features, such as small variations in the background, are also amplified, resulting in a noisy image. Figure 6d, on the other hand, is produced with an adaptive neighborhood of $t = 64$ and $s = 8$. This image is quite similar to the FFHE image, except that many more details are perceived. The images in Figs. 6b and 6c are

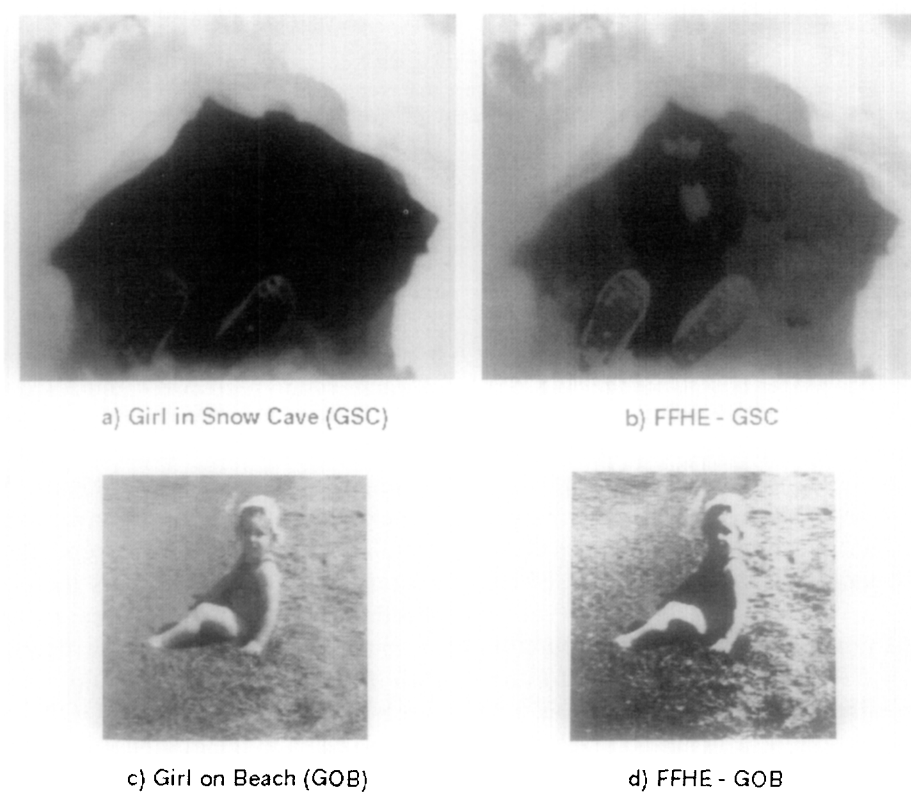

FIG. 3. The input images and FFHE images. (a) The girl in snow cave image (GSC; 288 × 240 pixels, 256 gray levels) and (c) the girl on beach image (GOB; 176 × 176 pixels, 256 gray levels). (b) The FFHE-enhanced GSC image. The image is significantly clearer; however, further improvement may be possible. (d) The FFHE-enhanced GOB image. The image is not improved: increased contrast results in a loss of detail in the darker parts of the image.

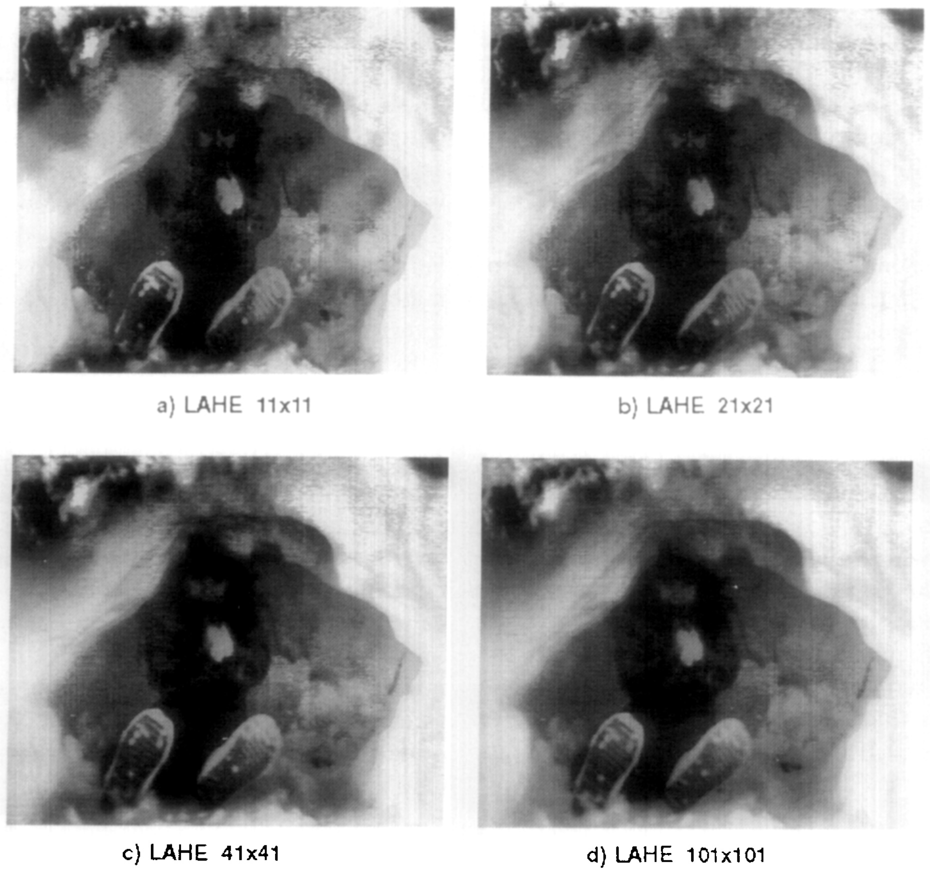

FIG. 4. Local-area histogram equalization (LAHE) of the GSC image using window sizes of 11×11 (a), 21×21 (b), 41×41 (c), and 101×101 pixels (d). The size of the window determines the type of detail in the image that is enhanced. The girl's face is most clearly seen in the image using the 41×41 window (c), but the best overall image appears to be the version enhanced using the 101×101 window (d).

produced using intermediate values for the foreground and background ($t = 16$, $s = 8$ and $t = 16$, $s = 5$).

In Fig. 7, the ANHE method is applied to the GOB image. The values of the t and s parameters are the same as those used in Fig. 6. The image in Fig. 7a is a rather noisy image much like the LAHE 11×11 image (Fig. 5a). There may be some debate as to which image best displays the details of the scene. The images in Figs. 7b and 7c appear to be less noisy than that in Fig. 7a and show the outline of the girl's arm and the pattern on her swim suit, but the details of her face are somewhat less clear. The image in Fig. 7d looks very much like the FFHE image in Fig. 3d.

Table 1 indicates the time required for the computation of the images in Figs. 4, 5, 6, and 7. The LAHE algorithm requires an exponentially increasing computation time as the window size is increased. The ANHE algorithm, on the other hand, has a steady decrease in computational time as the size of the first layer of the adaptive neighborhood is increased. The longest computational time is taken by the parameter combination of $t = 0$, $s = 16$.

The last two columns in Table 1 indicate the percentage of redundant seed pixels in the GSC and GOB images. It is not necessary to find an adaptive neighborhood for the redundant seed pixels as these pixels are in the first layer (the foreground) of the adaptive neighborhood of some other seed pixel and have the same gray-level value as that of the seed pixel. There will be a high degree of variability in the percentage of redundant seed pixels in images as this factor is highly dependent upon the contextual information in a given image. One can, however, observe the general trend that as the first layer tolerance (the parameter t) is increased, the percentage of redundant seed pixels increases, and the time required to compute the output image decreases. For example, with the parameters $t = 64$ and $s = 8$, about 98% of the pixels in both the GSC and GOB images are redundant seed pixels, and the computation time for these images is less than $\frac{1}{6}$ the time required when the parameters are $t = 0$ and $s = 16$. With high percentages of redundant seed pixels, computation times for ANHE are in the range of those presented by Leszczynski and Shalev [7] for their approximation to the LAHE algorithm.

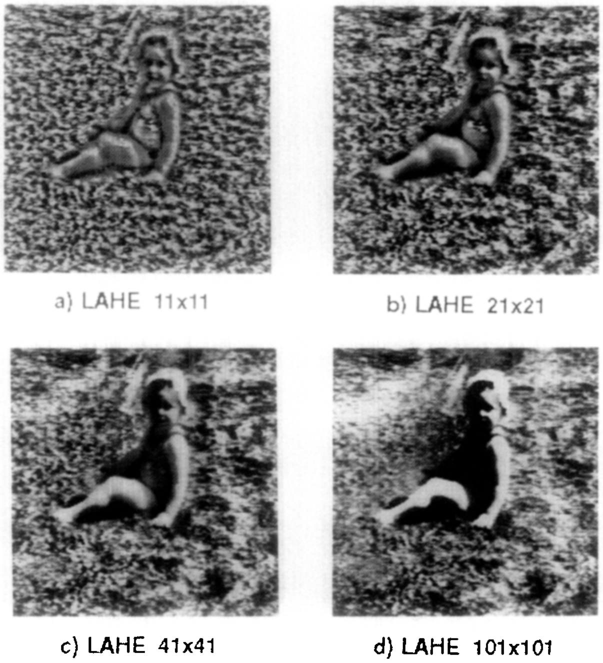

FIG. 5. LAHE of the GOB image using window sizes of 11 × 11 (a), 21 × 21 (b), 41 × 41 (c), and 101 × 101 pixels (d). The 11 × 11 window (a) produces the most clear enhancement of the girl's face, while the 21 × 21 (b) and 41 × 41 (c) versions most clearly show the outline and pattern of the swim suit and differentiate the right arm from the girl's shadow.

4. DISCUSSION

The basic idea behind both LAHE and ANHE is to equalize the histogram of the appropriate contextual region of an image to enhance detail. If the contextual region contains multiple objects of low visual contrast, then the contrast of these objects can be increased with respect to one another and with respect to the background through histogram equalization. The LAHE method involves arbitrarily selecting the size of the sliding window over which histogram equalization is performed. The LAHE method works well when in fact the sliding window incorporates a number of objects with poor visual contrast. However, if the window is too small and does not contain multiple objects, or if the window is too large and the objects of interest do not represent a significant portion of the local-area histogram, then LAHE fails. In ANHE, on the other hand, the size and the shape of the region over which the histogram is equalized are determined from the actual details and content of the image, thereby circumventing this problem to some extent. The approach adopted is to assume that the image is made up of a number of relatively uniform regions and that there are objects within these regions that are to be enhanced. The adaptive neighborhood is grown so that it can incorporate the sizes and the shapes of these regions. Thus ANHE is based primarily on visually perceivable detail in the image, assuming of course that appropriate values of t and s have been used.

The parameter s in ANHE is used to incorporate pixels into the adaptive neighborhood that are different from the pixels in the first layer (foreground) of the adaptive neighborhood. The parameter s provides a mechanism for mediating the gray-level change introduced by ANHE. For example, in Fig. 6d (GSC, $t = 64$, $s = 8$), the adaptive neighborhood determined for most of the seed pixels that are inside the snow cave is the entire snow cave. By setting $s = 8$, some bright pixels from outside the cave are included in the adaptive neighborhood, restricting the output values of the equalization procedure to levels below the maximum brightness (gray level, 255).

Rehm and Dallas [11] have shown that the LAHE method produces an edge artifact at points at which the sliding window crosses a sharp natural boundary in an image. They argue that this artifact is due to the rapid change in the pixel transformation as the window crosses the boundary. To correct this problem, they recommend controlling the abrupt change in the equalizing transformation by subtracting a very smooth version of the image prior to application of the LAHE method. In ANHE, this problem of a change in the equalizing transform when the seed pixel is in one region but the window contains significant values from outside that region is nonexistent. When seed pixels are inside one region, they will all have the same transformation. As just indicated, however, some contribution from outside the region is useful for maintaining relative gray levels between the regions, and this

TABLE 1
Computation Times for the LAHE and ANHE Methods

Method	Size	GSC time (s)	GOB time (s)	GSC % redun pixels	GOB % redun pixels
LAHE	11 × 11	370	229	—	—
	21 × 21	536	329	—	—
	41 × 41	1091	671	—	—
	101 × 101	4121	2385	—	—
ANHE	$t = 0$, $s = 16$	1657	1497	48	24
	$t = 16$, $s = 8$	592	925	80	59
	$t = 16$, $s = 5$	441	730	80	59
	$t = 64$, $s = 8$	257	173	99	98

Note. Time required for the computation of the images in Figs. 4, 5, 6, and 7. The LAHE algorithm requires an exponentially increasing computation time as the window size is increased, while the ANHE algorithm has a steady decrease in computational time as the size of the first layer of the adaptive neighborhood is increased. The last two columns contain the percentage of redundant seed pixels in the GSC and GOB images. The programs were run on a Sun 3/60, 20-MHz computer, operating under SunOS UNIX 4.1.1.

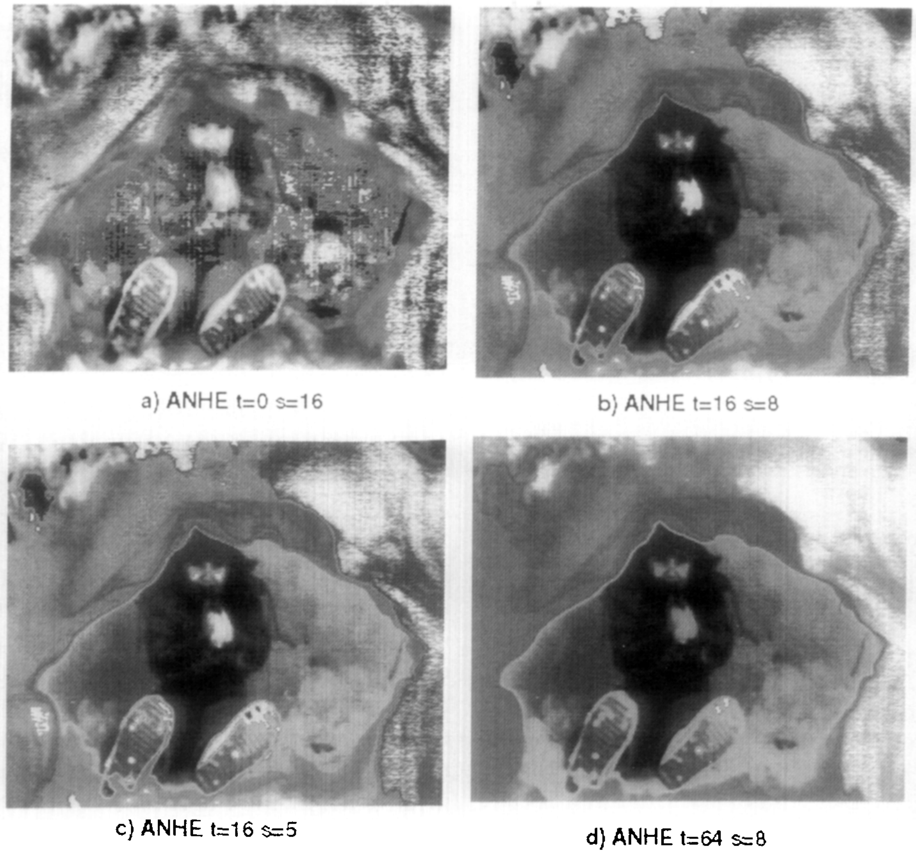

FIG. 6. Adaptive-neighborhood histogram equalization (ANHE) of the GSC image using various combinations of t and s: $(t, s) = (0, 16)$ (a); $(t, s) = (16, 8)$ (b); $(t, s) = (16, 5)$ (c), $(t, s) = (64, 8)$ (d). The ANHE method with $(t, s) = (64, 8)$ (d) clearly shows more details in the GSC image than the other methods presented.

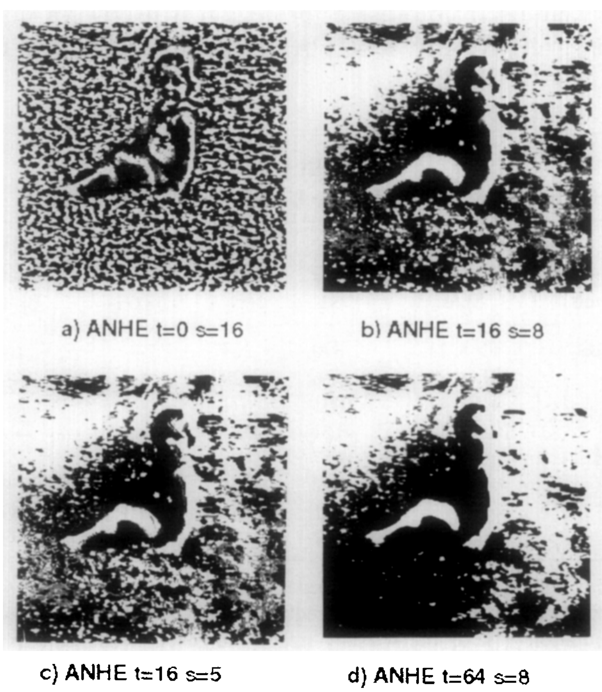

FIG. 7. ANHE of the GOB image using various combinations of t ans s (same as those given in the legend to Fig. 6). The ANHE method with $(t, s) = (16, 8)$ shows the details of the girl's face and arm without significantly transforming the background beach area.

is achieved with the s parameter. The inability to maintain relative gray levels is a limitation of the LAHE method from which ANHE does not suffer if the t and s parameters are properly chosen.

A second problem with LAHE, to which ANHE is also susceptible, is the undesired enhancement of noise. This typically is a problem when the sliding window contains a region that is relatively uniform, and noise values are just above or just below the uniform gray levels. Due to the large enhancement of the relatively uniform region provided by histogram equalization, the noise pixels are also enhanced significantly. Pizer *et al.* [5] suggested that noise enhancement can be restricted by limiting the slope of the enhancement transformation. In other words, by limiting the degree to which the uniform region is redistributed over the gray scale of the output image, the degree to which noise is enhanced may also be limited. This approach, used by Rehm and Dallas [11] in their investigation of boundary artifacts in LAHE, may be directly applied to ANHE also.

The approach adopted in this work of transforming the gray level of a seed pixel on the basis of the characteristics of its adaptive neighborhood is along the lines of our previous work with adaptive neighborhood contrast enhancement. Gordon and Rangayyan [12] first suggested this approach for the enhancement of contrast in images and others [9, 10, 13–15] have continued this work. Their viewpoint, however, was somewhat different from that used in this paper: they attempted to identify distinct objects using adaptive neighborhoods, while we attempt only to identify uniform regions in which objects can be made more perceivable by histogram equalization. This difference in viewpoint leads to variations in the way in which the adaptive neighborhood is calculated.

5. CONCLUSION

In this report, a new method of adaptive-neighborhood histogram equalization is presented. The method involves identifying contextually related regions in an image and applying histogram equalization to these regions. The ANHE method is compared with two well-known methods, full-frame histogram equalization and local-area histogram equalization. It is shown that for certain types of images, the FFHE and LAHE methods produce results poorer than those of the ANHE method. Using natural redundancies, processing times for the ANHE method can be brought into the range of the sample and interpolate approximations of the LAHE method.

ACKNOWLEDGMENT

This project is supported by grants from the Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

1. E. H. Hall, Almost uniform distributions for computer image enhancement, *IEEE Trans. Comput.* **C-23**(2), 1974, 207–208.
2. W. Frei, Image enhancement by image hyperbolization, *Comput. Graphics Image Process.* **6,** 1977, 286–294.
3. D. J. Ketchum, Real-time image enhancement techniques, *Proc. SPIE/OSA* **74,** 1976, 120–125.
4. S. M. Pizer, Intensity mappings for the display of medical images, in *Functional Mapping of Organ Systems and Other Computer Topics* (P. D. Esser, Ed.), 11th Annual Symposium on the Sharing of Computer Programs and Technology in Nuclear Medicine, pp. 205–215, 1981.
5. S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Geer, B. tar Haar Remeny, J. B. Zimmerman, and K. Zuiderveld, Adaptive histogram equalization and its variations, *Comput. Vision Graphics Image Process.* **39,** 1987, 355–368.
6. K. W. Leszczynski and S. Shalev, Digital contrast enhancement for online portal imaging, *Med. Biol. Engrg.* **27,** 1989, 507–512.
7. K. W. Leszczynski and S. Shalev, A robust algorithm for contrast enhancement by local histogram modification, *Image Vision Comput.* **7**(3), 1989, 205–209.
8. R. C. Gonzalez and P. Wintz, *Digital Image Processing,* 2nd ed., Addison–Wesley, Reading, MA, 1987.
9. W. M. Morrow, *Region-Based Image Processing with Application to Mammography,* M.Sc. thesis, The University of Calgary, 1990.
10. W. M. Morrow, R. B. Paranjape, R. M. Rangayyan, and J. E. L. Desautels, Region-based contrast enhancement of mammograms, *IEEE Transactions on Medical Imaging,* 1992, in press.
11. K. Rehm and W. J. Dallas, Artifact suppression in digital chest radiographs enhanced with adaptive histogram equalization, *Proc. SPIE* **1092,** 1989, 220–300.
12. R. Gordon and R. M. Rangayyan, Feature enhancement of film mammograms using fixed and adaptive neighborhoods, *Appl. Opt.* **23**(4), Feb. 1984, 560–564.
13. A. P. Dhawan, G. Buelloni, and R. Gordon, Enhancement of mammographic features by optimal adaptive neighborhood image processing, *IEEE Trans. Med. Imaging* **MI-5**(1), Mar. 1986, 8–15.
14. A. P. Dhawan and E. LeRoyer, Mammographic feature enhancement by computerized image processing, *Comput. Methods Programs Biomed.* **27**(1), 1988, 23–35.
15. A. Beghdadi and A. Le Negrate, Contrast enhancement technique based on local detection of edges, *Comput. Vision Graphics Image Process.* **46,** 1989, 162–174.