DIGITAL IMAGE HALFTONING BY NOISE THRESHOLDING

Anamitra Makur

Electrical Communication Engineering
Indian Institute of Science
Bangalore 560012, India
amakur@ece.iisc.ernet.in

M. R. Raghuveer

Electrical Engineering
Rochester Institute of Technology
Rochester, NY 14623-0887, USA
mrreee@ritvax.isc.rit.edu

ABSTRACT

Approaches for digital halftoning of images using dithering threshold the input image with additive dithering noise. The paper presents a technique which thresholds the noise directly. The threshold is modified at each step such that the expected value of the output is equal to the input pixel's gray value. Further, error feedback is used to correct the threshold. Tests on images show the method's ability to retain features in the high frequency regions such as edges as well as low frequency features such as slow variations in intensity.

1. HALFTONING

Halftoning is the process of converting a gray scale image into a binary image that is perceived to retain its gray levels. A key component of multimedia efforts has been the need to display images on different types of output devices. In those instances where the devices have just binary output capability such as various printers and screen displays, it is natural to use halftoning to render gray scale images.

Since most of the image energy is present in the lower frequencies, a halftone that looks like the original image would have a similar low frequency spectrum. Consequently a halftone contains much more energy in the high frequency region than the original image. The idea here is that high frequency noise is less bothersome to the eye while contributing to the breaking up of undesirable patterns. One way to introduce such energy is to add high frequency noise to the image and then to threshold it to produce a binary output. This process is called dithering.

The error diffusion method of digital halftoning provides a useful benchmark. References [1] and [2] include a discussion of some of the advantages of the error diffusion process. For a halftone of a specified quality, the error diffusion method requires a smaller dot density than methods invented before. This process derives its name from the fact that the error in converting a gray value to a binary value (say 0 for black and 255 for white) is distributed to neighboring pixels. Thus when decision is to be made at a particular pixel whether a black or a white dot is to be placed, it is not just the gray value of the original image at that pixel which is compared with a threshold but it is the sum of the original gray value and the weighted average of past errors that is so compared. It can be shown that error diffusion tends to place dots such that with a constant input gray value, the average value of the halftone equals the input gray value [2]. The error diffusion method has some drawbacks such as "worm" artifacts and edge-blurring. Investigators have looked into alternative schemes to overcome these defects. See [2] [3] for example.

2. NOISE THRESHOLDING

In this paper we develop a noise based approach that provides an alternative to conventional dithering. Rather than thresholding the sum of the image and noise (as in dithering), or thresholding the image (as in error diffusion), we threshold the noise itself. Figure 1 shows a block diagram describing the approach. The noise process (through a shaping filter) is thresholded to produce the halftoned output. The threshold value depends only on the (feedforward filtered) image in the open loop approach. In the closed loop approach the threshold is modified taking into account the error between the (feedback filtered) halftone and the (feedforward filtered) input image. Let X denote the noise sample, I denote the input image ($0 \leq I \leq 1$), and H denote the halftone. In the open loop approach we set the threshold to that value which makes the expected value of the output equal to the input gray value. Let us call the corresponding threshold T_o. Since H takes only values 0 or 1,

$$\mathrm{E}\{H\} = p(H=1) = \int_{T_o}^{\infty} f_X(X)\, dX$$

where $f_X(X)$ is the pdf of the noise process. Therefore,

$$T_o = F_X^{-1}(1-I)$$

where $F_X(X)$ is the cdf of the noise process. For an uniform pdf from 0 to 1, it simplifies to $T_o = 1 - I$. For a gaussian pdf with mean $\frac{1}{2}$ and variance 1, it simplifies to

$$T_o = \begin{cases} \frac{1}{2} + \mathrm{erf}^{-1}(1-2I) & \text{if } 0 \leq I \leq \frac{1}{2} \\ \frac{1}{2} - \mathrm{erf}^{-1}(2I-1) & \text{if } \frac{1}{2} \leq I \leq 1 \end{cases}$$

work done at the Department of Electrical Communication Engineering, Indian Institute of Science, Bangalore

0-7803-2431-5/95 $4.00 © 1995 IEEE

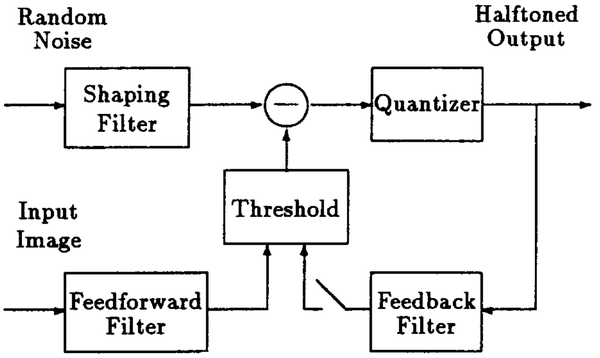

Figure 1: Block Diagram of the Noise Thresholding Approach

and for a symmetric triangular pdf extending from 0 to 1, the threshold is given by

$$T_o = \begin{cases} 1-\sqrt{\frac{I}{2}} & \text{if } 0 \le I \le \frac{1}{2} \\ \sqrt{\frac{1-I}{2}} & \text{if } \frac{1}{2} \le I \le 1 \end{cases}$$

The open loop approach, however, fails to produce reliable halftone for some given image because there is too much high frequency noise.

In the closed loop approach the threshold T is expressed as

$$T = f(T_o, e)$$

where $e = I - H$, the error between the halftone and the input. The idea is that the halftone H tracks the image I, and the threshold is changed in a manner such that the error e is minimized. One choice of the function that we have used in our simulations is $f(T_o, e) = T_o - e$. This in fact is a corrective measure by feeding back (diffusing) the error. Note that for the uniform pdf case, T is the value of the threshold that would have produced an expected halftone value of $I - e$, which means that the error of the previous instance has been diffused (with weight 1). The closed loop approach is expected to produce much better halftones due to the presence of feedback.

3. SHAPING FILTERS

One advantage the proposed method is its very general nature since it offers a wide choice of the three filters: noise shaping filter, feedforward filter and feedback filter. The noise shaping filter may be used to produce any type of coloured noise. Use of blue noise is popular in halftoning [4], and the same could be achieved by choosing a proper filter. For example, in some of our simulations we used the following high-pass filter: [−0.0156 0.0938 −0.2344 0.3125 −0.2344 0.0938 −0.0156]. It is possible in principle to use a two-dimensional filter for shaping a random noise field.

The feedforward filter offers a choice about which part of the input image spectrum is more important to be retained. Since the dot density of the halftone is equal to the pixel density of the image, one should not expect to retain all parts of the input spectrum. Instead, we may choose to put more emphasis on the low frequency spectrum. A low-pass feedforward filter will produce I, the local average of the gray values, which is expected to be retained by the halftoning process. The popular error-diffusion coefficients of size 2×3 and 3×5 are 2-D causal FIR low-pass filters, and we used the latter one for our simulations. Note that the feedforward filter may be non-causal and IIR.

For the closed loop approach, the feedback filter has the function on the output halftone as the feedforward filter had on the input image. Therefore, a natural choice is to have an identical filter for both feedforward and feedback path. This would imply that the system tries to produce a halftone having the part of the spectrum selected by the filter to be nearly identical to that of the gray scale image. However, this filter has to be causal, and it involves little implementational complexity since it operates on a binary input. These differences may lead to different choices of the feedforward and the feedback filters. Our simulations use identical feedforward and feedback filters.

4. RESULTS

The proposed noise thresholding method of halftoning is very easy to implement. With the closed loop approach it produces halftones which bear greater fidelity to the original than the error diffusion and the conventional dithering approaches.

An indication of the claims above can be seen in figures 2 and 3. While figure 2 shows M31, the Andromeda galaxy, figure 3 shows the popular face of Lenna. In each case, figure a shows the original (300 dpi halftone), figure b shows the error diffusion halftone, figure c shows a halftone generated by pseudo-random dithering, and figure d shows a halftone generated by noise thresholding. While all halftones for figure 2 have been printed with one dot per pixel of the original image at a spacing of 100 dpi, all halftones for figure 3 are with one dot per pixel at 75 dpi. The proposed method, although noise-based, produces the least noisy halftone while providing a good reproduction of features. For example, if we were to look at the central white region in figure 2, of the three methods shown, the proposed method provides the most faithful rendering in terms of gray level as well as the region's boundary.

5. REFERENCES

[1] E. Barnard, "Optimal error diffusion for computer-generated holograms," *J. Opt. Soc. Am. A*, vol.5, no.11, pp.1803–1817, November 1988.

[2] J. Sullivan, R. Miller and G. Pios, "Image halftoning using a visual model in error diffusion," *J. Opt. Soc. Am. A*, vol.10, no. 8, pp. 1714–1724, August 1993.

[3] S. Kollias and D. Anastassiou, "A unified neural network approach to digital image halftoning," *IEEE Trans. Sig. Proc.*, vol.39, no.4, pp.980–984, April 1991.

[4] T. Mitsa and K. J. Parker, "Digital halftoning technique using a blue-noise mask," *J. Opt. Soc. Am. A*, vol.9, no.11, pp.1920–1929, November 1992.

Figure 2a: Original Andromeda Galaxy

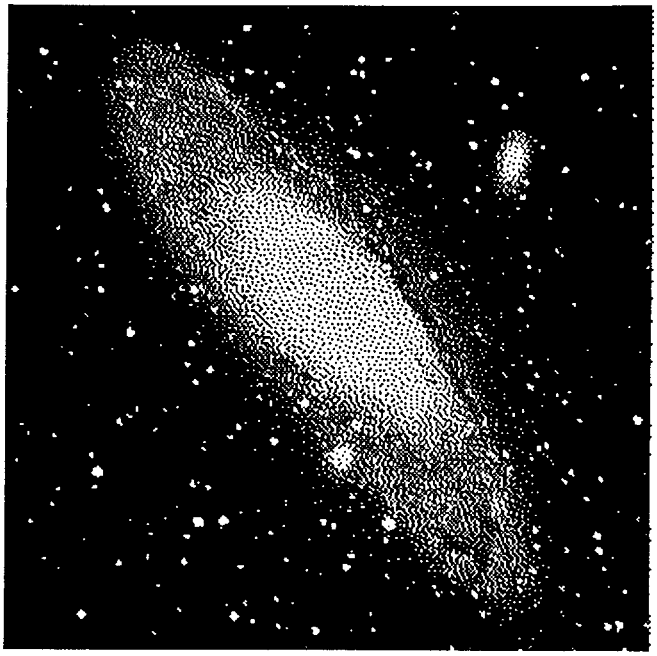

Figure 2b: Error Diffusion

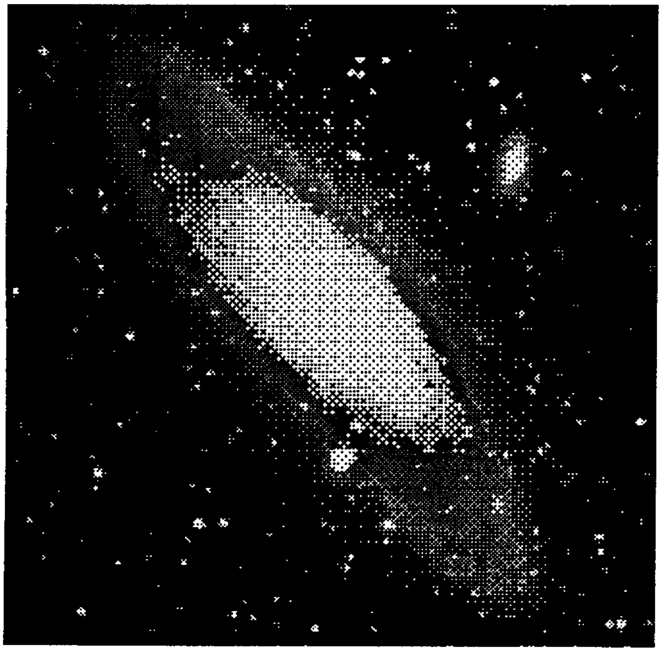

Figure 2c: Dithering

Figure 2d: Noise Thresholding

Figure 3a: Original Lenna

Figure 3b: Error Diffusion

Figure 3c: Dithering

Figure 3d: Noise Thresholding