COMPUTER GRAPHICS AND IMAGE PROCESSING 6, 184–195 (1977)

NOTE

Image Enhancement by Histogram Transformation *

ROBERT HUMMEL

Computer Science Center, University of Maryland, College Park, Maryland 20742

Received January 7, 1976

A number of simple and inexpensive enhancement techniques are suggested. These techniques attempt to make use of easily computed local context features to aid in the reassignment of each point's gray level during histogram transformation.

The visual contrast of a digitized image can be improved by adjusting the grayscale so that the histogram of the resulting image is flat. This increase in visual contrast is an artificial enhancement due to a more judicious choice of the relative intensities represented by the quantization levels of the original image. Although the process provides no enhancement in an information-theoretic sense, applying the appropriate transformation can result in a remarkable increase in visual clarity.

As shown in an earlier paper [1], histogram modification can be accomplished quite simply by using a table look-up transformation applied to each point of the picture. The transformation which comes closest to providing a uniform distribution can be obtained from a scaled version of the cumulative distribution function, i.e., the integral of the original histogram. If we require this transformation to be single valued, then gray-level bins of the original histogram can only be merged together, but no bin can be broken up. Thus if the original image has a histogram with large peaks, the transformed histogram will have only a very approximately flat histogram. Nonetheless, contrast is improved because gray-level bins with large numbers of points are moved farther apart, thereby increasing the discriminability between the corresponding constant gray-level sets. We will refer to this kind of (single-valued) transformation as *histogram equalization*. Examples are shown in Fig. 1.

Several different techniques have been suggested to modify this transformation process so as to obtain pictures with flatter histograms. Any such modification involves defining a multiple-valued transformation. That is, in order to fill the gray-level bins with exactly the right numbers of points, some bins in the original image will have to be broken up. Suppose that there are n points in the original

* The support of the U.S. Air Force Office of Scientific Research, under Contract F44620-72C-0062, is gratefully acknowledged, as is Shelly Rowe for her help in preparing this paper.

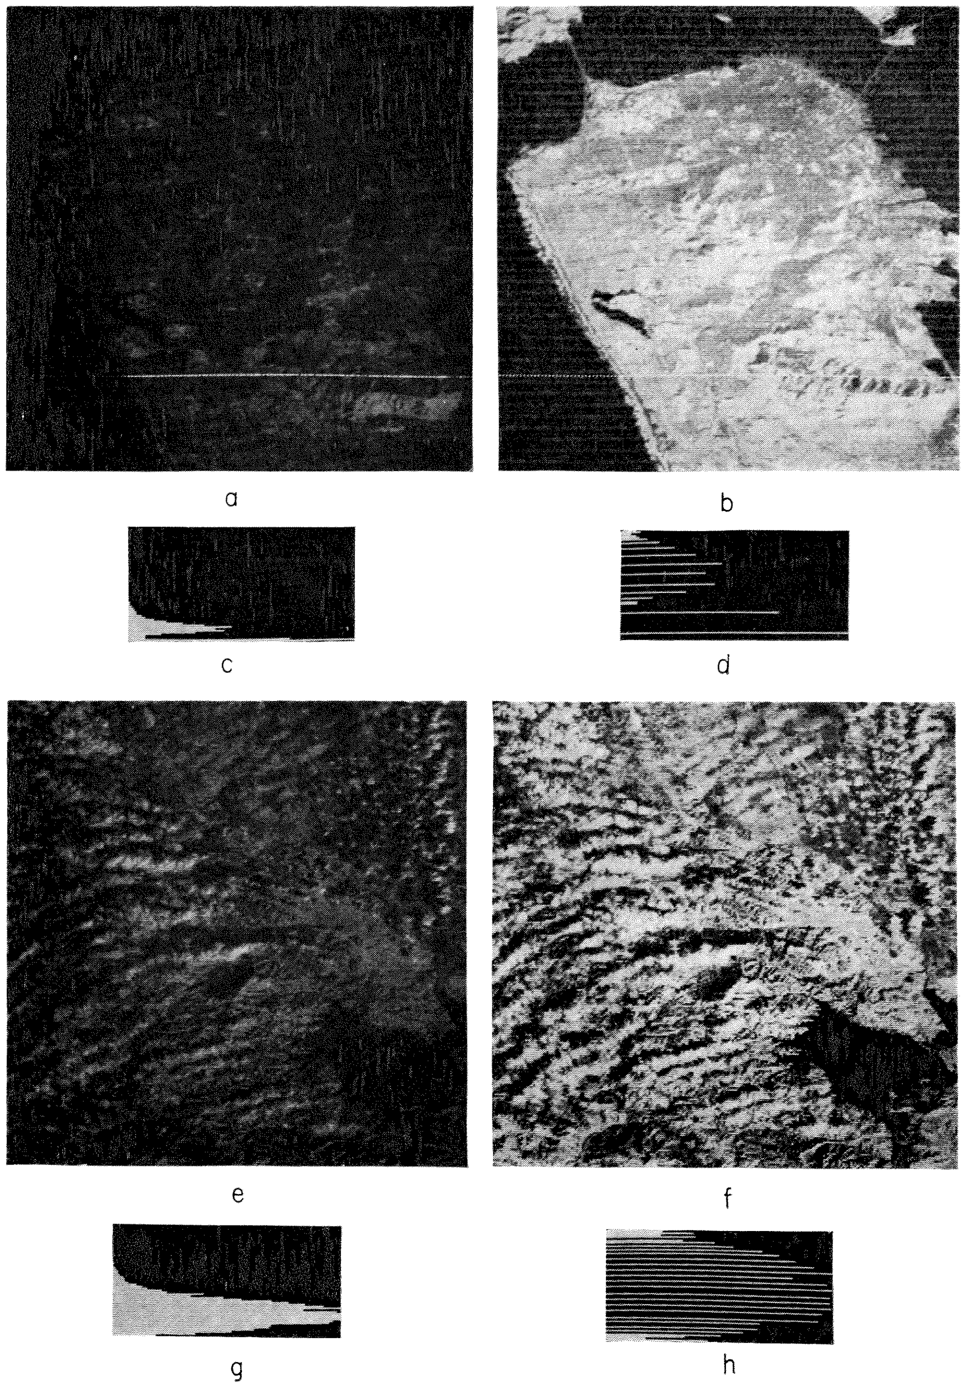

FIG. 1. Examples of histogram equalization: (a) Original. (b) Histogram equalized version of (a). (Note enhancement of quantization banding error in the bay and ocean regions.) (c) Histogram of (a). (d) Histogram of (b). (e) Original. (f) Histogram equalized version of (e). (g) Histogram of (e). (h) Histogram of (f).

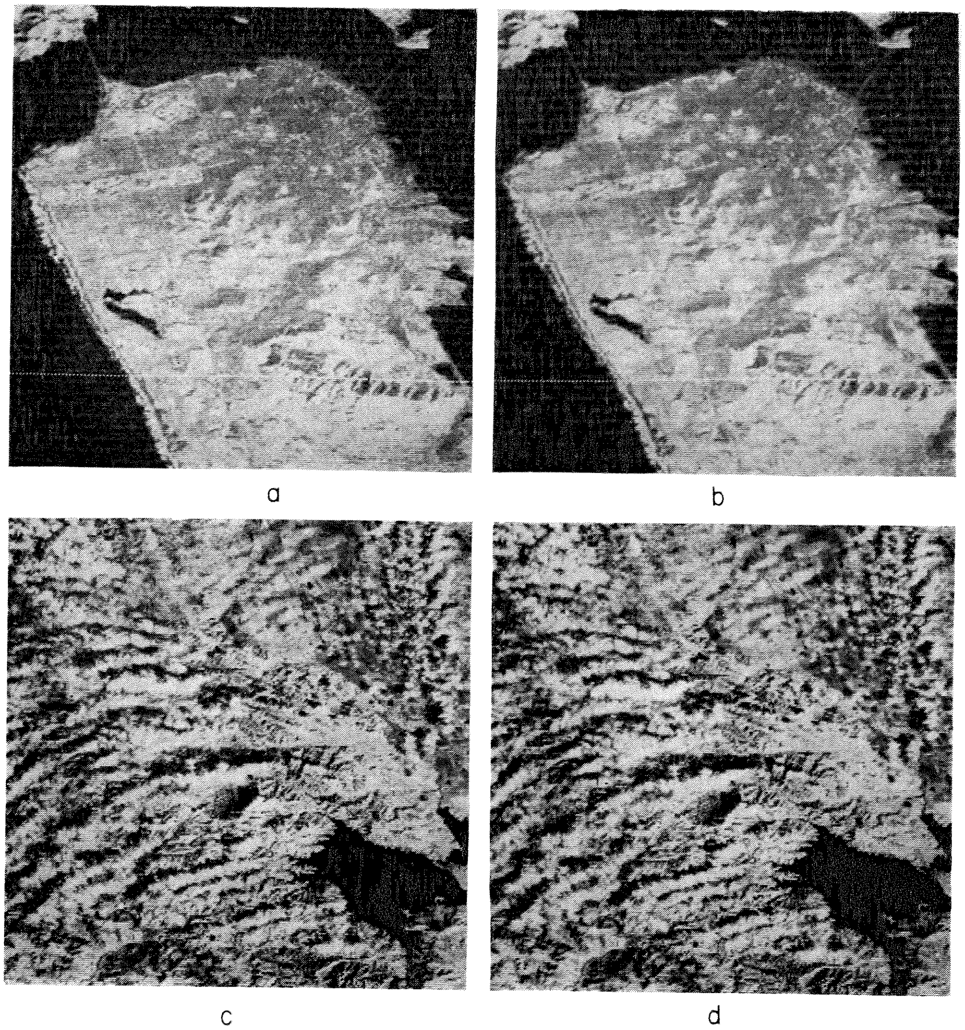

FIG. 2. Examples of histogram flattening: (a) Histogram flattened version of Fig. 1a using random selection tie breaking. (b) Histogram flattened version of Fig. 1a using neighborhood average tie breaking. (c) Histogram flattened version of Fig. 1e using random selection tie breaking. (d) Histogram flattened version of Fig. 1e using neighborhood average tie breaking.

image with gray level g, and that in order to create a new version with a flat histogram, the bin must be broken up so that n' points are assigned gray level g', and the remaining $n'' = n - n'$ points are assigned gray level $g' + 1$. One method of performing the necessary split is to choose these n' points randomly from among the n points with gray level g. A more sophisticated method is to do initial tie breaking on the average gray level in (say) the 3 by 3 neighborhood around each point, and then resolve any remaining ties by the random choice method. Thus, the n' points with the smallest neighborhood averages are assigned the gray level g', and the remaining n'' points with relatively greater neighborhood averages are assigned level $g' + 1$. Examples of pictures transformed by these two methods, which we will refer to as *histogram flattening*, are shown in Fig. 2.

Although the histograms of the pictures in Fig. 2 are exactly flat, the visual results are virtually indistinguishable from the enhancement in Fig. 1. Since histogram flattening is several times more costly than histogram equalization, it is clear that the tie-breaking routines used to produce Fig. 2 cannot be justified in terms of visual enhancement. This is not surprising, since the random selection and the neighborhood average methods of tie-breaking add no essential information, and in fact slightly degrade the information already present in the original image. In terms of the quantization levels represented in the final picture, tie-breaking is analogous to minor adjustments of the cutoff points between quantization levels. Although this can affect the appearance of the histogram drastically, the net effect on the appearance of the cumulative distribution function, which is probably a better indicator of visual contrast, is minimal.

By making use of local context information for each picture point, it is possible to further increase the visually extractable information in a digitized picture. In particular, preprocessing techniques which enhance edges and lines, reduce noise, and correct for illumination gradients can be combined with histogram equalization to yield the effect of a knowledge-based tie-breaking technique. That is, knowledge about real pictures can be used to define enhancement routines which use local features to effectively requantize an image to more levels than were present in the initial image. The number of points having a given gray level will thus become relatively small (when compared to the histogram values of the original picture) so that the histogram equalization transformation will provide more nearly flat histograms. More importantly, preprocessing followed by histogram equalization may invert the gray-level relationship between two points, or may split up constant gray-level sets into subsets which may not be assigned to consecutive gray levels. Depending upon the type of preprocessing used, it is possible to enhance the visual clarity of the image beyond the simple contrast adjustments of Figs. 1 and 2.

In the following paragraphs, we describe a number of inexpensive enhancement routines. Our aim is to provide a selection of programs which can be run automatically prior to displaying a picture. We emphasize programs that do not require many parameter adjustments, but enhance the visual perceptibility of important features in a broad class of pictures. Furthermore, we require that the processing times of these enhancement routines be extremely short. This precludes the use of Fourier techniques, of large-scale feature detectors requiring more than very local computations, and of other higher-level constructs to guide the processing.

The first type of preprocessing enhancement operation we will consider is automatic photometric correction. Although histogram equalization enhances the contrast both in fine detail and in gradual intensity gradients, usually our main interest is limited to fine structures such as edges, lines, and local texture elements. The enhancement of gradual intensity gradients, caused by illumination peculiarities or by a nonuniform photometric response in the imaging system, limits the amount of enhancement possible in small neighborhoods where similar gradients

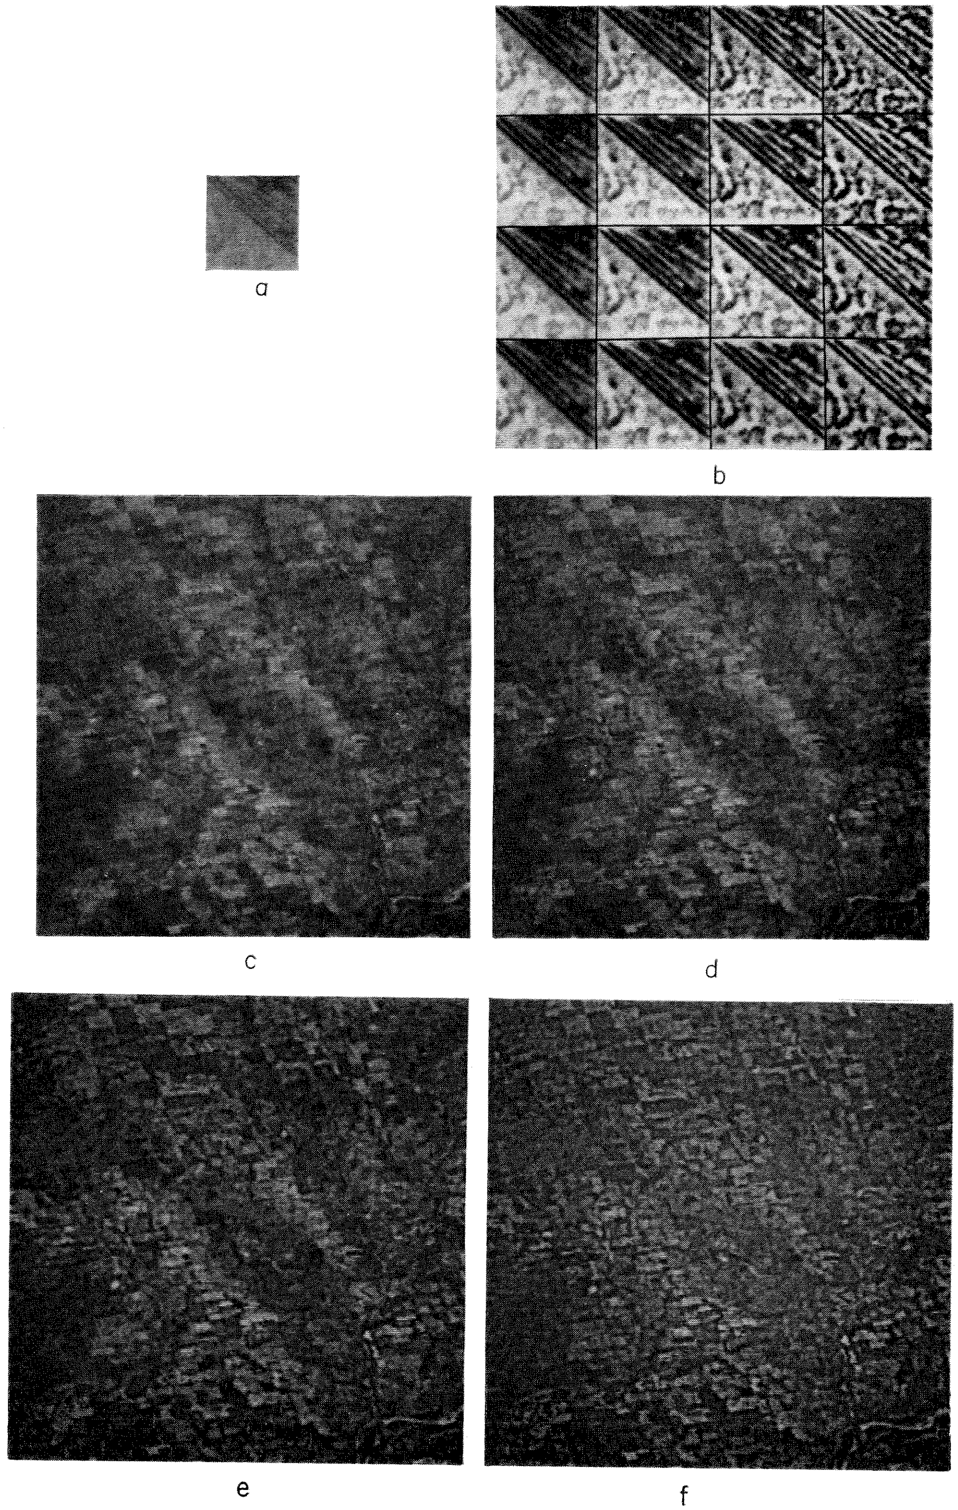
a
b
c
d
e
f

are not present. For example, if a shadow across a picture causes one region of the picture to be much darker than other regions, histogram equalization of the shadow region alone will bring out far more detail than equalization of the entire picture. However, by using some simple neighborhood properties such as local average gray level, or by analyzing the local histogram, a preprocessor can be designed to recognize and correct for gradual gradients or large shadow regions. Afterwards, the final histogram equalization transformation can be applied to the entire picture. From another viewpoint, we are seeking an inexpensive method to enhance high-frequency information, and attenuate low spatial frequencies.

We investigate two methods for automatic photometric correction. The first method, local mean correction, attempts to adjust the local average gray level at each point toward the global mean. Let μ_0 denote the global mean, and suppose that the average gray level in an m-by-m neighborhood of a pixel is μ. We lighten the pixel's gray level g if μ is darker than the overall average μ_0, and darken g if μ is too light. Specifically, we give the pixel a new gray level $g - \alpha(\mu - \mu_0)$, where α is a constant $(0 \leq \alpha \leq 1)$. The two parameters for this scheme are the neighborhood size m, and the correction coefficient α. Figure 3 shows some pictures which were enhanced in this manner for various values of m and α. As expected, higher values of α increase the enhancement of lines and edges as values at points in the interior of large regions are pushed closer toward the mean gray level. When the neighborhood size as specified by m is smaller than the size of the predominant regions in the picture, such as the fields in the second example of Fig. 3, then high values of α destroy the interior gray levels while enhancing edges bordering the predominant regions, as in Fourier dominant high-pass filtering.

A second method, called local histogram modification, transforms each point according to the histogram equalization transformation defined by the histogram of a square neighborhood of the point [2]. As with the previous scheme, dark and shadow regions tend to be lightened by this process. Local histogram modification retains the advantage of histogram equalization on small regions, but essentially photomosaics these regions into a single large region by computing a separate transformation for each point. Although the resulting pictures are likely to be quite contrasty, their histograms are not necessarily flat, and so a small amount of additional enhancement can be achieved by following the local histogram modification by global histogram equalization. Figure 4 shows some pictures enhanced by local histogram modification for various neighborhood sizes. Unlike local mean correction, local histogram modification provides no parameters to adjust the degree of photometric correction. Accordingly, the results of Fig. 4 correspond to complete photometric correction, as in the case of $\alpha = 1$ in Fig. 3. Our examples assume that the local transformation should be based on a histogram equalization strategy. However, certain classes of pictures might be better enhanced if the local histogram transformation at each point were defined to yield

FIG. 3. Examples of local mean correction: (a) Original terrain sample. (b) Enhancements of (a) using local mean correction for various neighborhood sizes and correction factors. Vertical axis (top to bottom): size $m = 7, 9, 11, 13$. Horizontal axis (left to right): factor $\alpha = 0.25, 0.5, 0.75, 1$. (c) Original. (d) Enhanced version of (c), $m = 13$, $\alpha = 0.3$. (e) Enhanced version of (c), $m = 13$, $\alpha = 0.6$. (f) Enhanced version of (c), $m = 13$, $\alpha = 0.9$.

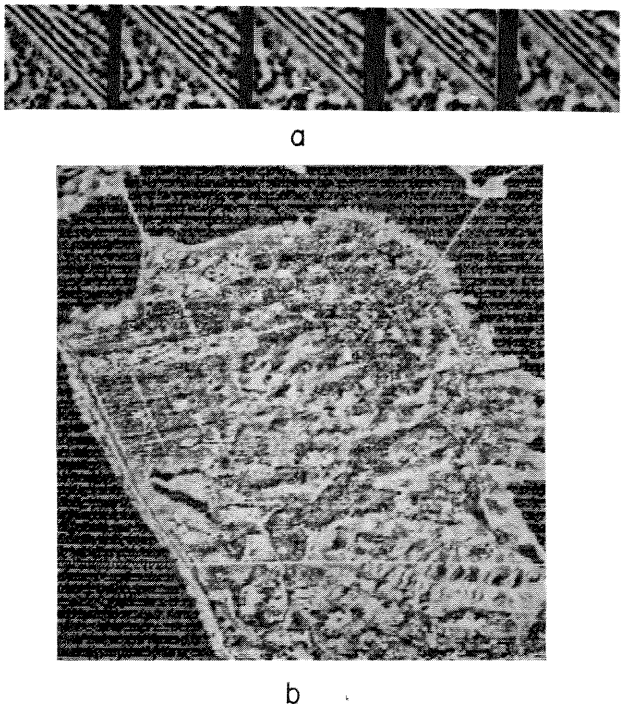

FIG. 4. Examples of local histogram modification: (a) Enhanced version of Fig. 3a using local histogram modification for various neighborhood sizes. (b) Enhanced version of Fig. 1a using local histogram modification, $m = 13$.

something other than a uniform distribution. In this way, one could modulate the degree of local enhancement by adjusting parameters defining the desired resultant local histogram.

Computation of local averages is very cheap, since once a square m-by-m has been summed, the sum for each successive consecutive square can be computed using only m additions and m subtractions. Local histogram modification requires somewhat more work. The local histogram at each point can be determined from the histogram by $2m$ table look-ups, m additions, and m subtractions. However, the transformation which determines the new gray level of the central point depends on the integral of the histogram, which can be computed by a series summation from the bottom, or a series subtraction from the top. That is,

$$\begin{aligned} T(z) &= N[\sum_{k=1}^{z} p(k) + \tfrac{1}{2}p(k+1)] \\ &= N[1 - \sum_{k=N,\,z+1,\,-1} p(k) - \tfrac{1}{2}p(k+1)], \end{aligned}$$

where N is the number of gray levels and $p(k)$ is the local histogram value (see [1, Section 4]). Thus a maximum of about $\frac{1}{2}N$ addition operations are required to determine the new gray level once the local histogram is known. These extra computations are reflected in the processing times required for enhancement using local histogram modification.

A second kind of preprocessing which can precede global histogram equalization (and which may also be combined with automatic photometric correction) involves deblurring by using the Laplacian. It is well known that if a small multiple of the digital Laplacian is subtracted at each point of a picture, the result will tend to enhance edges, lines, and spots by creating artificial Mach bands. The Laplacian picture may be computed at each point according to the formula

$$l_{ij} = \sum_{i'=i-1}^{i+1} \sum_{j'=j-1}^{j+1} g_{i'j'} - 9 \cdot g_{ij}.$$

The complete preprocessing step requires the computation of the picture $g_{ji} - \beta \cdot l_{ij}$, which is the same as $(1 + 9\beta)g_{ij} - \beta s_{ij}$, where s_{ij} is the sum of the pixels in the 3-by-3 neighborhood centered at (i, j). Figure 5 shows pictures enhanced in this fashion for various values of β followed by the standard histogram equalization algorithm. Laplacian deblurring works best when the edges and lines which are the object of enhancement delineate well-defined regions corresponding to distinct, broad features in the original picture. When the lines and edges are themselves the prominent features, deblurring does not work nearly as well (Fig. 5d).

One major difficulty with Laplacian deblurring techniques is that noise points are enhanced even more than edges. If we had a way of identifying noise points, then we could replace these points by neighborhood average values (e.g., by the averages of their eight neighbors), which would accomplish a considerable amount of noise cleaning. Noise points are often characterized by large Laplacian values. However, because the actual gray-level difference between a noise point and its neighbor points is often small compared to the gray-level steps at the strongest edges in a picture, it becomes impossible to distinguish edge points from noise points by Laplacian techniques alone. We could resolve this problem for isolated noise points by examining neighborhood Laplacian values, or by using gradient operators initially to identify edges. However, we may also note that local histogram modification enhances noise extremely well if the neighborhood size is small. If the Laplacian picture is computed *after* applying a local histogram modification routine, then the points whose Laplacian values have greatest magnitudes will correspond to noise points, but not to edges (if the threshold is high enough). In Fig. 6 we show the result of eliminating the noise points identified in this manner.

It is a simple matter to combine a number of preprocessing steps to enhance a picture. If no noise reduction is necessary, one can begin with local mean correction and Laplacian deblurring before applying the final histogram equalization routine. If noise cleaning is to be included, one should begin with local histogram modification, followed by computation of the Laplacian, and analysis of the Laplacian values to determine noise points. Values at noise points should be replaced by local averages, whereas values at edge points (having intermediate

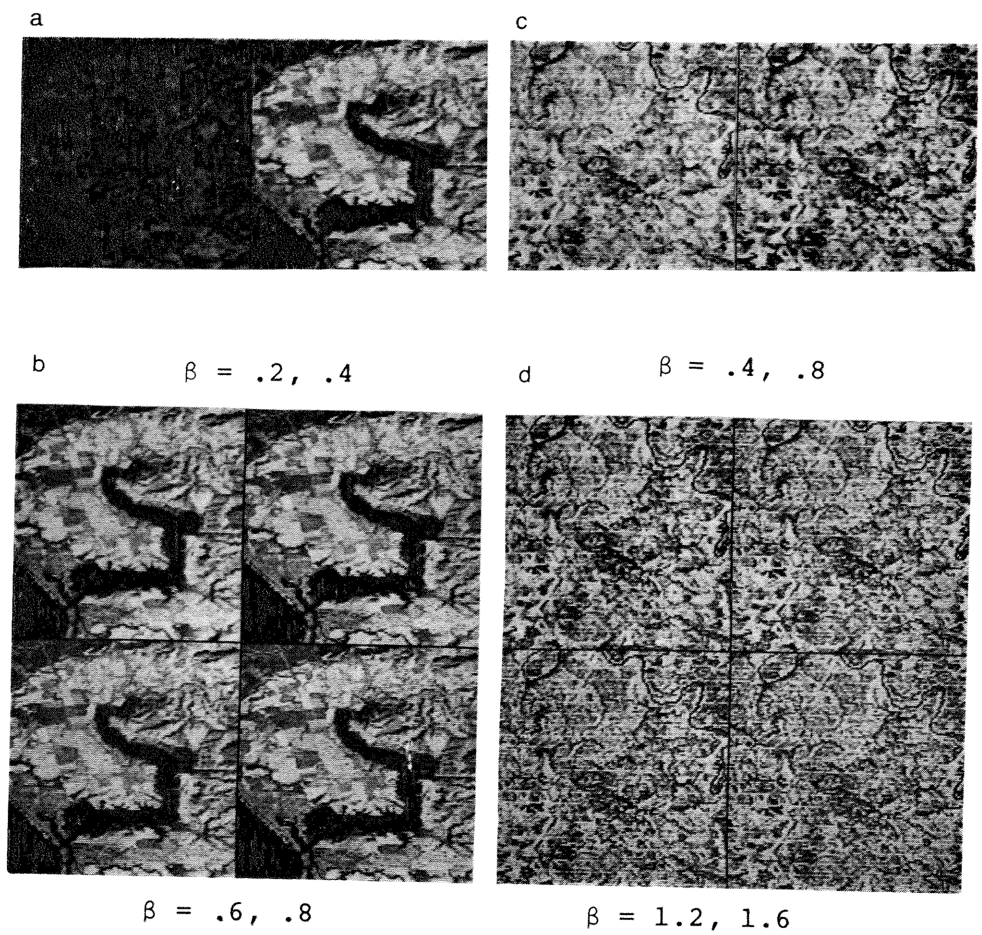

FIG. 5. Examples of Laplacian deblurring: (a) Original and histogram equalized versions. (b) Enhanced versions of (a) using Laplacian deblurring for various correction factors β. (c) Original and histogram equalized versions. (d) Enhanced versions of (c) using Laplacian deblurring for various correction factors β.

Laplacian magnitudes) should be adjusted according to the deblurring formula. All other points can either be left unadjusted, or replaced by values based on a blurred version of the image (in order to smooth regions and soften false quantization edges).

In Figs. 1 through 6 we have assumed that the most desirable histogram output is a flat distribution. This assumption is motivated by the knowledge that a uniform distribution makes best use of the information transmission capabilities of the available gray levels. However, there is no reason to expect the flat histogram to yield a visually optimal version of any image. Indeed, for a picture whose initial histogram is bimodal, with the crucial information contained in the points represented by the smaller peak, histogram flattening will often destroy visual perceptibility of the image. On the other hand, if histogram flattening improves the appearance of an image because gray levels are moved farther apart, we might expect histograms with peaks at the ends and a dip in the middle to yield even better enhancement results than can be obtained for pictures having unimodal histograms. Figure 7 shows a single image transformed according to a number of different modification strategies. Versions a and b have histograms obtained from a truncated portion of a normal distribution, and versions c and d have catenary histograms obtained from different sections of a hyperbolic cosine function.

Because the enhancement effect of a histogram modification transformation depends upon the initial image, some have suggested an interactive approach to gray-level adjustment [3]. In [1], we suggested that the optimal output histogram might be a solution to a differential equation which depends upon the

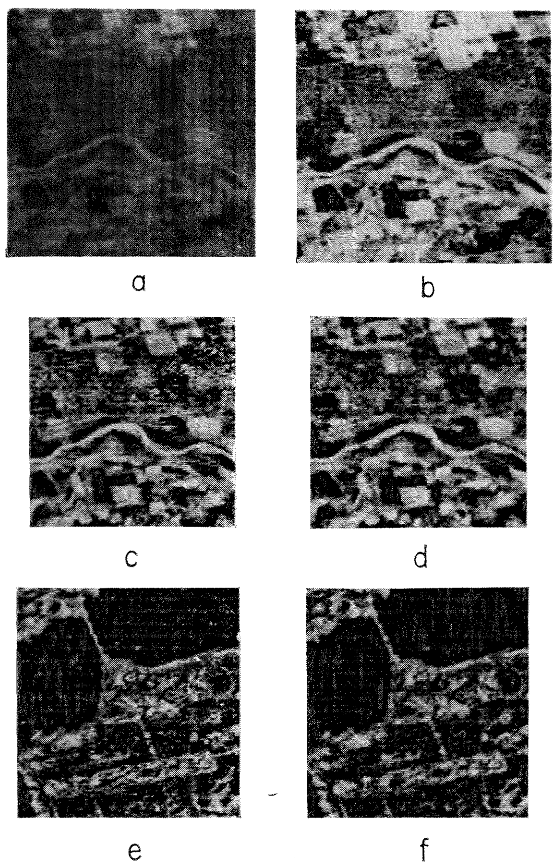

FIG. 6. Noise cleaning by Laplacian thresholding: (a) Original. (b) Histogram equalized version of (a). (c) Enhancement of (a) using local histogram modification. (d) Noise cleaned version of (c) using Laplacian thresholding. (e) Enhancement of a portion of Fig. 1a using local histogram modification. (f) Noise cleaned version of (e) using Laplacian thresholding.

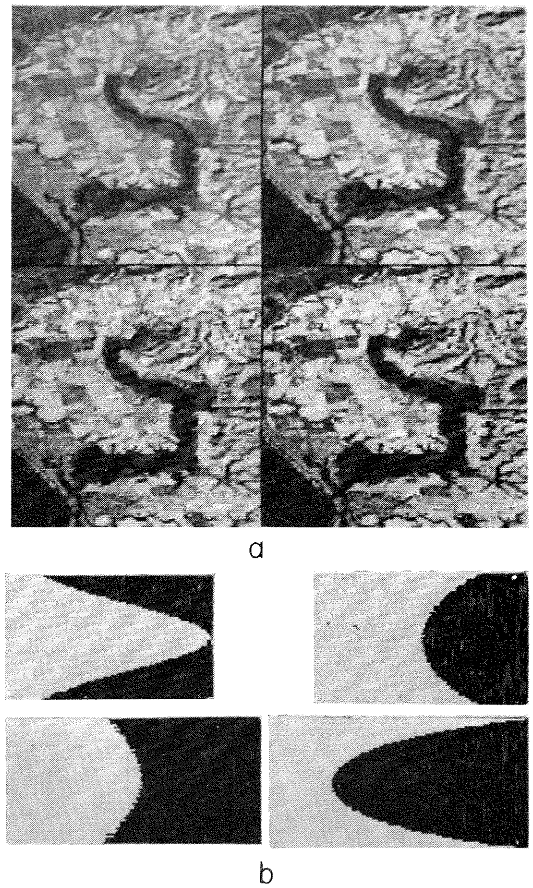

FIG. 7. Histogram modification using normal and catenary distributions: (a) Enhancements. (b) Respective histograms.

initial histogram. Nonetheless, for images which have poor contrast and a large amount of hidden detail (such as LANDSAT images), it is not unreasonable to ask for an overall best display histogram.

In another recent report [4], it has been argued that the transformation which flattens the distribution of the outputs of the retinal receptors will provide the best visual enhancement. Assuming a logarithmic response of the human visual system, the resulting transformation will transform the histogram into a section of a scaled version of the rectangular hyperbola, in which the number of points having a given intensity falls off inversely with the intensity level. This method seems to work well for satellite pictures with inherently poor contrast, due to their relatively small amount of detail. The rate of dropoff of the histogram, that is, the intensities represented by the quantization levels, and the gray level of the darkest (most populated) level, must be adjusted according to the amount of detail in the picture. Once again, an interactive approach is indicated.

REFERENCES

1. R. A. Hummel, Histogram modification techniques, *Computer Graphics and Image Processing* **4**, 1975, 209–224.
2. D. J. Ketcham, R. W. Lowe, and J. W. Weber, *Image Enhancement Techniques for Cockpit Displays*, TR-P74-530R, Display Systems Laboratory, Hughes Aircraft Co., Culver City, Calif., 1974.
3. R. C. Gonzalez and B. A. Fites, Gray-level transformations for interactive image enhancement, presented at The Second Conference on Remotely Manned Systems, June 1975.
4. W. Frei, *Psycho-visual Image Enhancement*, TR-560, Image Processing Institute, University of Southern California, Los Angeles, Calif., 1975.