Journal of Electronic Imaging 1(2), 143-151 (April 1992).

Inverse halftoning

Christopher M. Miceli
Kevin J. Parker
University of Rochester
Electrical Engineering Department
Rochester, New York 14627

Abstract. *Inverse halftoning is the method by which an approximation of a gray-scale image is reconstructed from a binary, halftoned version of the original. Several inverse-halftone algorithms are described, including a three-level cascade algorithm. We demonstrate that* a priori *knowledge of the halftone technique is not essential, but can be used if available. Finally, we demonstrate the results of applying inverse-halftone operations to both computer synthesized and photographic images.*

1 Introduction

Binary digital halftoning is the process of transforming an n-bit gray-scale image into a 1-bit binary image perceived to contain gray scale. Digital inverse halftoning is the reconstruction of a gray-scale image from its halftone rendering.

The desire to create an inverse halftone of a 1-bit image occurs in several simple and advanced applications. In some instances, an individual may wish to sharpen, enlarge, or rehalftone an image, but the only source available is a halftone rendering.

Most printed materials are produced using halftone techniques. Recently, many advances in digital scanning devices have taken place, including the introduction of digital duplicating devices. The ability for such devices to process halftone images, including the ability to convert between halftone algorithms, would be highly advantageous. The importance of manipulating halftone images is equally valuable to the desktop publishing industry. Elementary image processing operations such as filtering, decimation, interpolation, sharpening, etc., which are routinely implemented on gray-scale images, are not easily performed on halftone images. Methods for the processing of binary text[1] and binary images[2] have been discussed in the literature. These algorithms are limited to binary characters and are not applicable to high-frequency distributions of black and white, which produce the illusion of gray. There has been some interest in enhancement of halftone images,[3] but in general such operations require special care and may have limited performance.

In principle, the halftone (dither) procedure can be viewed as a thresholding operation. The original gray-scale image is passed through some sequence of operators resulting in the assignment of a 1 or a 0 (absence or placement of a black picture element). The inverse procedure can be seen as an inverse-quantization operation. Gray levels are derived by processing the binary image with some inverse-quantizing criteria (discussed in the following sections), resulting in the assignment of a discrete gray level. In a more complex treatment of the general problem, the steps of halftoning and inverse halftoning can be considered a lossy compression/decompression technique or an encryption/decryption technique.

Few previous reports in the open literature involve the reconstruction of gray-scale images from binary images.[4-8] These range from simple low-pass filtering to complicated neural networks. In general, low-pass filtering that is sufficient to smooth the impulses found in binary images will also blur sharp features. Hence, reconstructions based on low-pass filtering either remain grainy or become visually displeasing because of blurring. Neural networks require training and may not be optimal over many different halftone techniques. Furthermore, existing neural network tech-

Paper 91-039 received Dec. 27, 1991; revised manuscript received March 25, 1992; accepted for publication March 27, 1992.
1017-9909/92/$2.00.

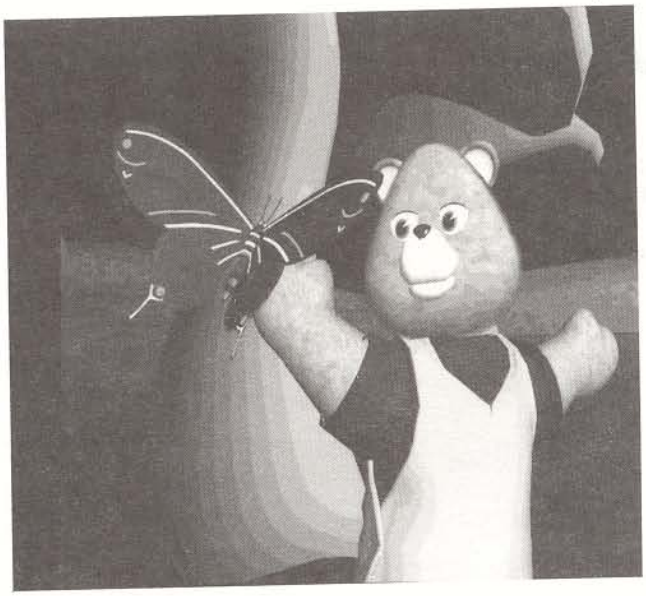

Fig. 1 Original gray-scale image of the "bear" image.

niques do not explicitly take advantage of all the *a priori* constraints on the gray-scale image (and power spectrum) or the nature of the halftone mask when its structure is available.

This paper begins with a brief review of the halftone process. We then turn our attention to solving the inverse-halftone problem. When a known halftone screen, or mask, is used in the halftone process, we show that knowledge of the mask can be used to generate estimates. However, we show that knowledge of the halftone process is not essential. Finally, we define a three-step cascade of operations leading to a gray-scale estimate of a halftone image.

Results are qualitatively expressed and compared using both photographs and data comparison plots. Data comparison plots consist of the overlay of original gray levels (solid line) and reconstruction gray levels (stars). The data are extracted from a horizontal scan line of the image in Fig. 1. The line is 300 from the top (out of 512) and cuts across discrete gray levels of the mushroom stalk. In the original gray levels (solid line), noticeable impulses are associated with each change in gray level. These impulses are the result of *sharpening* the edges in the image prior to halftoning. Results are quantitatively expressed in terms of their normalized mean square error (mse) defined as

$$\text{mse}=\frac{\sum_i\sum_j[g(i,j)-g'(i,j)]^2}{\sum_i\sum_j g(i,j)}\ , \tag{1}$$

where g represents the original gray-scale image, and g' represents the reconstruction. The mse is a simple quality measurement between the original and the distorted picture, which is commonly used in the literature.[7–9]

The images used to evaluate the presented inverse halftone techniques are "bear" (a 512×512 pixel computer-generated image) and "girl" (a 512×512 pixel "natural" image). All image processing is performed on a SUN SPARC station using single precision arithmetic. All images are processed in an X-window environment as 8-bit raster files. Images were converted from raster format into PICT format and transferred to a Macintosh computer where they were displayed and photographed using a Lasergraphics LFR plus camera recorder.

The following conventions are used throughout this paper. A gray-scale (multilevel) image is defined as a two-dimensional light intensity function $g(i,j)$, where i and j denote discrete pixel coordinates, and the value of g at any point (i,j) is proportional to the gray level of the image at that point. An 8-bit gray scale is used, providing $L=256$ possible gray levels lying in the interval $[0, L-1]$ where gray level $l=0$ is considered black and $l=255$ is considered white. A halftone (binary or bilevel) image is defined similarly, with the exception that its composition is 1 bit in nature. Hence, a halftoned image only contains two gray levels $[b_0,b_1]$, where b_0:b_1 may be considered either 0:1 or 0:255 corresponding to black/white. Preceding all reconstructions, it is assumed that an $M\times N$ pixel, 8-bit gray-scale image (of real and positive integers) has been halftoned into an $M\times N$, 1-bit binary image. The goal is to reproduce estimates of the original $M\times N$ 8-bit image that are

1. visually acceptable (edge and flat image regions are accurately reproduced and free of obvious or annoying artifacts
2. acceptable for elementary image processing algorithms
3. capable of allowing conversion between halftone methods.

We begin with a brief review of halftoning.

2 Halftoning

Halftoning is the process of transforming continuous gray-scale information into binary information perceived to contain continuous tone. This problem arises in many forms of media transfer, from graphic arts to facsimile machines. The binarization algorithm takes as input the sample $g(i,j)$ and generates the sample $b(i,j)=0$ or 1 for the binary image. This thresholding decision can be stated

$$b(i,j)=\begin{cases}0 & \text{if } g(i,j)\leq t\\ 1 & \text{if } g(i,j)>t\end{cases}\ , \tag{2}$$

where $g(i,j)$ is the gray level of a pixel with coordinate (i,j) and t is the threshold value, an element in $[0, L-1]$. The resulting image is referred to as a halftone image.

A complete analysis of halftone techniques is beyond the scope of this paper. Digital implementation of the halftone process is supported by a wealth of methods thoroughly discussed in several reviews and surveys.[10–14] Several existing techniques include ordered dither,[15] error diffusion,[16–18] and blue noise mask thresholding.[19] Figure 2 is a halftone rendition of Fig. 1 generated using a blue noise mask.[19]

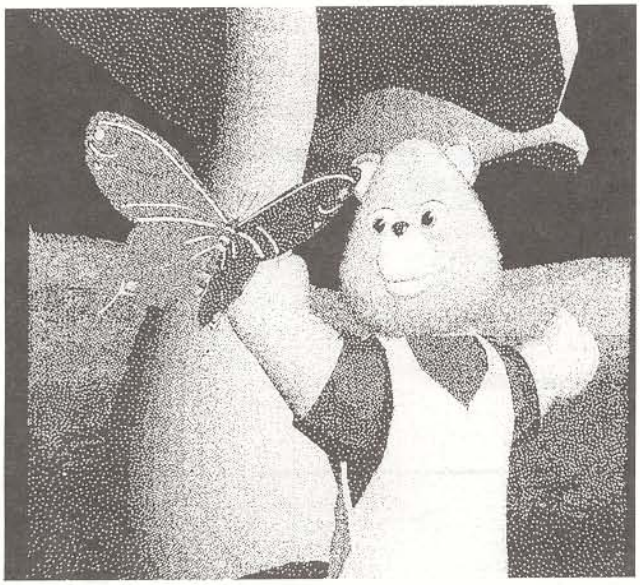

Fig. 2 Halftone of "bear" derived using a blue noise mask.

3 Inverse Halftoning

3.1 *Binary Information Only*

First, we investigate the process of reconstructing a gray-level image when the binary halftone image is the only information available.

Most halftone techniques are designed to reproduce accurately the average gray level of a "flat" or uniform region. That is, the proportion of 1's to 0's in an $n \times n$ region is directly related to the gray level of that region. A dithered image should contain, on the average, no error in the dc component. The most straightforward estimate of the gray-scale image, using only the halftone image, is a localized average of the halftone image. That is, if an image is a constant gray level g, where $0 \leq g \leq 1$, it is assumed that the binary image $b(i,j)$ has the property that

$$g(i,j) = \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} b(i,j) \ . \tag{3}$$

Thus, a simple localized estimate $g'(i,j)$ is given by

$$g'(i,j) = \frac{1}{(M+1)^2} \sum_{i=-M/2}^{M/2} \sum_{j=-M/2}^{M/2} b(i,j) \ . \tag{4}$$

This neighborhood approach effectively quantizes the binary image such that, for a window of dimension $n \times n$, $(n \times n)+1$ quantized levels are produced in the estimate image. A fundamental problem with using spatial low-pass filters on halftone images is the undesirable result involving the loss of edge information for increased approximations of the mean.

These results suggest the use of adaptive measures of local statistics in the simplest form. This can include run-lengths of 0's and 1's along any directional line in the halftone image. For run-lengths along rows and all b's $= 0$,

$$\text{if } b(i,j) = b(i,j+1) = \ldots = b(i,j+J) = 0 \ ,$$

$$\text{then } g'(i,j) = g'(i,j+1) = \ldots = g'(i,j+J) = \frac{1}{J+1} \ . \tag{5}$$

For example, on the average, a level of $g = 1/5$ will have a dispersed dot halftone representation of four 0's and one 1. Thus, a string of four zeros (bounded on both sides by 1's) is assumed to represent $g = 1/5$. There is no guarantee of any discernable concentration of 1's or 0's at (near) an edge. However, a change in the statistical appearance of the 1's and 0's will occur at (near) an edge. Since run-lengths are likely to terminate at high-contrast boundaries, this estimate will not necessarily blur image features. The run-length approach will, however, quantize the $g \leq 1/2$ estimate into integer ratios of 1/2, 1/3, 1/4, 1/5, etc. Where 1's represent the majority, a similar analysis handles the case for $g \geq 1/2$:

$$\text{if } b(i,j) = b(i,j+1) = \ldots = b(i,j+J) = 1 \ ,$$

$$\text{then } g'(i,j) = g'(i,j+1) = \ldots = g'(i,j+J) = \frac{J}{J+1} \ . \tag{6}$$

where estimates are quantized into the integer ratios of 1/2, 2/3, 3/4, 4/5, etc. Equations (5) and (6) apply to rows (g'_r), but can easily be reindexed and applied to columns (g'_c), the results of which can be combined as a mathematical average

$$g'(i,j) = \frac{g'_r(i,j) + g'_c(i,j)}{2} \ . \tag{7}$$

Figure 3 contains a block diagram of our adaptive binary run-length (ABRL) for inverse halftoning. This approach could also be applied to two-dimensional regions, as opposed to simple rows and columns.

Figure 4 contains a reconstruction of the bear using ABRL. Figure 5 shows the corresponding data comparison plot. In comparing the ABRL estimates to the spatial low-pass estimates, we find that by using the ABRL inverse-halftone algorithm, edges are better preserved.

3.2 *Single Pixel Operations For Known Masks*

We now turn our attention to situations in which information can be derived from knowledge of a known halftone mask. Netravali and Bowen[5] approximated a gray-level image from its dithered image using information supplied by the halftone matrix. Their estimate of the gray-level image was achieved using a repeated, 4×4 dispersed dot dither matrix and a 3×3 neighborhood operation. We now show that, when a halftone mask is known, the range of possible gray values for the estimated gray level is reduced.

In principle, single-pixel estimates of the gray-scale image can be derived from single pixels of the binary image when the halftone mask $h(i,j)$ is known. This is because the basic decision rule for halftoning is stated as

$$b(i,j) = \begin{cases} 1 & \text{if } g(i,j) > h(i,j) \\ 0 & \text{if } g(i,j) \leq h(i,j) \end{cases} \ . \tag{8}$$

Thus, our single-pixel estimate (with h normalized, $0 < h < 1$) is described by

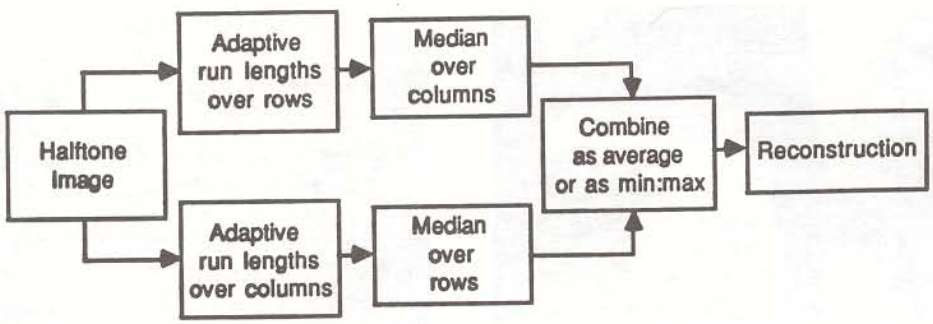

Fig. 3 Block diagram for ABRL algorithm.

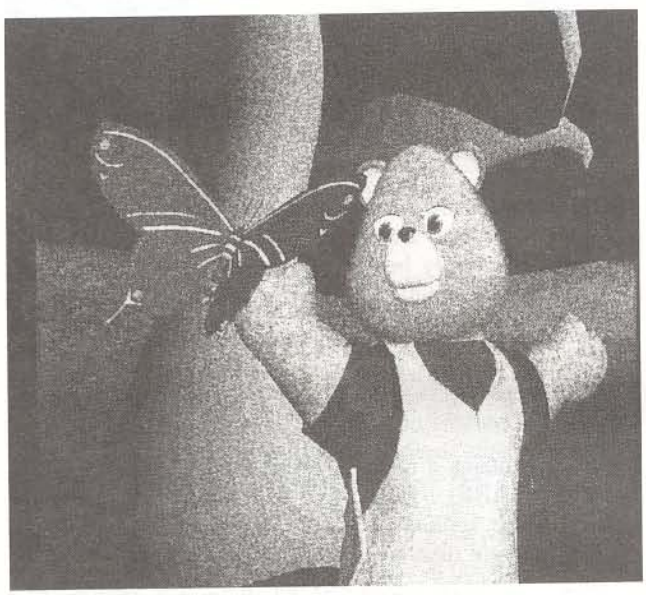

Fig. 4 Reconstruction using run-lengths combined as (row + col)/2.

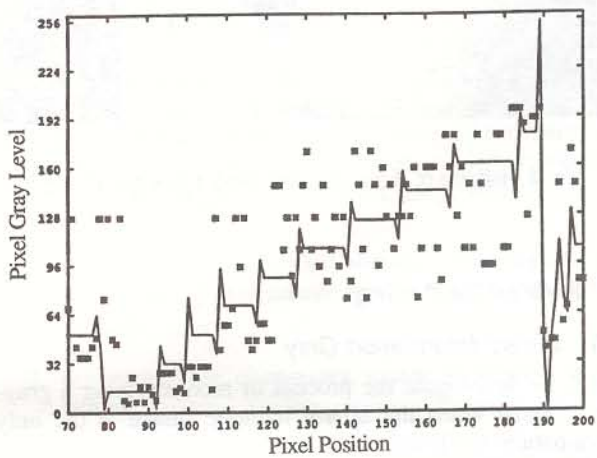

Fig. 5 Original gray levels (solid line) and reconstruction using run-lengths combined as (row + col)/2 (boxes) from scan line 300 of "bear."

$$g'(i,j) \Rightarrow \begin{cases} h(i,j) < g' < 1 & \text{if } b(i,j)=1 \\ 0 \leq g' \leq h(i,j) & \text{if } b(i,j)=0 \end{cases} , \tag{9}$$

where the estimate is taken to be the expected value of the given range. The precision of any estimate is linked to the range of numbers associated with that pixel. That is, if $b(i,j)$ is 1 and $h(i,j)$ is 0.6, then the range of possible values of $g'(i,j)$ is between 0.6 and 1.0, an undesirably large range. Without other information, the appropriate estimate would use the half range point, or 0.8 in this example.

Application of the single-pixel operation results in a range compression or "flatness" in the resultant image. This loss of dynamic range results from the use of the midpoint of the estimated range. That is, when $b(i,j)=1$,

$$g'(i,j) = \frac{1}{2}[1 - h(i,j)] + h(i,j) \ . \tag{10}$$

Thus, the expected value of $g'(i,j)$ for these pixels is

$$E[g'(i,j)] = \frac{1}{2}\{1 - E[h(i,j)]\} + E[h(i,j)] \ . \tag{11}$$

For linear, symmetric halftone masks $E[h(i,j)] = 1/2$, and

$$E[g'(i,j)] = \frac{1}{2}\left(1 - \frac{1}{2}\right) + \frac{1}{2} = 0.75 \ . \tag{12}$$

Similarly, when $b(i,j)=0$,

$$g'(i,j) = \frac{1}{2}[h(i,j)] \ , \tag{13}$$

$$E[g'(i,j)] = \frac{1}{2}\{E[h(i,j)]\} \ , \tag{14}$$

$$E[g'(i,j)] = \left(\frac{1}{2}\right)\left(\frac{1}{2}\right) = 0.25 \ . \tag{15}$$

Hence, in regions of the original image where $g(i,j)=0$, the average estimate obtained using single-pixel comparisons with the halftone is 0.25. Where $g(i,j)=1$, the average estimate is 0.75. This 50% gray-scale compression results in a "flat" or dynamically restricted estimate of the image. This compression can readily be rescaled.

Because this is a point operation, no additional blurring associated with neighborhood operations is introduced into

the reconstruction. The information obtained in performing the single-pixel operation is not limited to the estimated pixel values, $g'(i,j)$. Two additional constraints that may be used in further processing of the image are outlined below.

Constraint 1. A "goodness" or weight value, which distinguishes good estimates from poor estimates, can be assigned to each estimated pixel. For $g(i,j)$ approaching 0 and $g(i,j)$ approaching 1, the weight value increases as

$$\left.\begin{matrix}\forall\ b(i,j)=1 \cap h(i,j)\rightarrow 1.0\\ \forall\ b(i,j)=0 \cap h(i,j)\rightarrow 0.0\end{matrix}\right\}\Rightarrow(w\rightarrow 1) \tag{16}$$

and decreases as

$$\left.\begin{matrix}\forall\ b(i,j)=1 \cap h(i,j)\rightarrow 0.0\\ \forall\ b(i,j)=0 \cap h(i,j)\rightarrow 1.0\end{matrix}\right\}\Rightarrow(w\rightarrow 0)\ . \tag{17}$$

This information can be used to approximate the degree by which a pixel value is in error, and by how much its value should be adjusted.

Constraint 2. Range information is available for each estimated pixel. This constraint can be used to set boundaries, keeping a pixel from taking on impossible values when further processing is performed:

$$\begin{aligned}&\forall\ b(i,j)=1\Rightarrow h(i,j)<r<1\\ &\forall\ b(i,j)=0\Rightarrow 0\leq r\leq h(i,j)\ .\end{aligned} \tag{18}$$

Great increases in the precision of the (local) estimate can be obtained when two or more pixels of the binary image are used in conjunction with the corresponding pixels of the halftone mask. That is, assuming that the image is slowly varying in a region R, and the $b(i,j)$ are 1, we have

$$\forall\ b(i,j)_{\in R}=1\Rightarrow\max\{h(i,j)_{\in R}\}<g'<1\ . \tag{19}$$

Analogous rules can be derived for the case where the $b(i,j)$ are 0,

$$\forall\ b(i,j)_{\in R}=0\Rightarrow 0\leq g'\leq\min\{h(i,j)_{\in R}\}\ . \tag{20}$$

For mixed results, where b's are 0 and 1 within a neighborhood, we have

$$\begin{aligned}&\forall\ b(i,j)_{\in R}=1\Rightarrow x_1=\max\{h(i,j)_{\in R}\}\ ,\\ &\forall\ b(i,j)_{\in R}=0\Rightarrow x_0=\min\{h(i,j)_{\in R}\}\ ,\\ &\min(x_0,x_1)\leq g'\leq\max(x_0,x_1)\ ,\end{aligned} \tag{21}$$

providing new minimum and maximum bounds on g'. Note that if g is truly uniform within the region, then x_0 will be greater than x_1.

To demonstrate the advantage of the adaptive maximum and minimum operations, let us assume that the halftone mask is uniformly distributed between 0 and 1. We wish to determine the probability distribution function for $\max\{h_1,h_2,\ldots,h_x\}$, assuming all h's to be independent. For the case $x=2$, the cumulative distribution function $F_{g'}(g')$ is given by

$$\begin{aligned}F_{g'}(g')&=P(G'\leq g')=P[\max(H_1,H_2)\leq g']\\ &=P(H_1\leq g'\cap H_2\leq g')\\ &=P(H_1\leq g')P(H_2\leq g')\\ &=F_{H1}(g')F_{H2}(g')\ ,\end{aligned} \tag{22}$$

and the probability distribution function (pdf) is found by taking the derivative

$$f_{g'}(g')=F_{H1}(g')f_{H2}(g')+F_{H2}(g')f_{H1}(g')\ . \tag{23}$$

For uniformly distributed h, the resulting pdf $f_{g'}(g')$ is more highly weighted toward 1.0. Assuming the h's are independent and identically distributed,

$$f_{g'}(g')=2F_H(g')f_H(g')\ , \tag{24}$$

and $E(g')=2/3$. This result naturally extends for the case of x random variables. As the number of pixels used increases, the expected values of g' more closely approach the original gray level, and dynamic range is less compressed. This argument reveals where the blue noise mask is superior to conventional screen halftones for inverse halftoning. In the blue noise mask, adjacent pixels are highly uncorrelated. In conventional masks there is a great degree of correlation between neighboring pixels. Hence, region size required to gain independent information is increased. Independent multiple-pixel estimates can be obtained from different configurations.

We have evaluated a number of window arrangements where the halftone screen was known. [20]For these groupings it is determined that a 2×2 pixel window provides maximum results regarding trade-offs between increased blurring and minimized gray-scale compression.

Information introduced using neighborhood operations with the blue noise mask can be summarized as follows:

1. Because neighborhood operations imply low-pass filtering, some degree of image blurring occurs, the amount of which is dependent on the size of the neighborhood.
2. Range compression is greatly reduced, thus the original range of gray-scale values is more closely preserved.
3. "Goodness," or weight criteria, and range information can no longer be applied with certainty as in the single-pixel case. This is because an estimated value is no longer associated with a given pixel, but rather with the statistics of the neighboring pixels.

While it is true that the total number of gray levels is reduced, it is common image processing knowledge that human eye sensitivity requires only 6 bits or 64 gray levels to simulate continuous shading and avoid contouring.[21–24] The 1-pixel operation meets this criteria. Assuming the mask is known, we can treat the 1-pixel reconstruction as a significant improvement gained from trivial operations. No additional blurring is added, while a better approximation to gray scale is achieved.

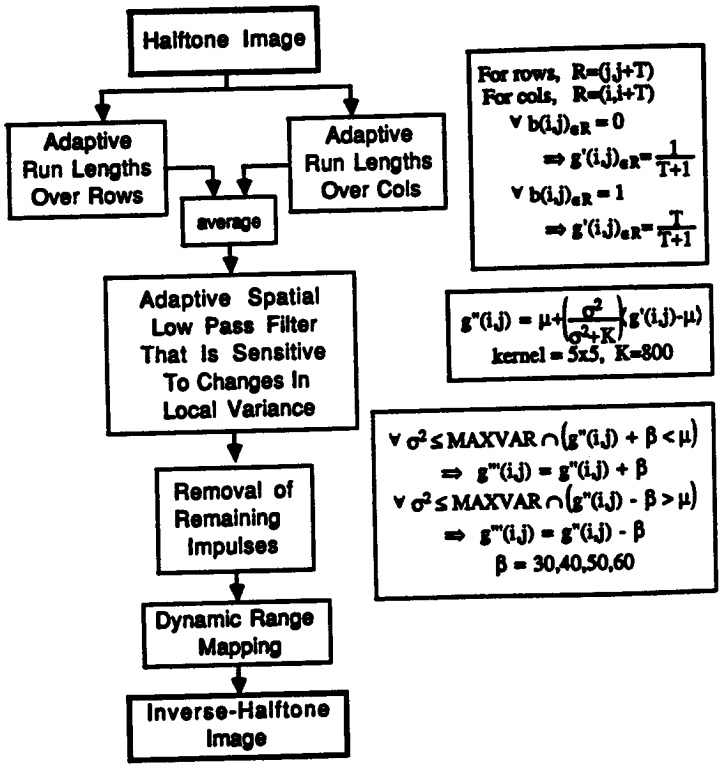

Fig. 6 Block diagram of cascaded inverse-halftone algorithm.

3.3 *Cascade Algorithm*

It is clear from the preceding discussion that a number of primary estimates of g' and postprocessed estimates of g'' can be generated. Several other estimation methods (i.e., iterative, adaptive low pass, and arithmetic mean over several results) are described in Miceli.[20] Each technique described above has unique features, and produces unique reconstructions of the gray-scale image. Some techniques perform better in terms of visual perception, others in terms of mse, but in many cases the differences are subjective, and no one algorithm stands out as exceptional. By combining different reconstructions or techniques, it is hoped that the "strong" features of one reconstruction will cancel or subtend the "weaknesses" associated with another. It is with this consideration that we approach the cascade algorithm.

We have determined a three-step sequence that produces visual results superior to the results of any one of the algorithms presented thus far. This sequence of operations can be applied to images halftoned by most dispersed dot halftone methods, and it works particularly well for blue noise mask and error diffusion techniques. The greatest attribute to this cascade is that no information about the halftone mask or method is required (other than that the halftone dots be dispersed).

The first stage involves the application of the previously described adaptive run-length algorithm to the halftone image. This produces an initial conversion of binary to gray without adding noticeable blur. Some dot structure seen in the halftone image is maintained. This is high frequency and not displeasing. The introduction of gray-scale information (with minimal blurring) provides an increased capacity for edge detection, which leads us to the second stage, smoothing by means of a local statistical-based smoothing algorithm. Since edges are more easily detected, such smoothing operators can yield better results than those operating on the halftone image directly. Lee's additive noise filter[25] is used with a 5×5 kernel and K set at 800 (see Fig. 6). Mu and sigma are the local mean and standard deviation calculated within the 5×5 kernel. For these values, an acceptable degree of smoothing occurs throughout the image. Acceptable implies that the filtered image does not contain artifacts associated with low-pass operations such as blotching or blurring. The third step is an impulse remover used to locate and make relative adjustments to any pixels not associated with edges, which deviate strongly (± 30 or greater) from the expected value of the neighborhood. Only a small percentage of the pixels are modified (in our examples, fewer than 2.0%). Figure 6 shows a block diagram representation of the cascade algorithm, where MAXVAR is an adjustable threshold on local variance, below which the algorithm adjusts pixel values.

Our algorithm as thus far described, does not accurately reproduce gray levels near the extremes. Reproducing the extreme gray-scale levels would require unusually long run-lengths. For an arbitrary number of consecutive black pixels, for example 10, the region may be associated with a small area of the image that is near solid black ($g=0$), but the limited area run-length quantizes the pixels to $1/(10+1)$, or $g'=23$. Hence, we remap the pixels near the extremes (0 and 255) with a simple exponential shift toward these extremes.

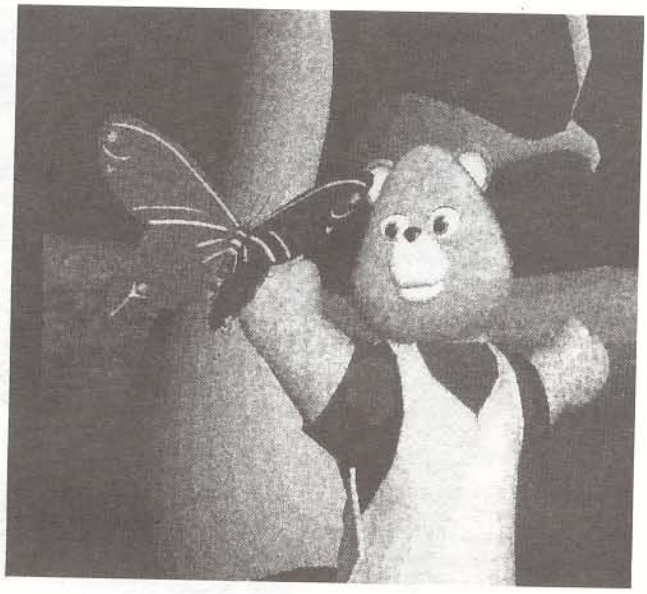

Fig. 7 Reconstructed gray-scale of the "bear" image, derived from the halftone rendering in Fig. 2 by means of the cascaded algorithm, which does not make use of mask information.

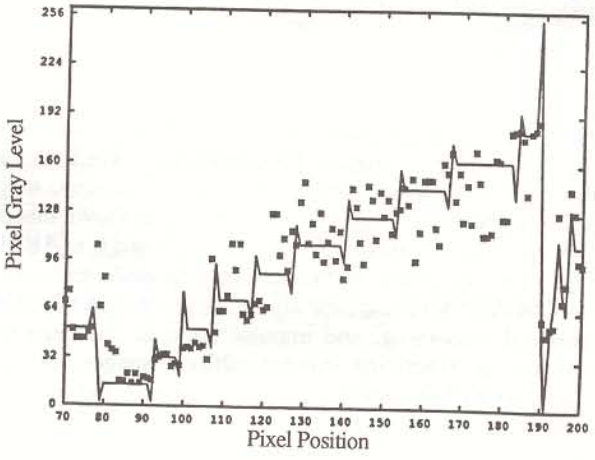

Fig. 8 Original gray levels (solid line) and reconstruction using the cascading of three different inverse-halftone algorithms (boxes) from scan line 300 of the "bear" image.

Figures 7 and 8 illustrate the gray-level reconstruction of the bear derived using this technique. Notice the "soft" quality of the reconstructed image, the accuracy of reproduction at the edge and flat image regions, and the lack of obvious and annoying artifacts.

This sequence of operations can be modified to include use of the mask information, if it is available. However, it is advantageous to be able to reconstruct without knowledge of the mask. For example, halftone images generated using error diffusion techniques, which do not use a mask, can be reconstructed by means of this technique. Also, the inclusion of mask information may lead to beating and aliasing problems if the mask is not properly realigned to the position it was at when halftoning was performed.

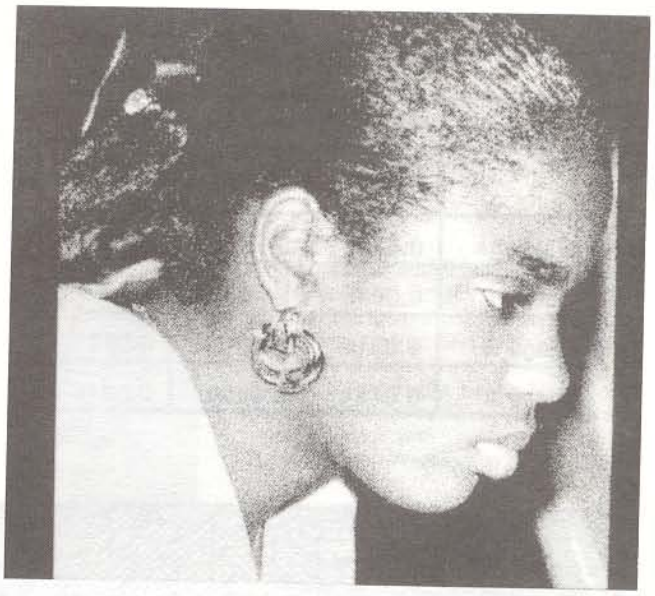

Fig. 9 A 512×512 pixel halftone image of "girl," derived using a blue noise mask.

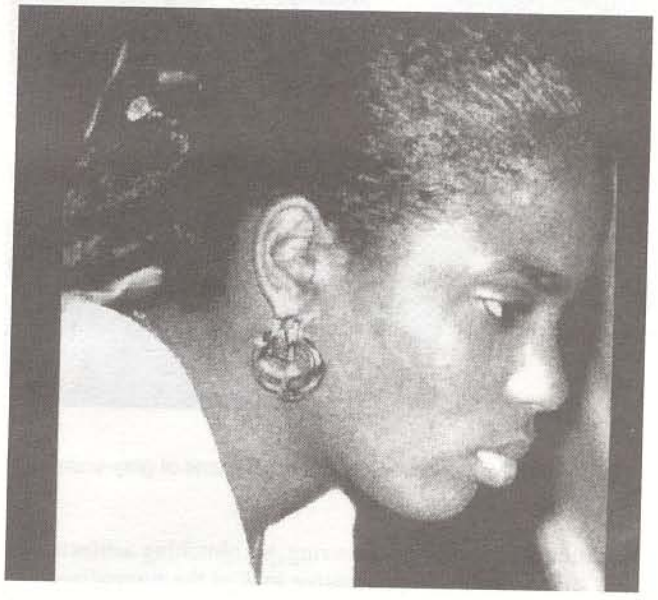

Fig. 10 A 512×512 pixel gray-level reconstruction of halftone "girl." Reconstruction derived using the cascade algorithm.

4 Results

At the beginning of this paper, we presented three goals that we believed inverse halftoning would play an important role in achieving. Those goals were to produce a visually acceptable gray-scale rendition of the halftone image suitable for continuous display devices, to allow elementary image processing algorithms to be applied to binary images, and to allow the conversion between halftone methods.

Figures 9 (halftone) and 10 (inverse-halftone) highlight the results of the cascade algorithm. Notice how the halftone image has been softened and appears more natural with little

Table 1 The mse associated with each stage of the cascade algorithm for the "bear," "girl," and "Lena" inverse-halftone reconstructions.

Normalized mse values for cascade algorithm			
	Bear	Lena	Girl
Halftone	0.6788	0.7645	0.3205
Stage 1	0.0949	0.0839	0.0394
Stage 2	0.0514	0.0413	0.0225
Stage 3	0.0464	0.0380	0.0216

Fig. 11 A 512×512 pixel clustered dot halftone of gray-scale "girl."

additional blur, false contouring, or blotching artifacts. Table 1 provides a comprehensive look at the normalized mse for both "bear" and "girl" at each stage of the cascade algorithm, as well as the well-known "Lena" image.

A 512×512 clustered dot halftone of the original gray-scale girl image is shown in Fig. 11. Figure 12 shows the results of converting the blue noise mask-generated halftone image of "girl" in Fig. 9 into a clustered dot halftone image by using the reconstructed gray-scale girl image shown in Fig. 10. Thus, it is possible, using the techniques described herein, to interconvert from dispersed dot to clustered dot techniques. However, when the original binary image is clustered dot, the interconversion is less successful because the inverse halftone reconstruction is usually of poorer quality.

5 Conclusions

A comprehensive investigation of digital inverse-halftone techniques has been presented. It was shown that for dispersed dot halftone techniques, information about the mask was not essential for generating reasonable estimates of the original gray-scale images. Fundamentally, when the only information available is the binary image itself, some neighborhood operation must be used. It has been shown that one such operation, adaptive run-lengths of 1's and 0's (ABRL), is particularly useful in facilitating gray-scale reconstructions. A three-level cascade algorithm comprised of ABRL, statistical smoothing, and impulse removal has been very effective in generating inverse-halftone images. In cases where the halftone mask structure is known, the mask can be used to yield improved gray-scale estimations of halftone images.

Fig. 12 Clustered dot halftone of "girl" derived by conversion from blue noise mask halftone of image in Fig. 9 (facilitated by inverse halftoning).

A general problem of halftone images is that they are not easily manipulated. Inverse halftoning has been found to be a plausible method for facilitating halftone manipulation. In some instances, manipulation complexity may depend on the exactness by which the original gray-scale image can be approximated. However, it should be restated that reconstructions do not need to be exact. The techniques presented here have been successful in allowing halftone images to be displayed as gray-scale images, allowing the application of elementary image processing algorithms to halftone images, and allowing interconversion of halftone techniques. We believe that the results presented here are not yet optimum, and that further investigation of the inverse-halftone process can lead to increased accuracy in, and capabilities for, halftone manipulation.

Acknowledgment

This work was supported by the University of Rochester Department of Electrical Engineering.

References

1. R. A. Ulichney and D. E. Troxel, "Scaling binary images with the telescoping template," *IEEE Trans. Pattern Anal. Machine Intell.* **PAMI-4**(3), 331–335 (May 1982).
2. I. T. Young, R. L. Peverini, P. W. Verbeek, and P. J. Van Otterloo, "A new implementation for the binary and Minkowski operators," *Comput. Graphics Image Proc.* **17,** 189–210 (1981).
3. S. Akira, "Image enhancement in a dithered picture," *Comput. Vision Graphics Image Proc.* **24,** 107–113 (1983).
4. C. N. Judice, "Method and arrangement for eliminating flicker interlaced ordered dither images," U.S. Patent 3,953,668 (April 1976).
5. A. N. Netravali and E. G. Bowen, "Display of dithered images," *Proc. SID* **22**(3), 185–190 (1981).
6. Y. Mita, S. Sugiura, and Y. Shimomura, "High quality multi-level image restoration from bi-level image," *Proc. Sixth Int. Congress Advances in Non-impact Printing Technologies Black & White and Color,* Orlando, Florida, October 21–26, 1990, pp. 235–236.
7. P. G. Roetling, "Unscreening stored digital halftone images," U.S. Patent 4,630,125 (1986).
8. Z. Fan, "Unscreening of stored digital halftone images by logic filtering," U.S. Patent 5,027,078 (1989).
9. J. O. Limb, "Distortion criteria of the human viewer," *IEEE Trans. Syst. Man Cybern.* **SMC-9**(12), 778–836 (Dec. 1979).
10. J. F. Jarvis, C. N. Judice, and W. H. Ninke, "A survey of techniques for the display of continuous tone pictures on bilevel displays," *Comput. Graphics Image Proc.* **5**(1), 13–40 (1976).
11. P. Stucki, "Image processing for document reproduction," in *Advances in Digital Image Processing,* pp. 177–218, Plenum Press, New York (1979).
12. J. P. Allebach, "Visual model-based algorithms for halftoning images," *Proc. SPIE* **310,** 151–158 (1981).
13. J. C. Stoffel and J. F. Moreland, "A survey of electronic techniques for pictorial reproduction," *IEEE Trans. Comm.* **COM-29**(12), 1898–1925 (Dec. 1981).
14. R. A. Ulichney, *Digital Halftoning,* MIT Press, Cambridge (1987).
15. B. E. Bayer, "An optimum method for two-level rendition of continuous tone pictures," *IEEE 1973 Intl. Conf. Communications,* Vol. 1, pp. 26–11 to 26–15 (1973).
16. R. Floyd and L. Steinberg, "An adaptive algorithm for spatial gray scale," *Proc. SID* **17**(2), 75–77 (1976).
17. R. A. Ulichney, "Dithering with blue noise," *Proc. IEEE* **76**(1), 56–79 (Jan. 1988).
18. D. Anastassiou and S. Kollias, "Digital image halftoning using neural networks," *Proc. SPIE* **1001,** 1062–1069 (1988).
19. T. Mitsa and K. J. Parker, "Digital halftoning using a blue noise mask," *Proc. SPIE* **1452,** 47–56 (1991).
20. C. M. Miceli, "Inverse halftoning," Master thesis, University of Rochester (submitted May 1991).
21. L. G. Roberts, "Picture coding using pseudo-random noise," *IRE Trans. Info. Theory* **IT-8,** 145–154 (1962).
22. A. K. Jain, *Fundamentals of Digital Image Processing,* Prentice Hall, Englewood Cliffs, NJ (1989).
23. R. C. Gonzalez and P. Wintz, *Digital Image Processing,* Addison-Wesley, Reading, MA (1977).
24. J. S. Lim, *Two-Dimensional Signal and Image Processing,* Prentice Hall, Englewood Cliffs, NJ (1990).
25. J. Lee, "Digital image enhancement and noise filtering by use of local statistics," *IEEE Trans. Pattern Anal. Machine Intell.* **PAMI-2**(2), 165–168 (Mar. 1980).

Christopher M. Miceli received the BS from the University of Buffalo, New York in 1989 and the MS from the University of Rochester, New York in 1991, both in electrical engineering. He is a member of both the Tau Beta Pi and Eta Kappa Nu associations. Miceli is currently an associate software engineer with General Electric Aerospace in the Washington, D.C., area.

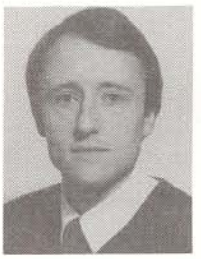

Kevin J. Parker received the BS degree in engineering science, summa cum laude, from SUNY at Buffalo in 1976. Graduate work in electrical engineering was done at MIT, with MS and PhD degrees received in 1978 and 1981. From 1981 to 1985 he was an assistant professor of electrical engineering at the University of Rochester; currently he holds the title of professor of electrical engineering and radiology. Dr. Parker has received awards from the National Institute of General Medical Sciences (1979), the Lilly Teaching Endowment (1982), the IBM Supercomputing Competition (1989), and the World Federation of Ultrasound in Medicine and Biology (1991). He is a member of the IEEE Sonics and Ultrasonics Symposium Technical Committee and serves as reviewer and consultant for a number of journals and institutions. He is also a member of the IEEE, Acoustical Society of America, and the American Institute of Ultrasound in Medicine. Professor Parker's research interests are in medical imaging, linear and nonlinear acoustics, and digital halftoning.