A Reprint from the

PROCEEDINGS

Of SPIE-The International Society for Optical Engineering

Volume 845

Visual Communications and Image Processing II

27–29 October 1987
Cambridge, Massachusetts

Medical image communication using halftone algorithms

Robert B. Goldstein, Eli Peli
Eye Research Institute of Retina Foundation
20 Staniford Street, Boston, Massachusetts 02114
and Harvard Medical School
Karl Wooledge
Image Analysis Laboratory
Tufts-New England Medical Center, Boston, Massachusetts

Medical image communication using halftone algorithms

Robert B. Goldstein, Eli Peli

Eye Research Institute of Retina Foundation
20 Staniford Street, Boston, Massachusetts 02114
and Harvard Medical School

Karl Wooledge

Image Analysis Laboratory
Tufts-New England Medical Center, Boston, Massachusetts

ABSTRACT

Physicians frequently need to see retinal, biopsy, and other images, produced originally by medical diagnostic equipment and then processed by computers. At present this is done using either time-consuming photography or expensive video equipment. We have used digital halftoning as a method for fast communication of images both on print and on computer terminal screens. This method quickly produces a good-quality halftone rendition of a grey-scale image. These images are suitable for display and printout on inexpensive devices that normally do not have grey-scale capability. The algorithm is based on a previously published error-propagation technique. We improved the algorithm by including a factor that accounts for the difference in size between light and dark points on various devices. The algorithm is extended to devices that have two bitplanes (VT240), and the execution and transmission times are reduced.

At Tufts-New England Medical Center in Boston, this program has been used in processing and reporting the results of muscle and nerve biopsies. At the Eye Research Institute in Boston, it has been used to report the results of retinal visual field mapping.

This technique has a wide range of applications. It allows "image processing" to be done on computers that have no traditional image-processing hardware. It allows several users to operate simultaneously on time-shared systems that have only a single image-processor. Images are displayed on 1 or 2 bitplane devices (LA50 printers or VT240 terminals). It allows image transmission over long distances -- replacing video communication equipment with RS232 cables and modems.

1. INTRODUCTION

A new method of communicating the results of clinically relevant processed images to the physician has been developed. This technique enhances available photographic and video methods and has certain advantages over these as well. Over the last 15 years image processing in the medical field has evolved into a commonplace tool.[1] Commercial systems are available for acquiring, editing, enhancing, and storing retinal images.[2] In addition, methods are being developed for transmitting images from a central, advanced medical imaging center to the physician's home or office.[3] These transmissions occur via satellite or land-line links, and require expensive CPU and video equipment for image display at the remote site. Transmission of a halftone image would eliminate the need for expensive equipment at the remote site.

We have developed a method that quickly produces a good-quality halftone rendition of a grey-scale image.[4] This technique has been used at two different facilities to report the results of processing three different kinds of images.

Traditional digital halftoning uses very high-resolution devices to create cells consisting of multiple pixels. In each cell a grey-scale is created by darkening varying numbers of pixels. The patterns of pixels are called "screens."[5] Our objective here is to simulate grey-scale on devices without high spatial resolution for which creating a "screen" is impractical.

Saghri, Hou, and Tescher[6] accomplished this by developing an error-propagation technique whereby the image is re-sized to match the number of cells in the output device. Each cell is then assigned an "on" or "off" value. The errors formed by this assignment are taken into account when assigning this binary state to neighboring cells.

In our clinical application we used this method to send the results of processed images directly to printers and terminals in the practitioner's office. In the process of implementing this technique we made the following enhancements:

1. We accelerated the program so that it could be used in a reasonable amount of time on a daily basis.
2. Because we found that the printed images were too dark, we modified the algorithm to account for the difference in size between the black and the white pixels on the various devices.
3. To take advantage of the additional bitplane of the VT240 terminals, we extended the algorithm to produce what we call a "quartertone" image -- an image containing pixels with 4 possible levels.

2. SAMPLE IMAGES

Figure 1 shows the image of a girl as seen on an image processor displayed with the full 256 available grey levels. This image is a 128 x 128 pixel image. The problem of digital halftoning is to render this image on devices that have a limited resolution and/or limited grey-scale capability. Figure 2 shows this same image printed on an LA50 printer, which has only 2 levels (on-off) but good resolution (500 x 1000 cells).

Figure 1. GIRL on an image processor

Figure 2. Halftoned GIRL on an LA50

Figure 3 shows the four different methods of display for viewing on a CRT that we have implemented. The method chosen for any particular application is a compromise between the desire to see various details versus speed of the rendition. The left side of the figure [panels (a) and (b)] shows simple thresholding techniques using 2 and 4 levels, respectively. The right side of the figure shows the image rendered via (c) the usual (2-level) halftone algorithm, and (d) the "quartertone" algorithm described in this paper.

3. THE ALGORITHM

A flowchart in Figure 4 shows that the process is logically divided into 2 tasks: (a) choosing and preprocessing the image for the "target device," and (b) performing the halftoning. The original image is expanded or compressed to fit into the number of "cells" available on the target device. Knowledge of the aspect ratios of the pixels on the target devices is built into the program (for example, the LA50 is 2:1), which is taken into account in the preprocessing step. The preprocessing details in the work of Saghri, Hou, and Tescher[6] consists of shrinking or expanding the image to match the dimensions of the target device.

When an expansion of the image is required, the Saghri, Hou, and Tescher[6] algorithm interpolates the grey levels to fill in the "added" pixels. We have accelerated up the process by NOT interpolating, but replicating the bounding pixels. This works because the error-propagation halftone process is a form of interpolation and adequately "fills in" the missing pixels. The expanded image has the dimensions of the target device and is now ready for the halftone process.

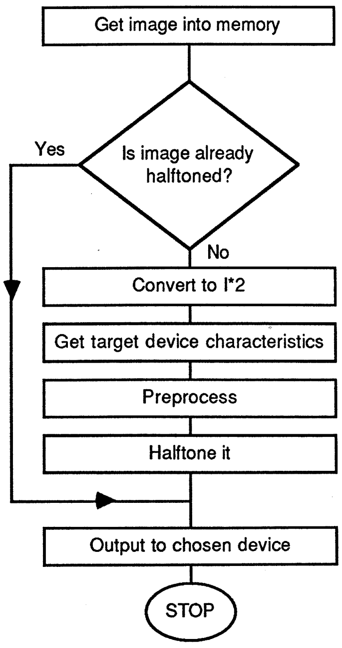

Figure 4. Flowchart of Halftone Program

(a) Thresholding with 2 levels

(c) Halftone with 2 levels

(b) Thresholding with 4 levels

(d) Quartertone

Figure 3. GIRL on a VT240

The halftone process consists of assigning a 0 or 1 to each cell. The quartertone process consists of assigning each cell to 0, 1, 2, or 3. Assigning a zero generates an "error" equal to the actual grey-scale value. Equivalently assigning a 1 (in the 2-level case) or a 3 (in the 4 level case) generates an "error" equal to the difference between the maximum and actual grey-scale values. When a cell is assigned a value, the errors introduced from assignments of previous pixels are considered. Following the notation of Saghri, Hou, and Tescher,[6] we define $E_g(m,n)$ as the total error generated, I(m,n) as the grey-scale value, and H(m,n) as the halftone (or quartertone) assignment at position (m,n). To calculate the error propagated to point (m,n), we have (eq. 10 of Saghri, Hou, and Tescher[6]):

$$E_p(m,n) = \sum_{i=1}^{I}\sum_{j=1}^{J} C_{ij} E_g[(m-i+1),(n-j+1)] \qquad \text{where} \qquad \sum_{i}\sum_{j} C_{ij} = 1 \ \text{ and } \ C_{11} = 0 \tag{1}$$

The C matrix is the error-distribution function and serves to give less weight to positions further from (m,n). C is device-dependent and is determined empirically. As illustrated in Figure 5, this matrix calculates the errors contributed by the upper-left neighboring pixels. The "C" in figure 5 is a 2 x 3 matrix; however, we usually use a 2 x 2 "C" matrix.

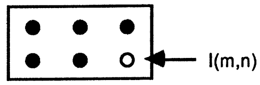

Figure 5. Pixels that contribute to error at (m,n)

To expand the algorithm for the 4-level display, we define:

$$\text{denom} = \text{maxval/levels} \qquad \text{where levels} = 2,4$$

and

$$t = I(m,n) + E_p(m,n).$$

For the usual case of maxval = 256, denom is 128 when levels = 2, and denom is 64 when levels = 4. H(m,n) is now assigned in the following way:

$$H(m,n) = \begin{cases} \left\lfloor \dfrac{t}{\text{denom}} \right\rfloor & \text{for } t < \text{maxval} \\ \text{levels} - 1 & \text{for } t \geq \text{maxval} \end{cases} \tag{2}$$

The square brackets, [ ], indicate the "floor function" in which the result of the division is truncated to the nearest integer lower than the actual quotient. H therefore gets set to 0 and 1 for 2 levels and 0-3 for 4 levels. $E_g(m,n)$ is then calculated:

$$E_g(m,n) = \begin{cases} t - \text{maxval} & \text{for } H(m,n) = \text{levels} - 1 \\ t - H(m,n) \cdot \text{denom} & \text{for } H(m,n) = 0 \ldots \text{levels} - 2 \end{cases} \tag{3}$$

For output with only 2 levels, $E_g(m,n)$ is assigned in the following way:

$$E_g(m,n) = t + K_{fact}, \quad \text{for } H(m,n) = 0 \tag{4}$$

K_{fact} accounts for the difference in size of a black versus white pixel. This problem exists both for printed output and for output displayed on a CRT. For a CRT, the lighted phosphors bleed into surrounding black areas. On the printed page, when a black dot is printed, it spreads ink over surrounding cells so a black pixel is larger than a white pixel. The problem is more noticeable for printed rather than CRT-displayed outputs. These size differences introduce additional errors. K_{fact} is introduced into eq. 4 to account for these additional errors and is determined empirically for each device. Presently, K_{fact} ranges from 0 to 1023.

4. CLINICAL APPLICATIONS

The halftone algorithm has been used to report three different types of medical images at two different facilities in Boston. At Tufts-New England Medical Center, it is used to report the results of muscle and nerve biopsies; at the Eye Research Institute, to report the results of retinal mapping.

Muscle biopsy slides are received for image analysis. Two types of fibers (light and dark) in these biopsies are measured. Examples of the types of parameters reported are number of fibers, average diameters, average areas, and descriptors of clustering and neighborhood relationships.[7] The physician receives the report on an LA50 printer in his or her office. The halftone image (Figure 6) is included in the report to provide a meaningful reference for the accompanying statistics. Fiber boundaries and the two types of fibers are easily distinguishable in this image. Panel (a) shows normal tissue and panel (b) shows diseased tissue.

Nerve biopsy slides have been received for image analysis. A sample of the processed image is shown in Figure 7. The accompanying report contains measurements of fiber density, total number of fibers, and histograms of myelin thickness and fiber diameters.

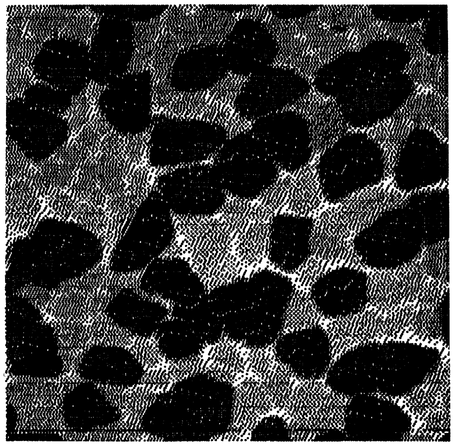

(a) Normal Tissue

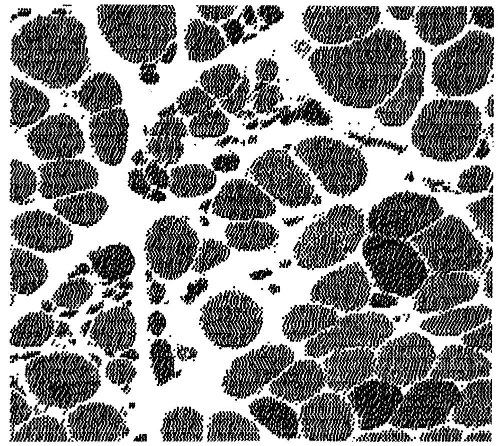

(c) Diseased Tissue

Figure 6. Halftoned Images of Muscle Biopsy

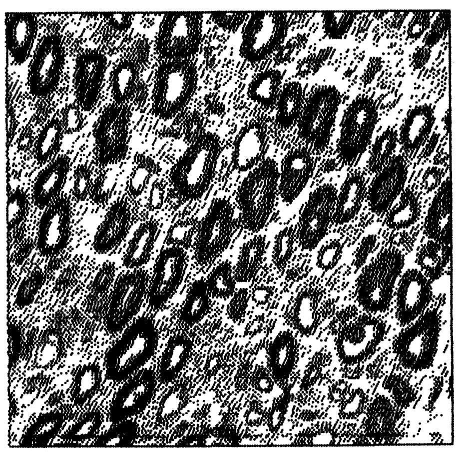

Figure 7. Halftoned Image of Nerve Biopsy

At the Eye Research Institute, automated perimetry mapping is performed directly on the retina using the Scanning Laser Ophthalmoscope.[8] A hard copy of the resultant retinal map showing the scotoma location is placed in the patient's record. The relation of the mapped visual-field loss to the visible retinal lesion aids the physician in deciding whether to treat the disorder with laser photocoagulation and in determining the optimal site to be treated. Figure 8 shows the map rendered via the halftone algorithm on an LA50 (Digital Equipment Corp., Maynard, MA.) printer. Figure 9 shows the map rendered using the default halftone screen on an Apple Laserwriter (Apple Computer, Inc., Cupertino, CA.). Although the Laserwriter output is clearly superior (due to 300-dots-per-inch resolution), the cost of a Laserwriter is an order of magnitude greater than that of an LA50. In the case of the retinal maps, the halftone output (regardless of the device) is preferable to photographic output because the former can be annotated more easily by the clinician and can be FAXed and reproduced via xerography. The quartertone image can be transmitted immediately to the laser room or office.

5. CONCLUSION

The presentation of images using halftones is slower and more cumbersome than viewing the monitor of an image processor. The halftone process provides less satisfactory hardcopy than some other available methods. However, even with those disadvantages, the halftone process has been used successfully in three different clinical applications. Its advantages are simplicity and low cost of the hardware required to present the output to the physician, ease of transmitting the images over existing communication hardware, and the continual need for paper hardcopy for medical files.

Future development of this technique will entail the following:

a. Continue to speedup the program.

b. Add an image-enhancement step to preprocess the image before halftoning. Preliminary experiments using adaptive enhancement[9] look very promising.

c. Add the capability to do interactive work. Many of these inexpensive terminals have a graphics cursor. The computer can relate the cursor position to the pixel coordinates of the actual image in the computer memory. This will give computers without an "image processor" additional image-analysis capability.

Clinicians need to see pictures as well as numerical results to make sound decisions regarding diagnoses and treatments. The halftone process is a valuable, flexible tool that enhances the available techniques for communicating retinal maps and other medical images to physicians.

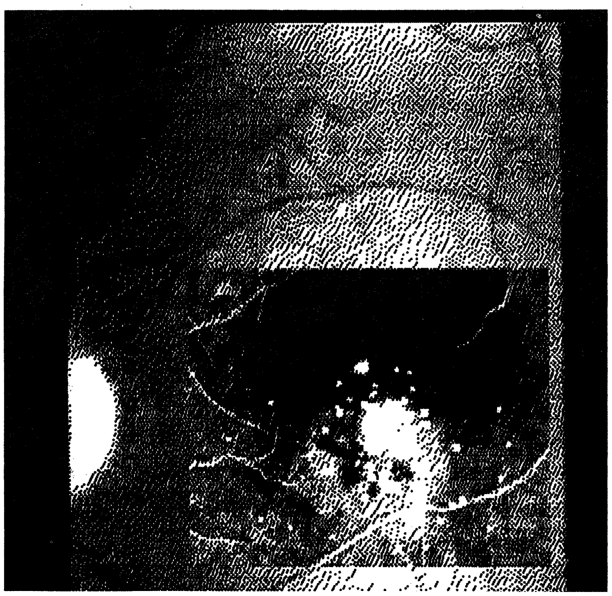

Figure 8. Retinal Map on an LA50

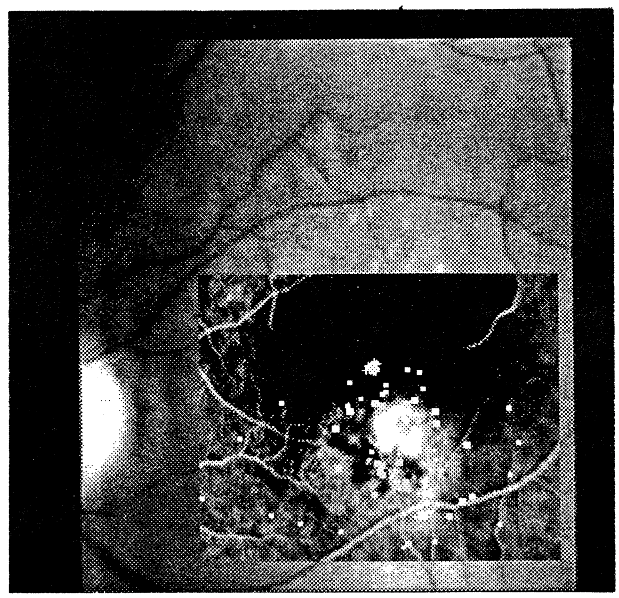

Figure 9. Retinal Map on an Apple Laserwriter

6. ACKNOWLEDGMENTS

Supported in part by a grant from the Alcoa Foundation and by R01-EY05957 from the National Institutes of Health. We wish to thank George T. Timberlake, Ph.D, Eye Research Institute for providing the retinal maps.

7. REFERENCES

1. J.W. Warnicki, "Development of an imaging system for ophthalmic photography," J. Biol. Photogr. 53(1), 9-17 (1985).
2. IS-2000 Ophthalmic Imaging System. PAR Technical Report. PAR Corporation, New Hartford, NY (1984).
3. R. Murphy, "Imaging: good medicine for hospitals' ills," Photonics Spectra, 79-84 (December, 1985).
4. R.B. Goldstein, E. Peli, and K. Wooledge, "Halftone - A program for converting grey-scale images to halftones ," in Proceedings of the Digital Equipment Users Society, San Francisco, pp. 209-211, Fall 1986.
5. D. Pelli, "Programming in postscript," BYTE 12(5), 185-202 (1987).
6. J.A. Saghri, H.S. Hou, and A.G. Tescher, "Personal computer based image processing with halftoning," Opt. Eng. 25, 499-504 (1986).
7. J.F. Brenner, D.J. Zahniser, D.W. Ziegelmiller, B.S. Adelman, T.L. Munsat, and W.G. Bradley, "An automated system for the analysis of muscle and nerve biopsies", Applied Optics, in press.
8. G.T. Timberlake, A.E. Jalkh, C.L. Trempe, and A.P. Cialdini, "Use of SLO retinal function maps in laser photocoagulation treatment of macular disease", Invest. Ophthalmol. Vis. Sci. 28(ARVO Suppl.), 49 (1987).
9. T. Peli and J.S. Lim, "Adaptive filtering for image enhancement," Opt. Eng. 21, 108-112 (1982).