J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*

IMAGE PROCESSING CONTRAST ENHANCEMENT

Contrast enhancement is a significant component of machine vision applications. The principal objective of contrast enhancement is to process an image so that the enhanced image is more suitable for further processing or viewing. For example, if an image with a large dynamic range is recorded on a medium with a small dynamic range, such as film or paper, the contrast and therefore the details of the image are reduced, particularly in the very bright and dark regions. Increasing the local contrast and reducing the overall dynamic range can significantly enhance the quality of the image. Figure 1 shows a low-contrast image enhanced significantly by a standard contrast-enhancement technique.

Contrast-enhancement methods are applied either during the image-acquisition process or after the digitized image is obtained. Following the contrast-enhancement step, most applications have image-processing, segmentation or feature-extraction steps. In this article, three types of contrast enhancement methods are discussed. The first set of methods is applied during the image-acquisition process. Here the electronic switches and modes of the camera will be studied, and automatic and manual methods of changing contrast as these switches and modes are changed will be discussed. These methods are useful for cameras that have the necessary electronic features. The second set of methods is based on processing the pixels of a digitized image to enhance its contrast. These methods are useful for a wide variety of cameras and are done after an image is acquired by a frame grabber. The third set of methods is based on processing the analog signal from the camera before it is digitized. These methods are fast, and use both analog and digital domain techniques. These methods are commonly referred to as the *hybrid contrast enhancement* methods.

Contrast Enhancement with Camera Eletronics

The most widely used cameras in machine vision applications are solid-state video cameras consisting of *charged-coupled devices* (*CCD*). Initially developed in AT&T Bell Laboratories in the 1970s, these cameras are now inexpensive, lightweight, and compatible with a wide variety of frame grabbers and video transfer protocols. CCD cameras come in both monochrome and color varieties. In addition to CCD cameras, there are also *charge-injection device* (*CID*) cameras. These were initially developed by General Electric for aerospace applications, and have wider dynamic range and greater resistance to blooming than common CCD cameras. However, CID cameras can be less compatible with frame grabbers and monitors used in most applications. In this article, only the features of CCD cameras will be considered.

Contrast Enhancement in the Digital Domain

Contrast enhancement in the digital domain is most common for machine vision applications (1,2,3,4,5). The methods consist of spatial domain techniques based on gray-level mappings, or frequency domain methods based on modifying the Fourier transform of the image. Contrast enhancement based on combinations of

Fig. 1. Image enhanced by a standard contrast enhancement method.

methods from these two categories is not uncommon. Common methods are: (1) gray scale transformation, which changes the gray scale in a uniform way throughout some region of the image; (2) histogram modification, which modifies the histogram of an image in a specified manner; (3) high-pass filtering and unsharp masking, which enhances some frequency components in the image; (4) homomorphic processing, which multiplies the image to an image formation model; and (5) adaptive modification of local contrast.

Contrast Enhancement in Digital and Analog (Hybrid) Domains

Contrast enhancement in the analog or hybrid domains are common in many commercial image processors (6,7,8). Enhancement is achieved by controlling the low and high offsets of the analog-to-digital converter (*ADC*) in the frame grabber. However, in most instances, the methods are manual, and depend on the user to provide the offsets. Moreover, the manual methods are designed to use the full dynamic range of the intensities in the image. Some automatic methods exist (7,8) in commercial applications, but are inaccurate due to the use of simple ADC models. More sophisticated models of the ADC have recently been proposed (9,10,11) for automatic contrast enhancement. These methods will be discussed and their use in real-time applications will be shown.

Contrast Enhancement with CCD Cameras

The CCD cameras are solid-state devices composed of discrete sensors arranged in a two-dimensional array. Most new CCD cameras are equipped with a number of features that allow automatic contrast adjustment. Among these features are: (1) electronic shutter speed control, (2) switches for conveniently adjusting the gamma mode, (3) scanning and transfer modes (frame or field), (4) automatic gain control (*AGC*) or manual gain control (*MGC*) mode and the camera amplifier gain, and (5) time delay integration (*TDI*).

Descriptions of Camera Modes. The gamma mode determines the relationship between scene brightness and the intensity of the video signal. When gamma mode is 1.0, the camera provides a linear relationship between scene brightness and the image intensities. The linear gamma mode is most commonly used in machine vision and commercial applications. The video-display gamma mode results in a video signal which is proportional to the scene brightness raised to the power of 0.45. This mode is approximately equal to the square

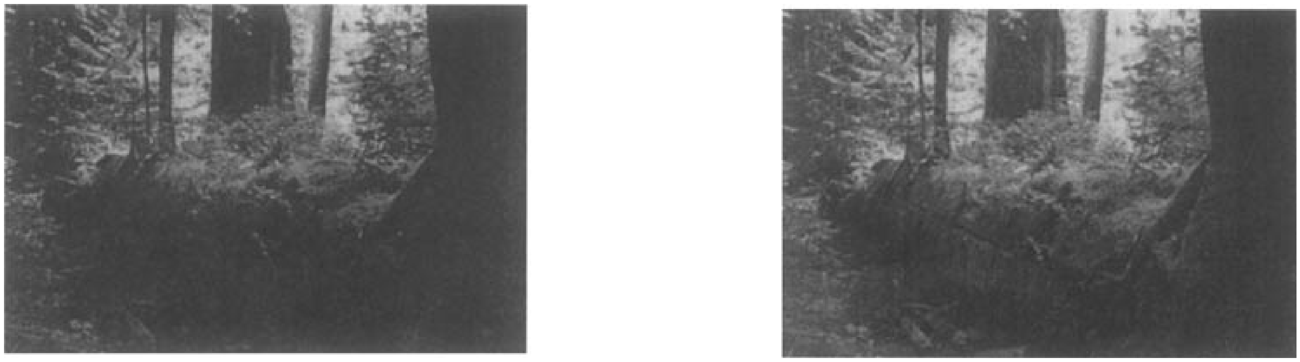

Fig. 2. Gamma correction with gamma = 0.45.

root of the scene brightness and is often referred to as the square-root mode. This mode is commonly used in video surveillance applications where the camera video is directly fed into a CRT-type video monitor, which, in turn, displays the video signal with a 2.2 power factor. Figure 2 below shows an image of a forest that is enhanced by using a gamma of 0.45. The enhanced image on the right shows considerably more details on the forest floor.

The camera gain control modifies the video signal based on the scene brightness and compensates for variations in the scene illumination and scene reflectance. The automatic gain control (*AGC*) allows automatic amplification of the video signal according to the scene reflectance, where the signal is amplified when the scene becomes darker and is reduced when the scene becomes brighter. The manual gain control (*MGC*) changes the gain manually and turns off the automatic signal conditioning feature.

CCD black-and-white cameras produce RS-170 video signals. There are two modes of scanning and transferring the horizontal pulses that make up the frames of the video signal. The frame mode is the standard interlace mode of horizontal line transfer. The field mode is the noninterlaced mode, and combines pairs of adjacent lines and then transfer them. Images in the field mode usually have a much more grainy appearance due to the lesser number of lines per image than the frame mode.

The field mode, however, is used with applications involving electronic shuttering for capturing fast-moving targets. In electronic shuttering, only a portion of the charge that has been accumulated over one field time is scanned. Usual time integrations are for 1 ms or less. Due to such short integration times, in some applications, a bright illumination source, such as a strobe light or a bright dc light is needed. Since in most CCD cameras there is a time delay between the capture of the two (even and odd) fields, the frame mode cannot be used to acquire fast-moving targets. Some new cameras are now available that allow one to use the frame mode for fast-moving targets.

Discussion. The methods in this category are convenient to use when the camera is equipped with these electronic features. They are fast and operate at the acquisition speed of the camera, which is usually 16.7 ms for a frame capture and 33 ms for a field capture. The usual problem of these methods is that they are not general enough for a variety of images, and depend on the availability of these electronic features in the camera hardware. This result is illustrated in Fig. 3, which shows the image of Fig. 1 enhanced by gamma correction with gamma = 0.45. Compare this result with the linear contrast enhancement result (see next section) in Fig. 1. It is clear that gamma correction does not yield the same performance as the digital domain technique.

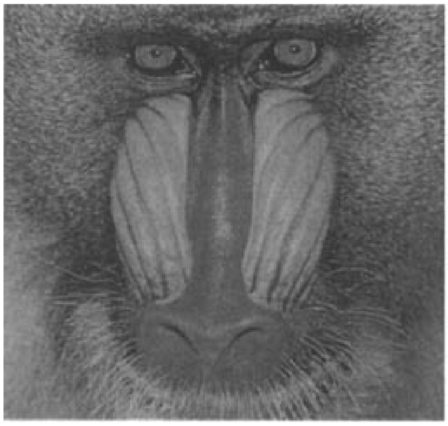

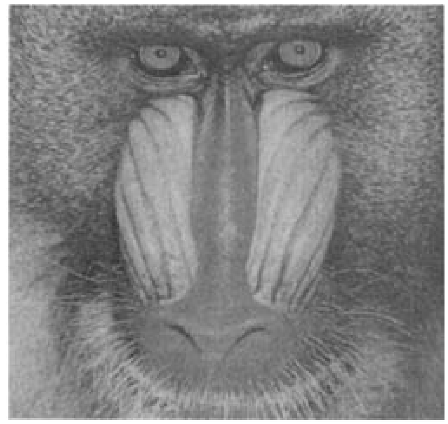

Fig. 3. Problems with gamma correction.

Contrast Enhancement in the Digital Domain

These techniques of contrast enhancement are based upon processing the digital image in order to enhance its contrast. The common methods of digital domain contrast enhancement are listed below.

Gray Scale Transformations. One commonly utilized technique of gray scale transformation is *linear contrast enhancement* (3,12). The objective is to utilize the full dynamic range of the output display medium to reveal the intensity variation present within an image. The technique is based on transforming each pixel intensity level in the input image into a new value in the output image by a linear transformation. This method is particularly well suited to images with Gaussian or near-Gaussian histograms, where all the intensity values fall within a narrow subset of the total intensity range. Figure 4 shows an 8-bit image which has gray values in the range [40,100]. These gray levels are linearly mapped to the range [0,255] to obtain an enhanced image in which the persons appear clearer. Let $I_{\min}$ and $I_{\max}$ be the lowest and highest gray values respectively in an image I. Then the transformation used to obtain the linear contrast enhanced image J is

$$J = \frac{255}{I_{\max} - I_{\min}}(I - I_{\min})$$

A second method of gray scale transformation automates the method of contrast enhancement. The *automated ends-in search* (12) technique allows the user to specify that a certain percentage of the pixels in the output image must be saturated either full black or full white. The method utilizes the cumulative distribution function to identify the intensity values at which the appropriate percentages will saturate when a linear contrast enhancement is performed.

Histogram Modification. The histogram of an image with gray values in the range [0, $L-1$] is a discrete function $p(r_k) = n_k/n$, where r_k is the kth gray level, n_k is the number of pixels in the image with gray level r_k, n is the total number of pixels in the image, and $k = 0, 1, 2, \ldots, L-1$. Loosely speaking, $p(r_k)$ gives an estimate of the probability of occurrence of gray level r_k. The histogram represents a global description of the appearance of an image, and manipulating the histogram results in contrast enhancement.

Methods of *histogram modification* are many, and they commonly belong to: (1) global methods, (2) adaptive methods, and (3) methods using structure. The global or stationary enhancement techniques involve gray level transformations based solely on the intensity of each pixel. Methods include histogram equalization

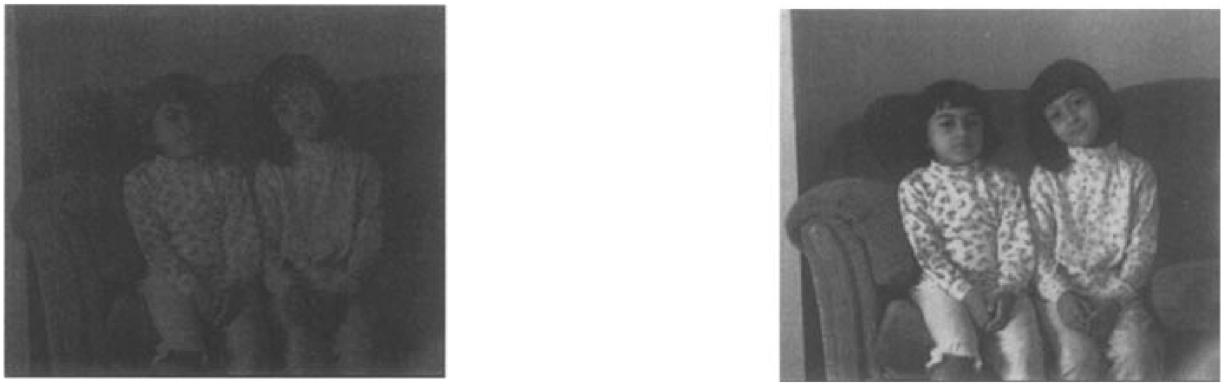

Fig. 4. Linear contrast enhancement.

(13), histogram hyperbolization (14), and histogram specification (2). Adaptive enhancement methods calculate the new intensity value for a pixel from its original value and some local image properties. Methods include unsharp masking (15), statistical difference filter (16), local histogram modification (17,18), pixel contrast measure (19), adaptive histogram equalization (*AHE*) (20), and contrast-limited AHE (21). The final category includes methods that use local image structures to change the enhancement calculation itself, or change the contextual region over which the calculations are done (22,23,24).

Histogram Equalization. In histogram equalization, the gray values are spread in the image "equally" over the entire range of gray levels so that the resultant histogram looks more or less like a uniform distribution. Following the notations used before, let L be the number of gray values in the image, $p(r_k)$ be the probability of the kth gray value, and a plot of $p(r_k)$ versus r_k be the histogram of the image. In order to obtain a uniform histogram of the image, the following transformation of each gray value r_k into s_k is used:

$$s_k = T(r_k) = \sum_{j=0}^{k} \frac{n_j}{n} = \sum_{j=0}^{k} p(r_j) \qquad \text{for } k = 0, 1, \ldots, L-1 \quad (1)$$

In order to illustrate the usefulness of histogram equalization, consider Fig. 5 which shows a 8-bit image of a dirt road. The bright original image has poor dynamic range and does not display the features on the road. The histogram shows that most of the pixels are concentrated around the gray value of 100. The histogram equalized image on the right displays considerably more details in the image, and the histogram is distributed evenly over all gray levels.

Histogram Specification. Since histogram equalization only allows one to obtain an approximation of a uniform histogram, with histogram specification one can specify a particular histogram shape. This method makes it possible to highlight certain gray level ranges in an image as desirable. Sometimes the ability to specify particular histogram shapes capable of highlighting certain gray level ranges in an image is desirable.

Consider an image with L gray levels, where $p_r(r_k)$ is the probability of the kth gray level, n_k is the number of times this level appears in the image, and n is the total number of pixels in the image. A plot of $p_r(r_k)$ versus r_k is the histogram of the original image. Let $p_d(d_k)$ be the desired histogram to obtain in the image. For this, histogram equalization is first utilized on the original image by applying Eq. (1); that is,

$$s_k = T(r_k) = \sum_{j=0}^{k} \frac{n_j}{n} = \sum_{j=0}^{k} p_r(r_j) \qquad \text{for } k = 0, 1, \ldots, L-1 \quad (2)$$

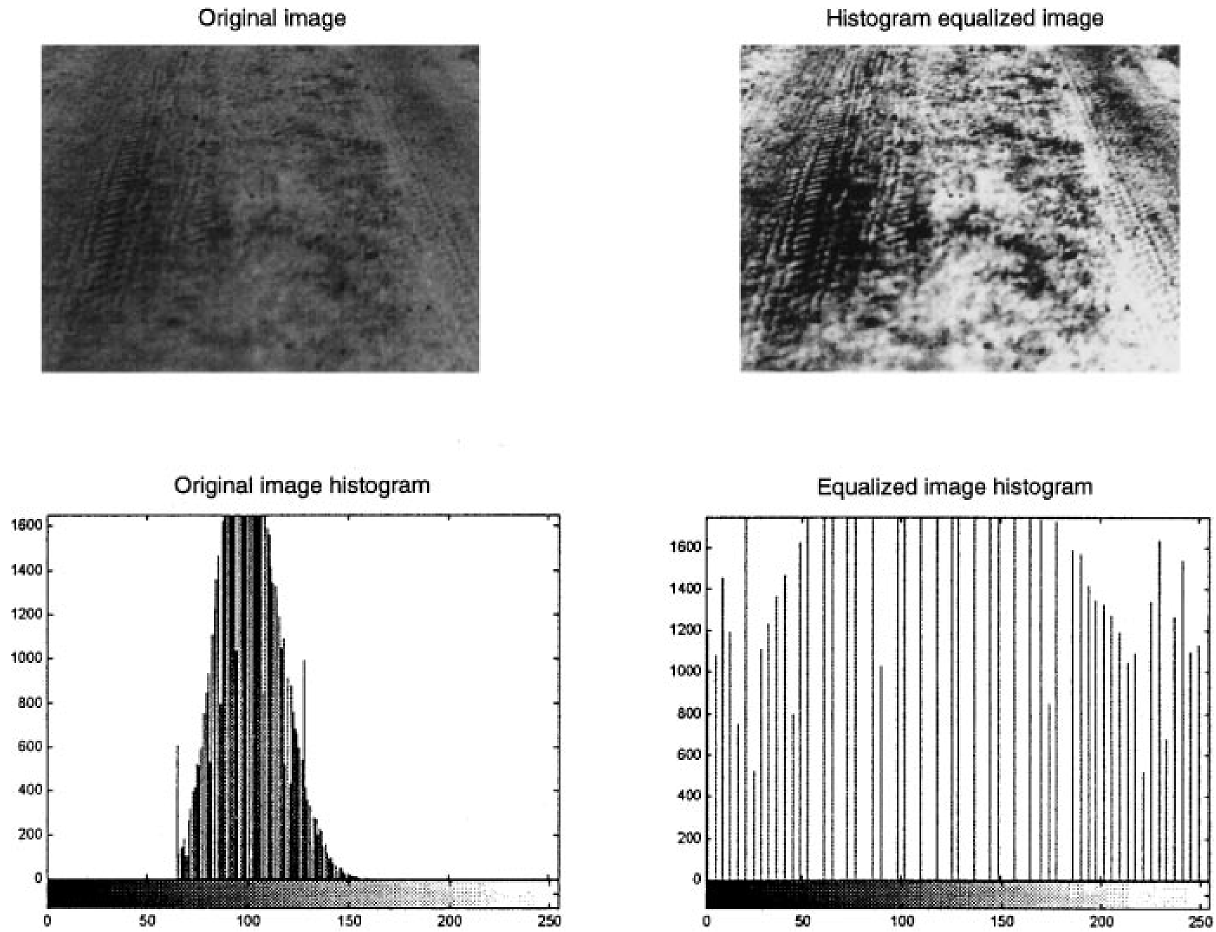

Fig. 5. Histogram equalization.

Next, the cumulative distribution function (CDF) $G(d_k)$ of the desired histogram $p_d(d_k)$ is obtained, as follows:

$$G(d_k) = \sum_{j=0}^{k} p_d(d_j) \qquad \text{for } k = 0, 1, \ldots, L-1 \tag{3}$$

An image with the desired histogram is obtained by applying the inverse transform G^{-1} to the histogram equalized image. Assuming that G^{-1} is single-valued, the procedure can be summarized as follows:

(1) Equalize the histogram of the original image to obtain gray levels s_k using Eq. (2).
(2) Specify the desired density function $p_d(d_k)$ and obtain the transformation function $G(d_k)$ using Eq. (3).
(3) Apply the inverse transformation function $d_k = G^{-1}(s_k)$ to the gray levels obtained in Step 1.

This procedure yields a processed version of the original image, with the new gray levels characterized by the specified density $p_d(d_k)$.

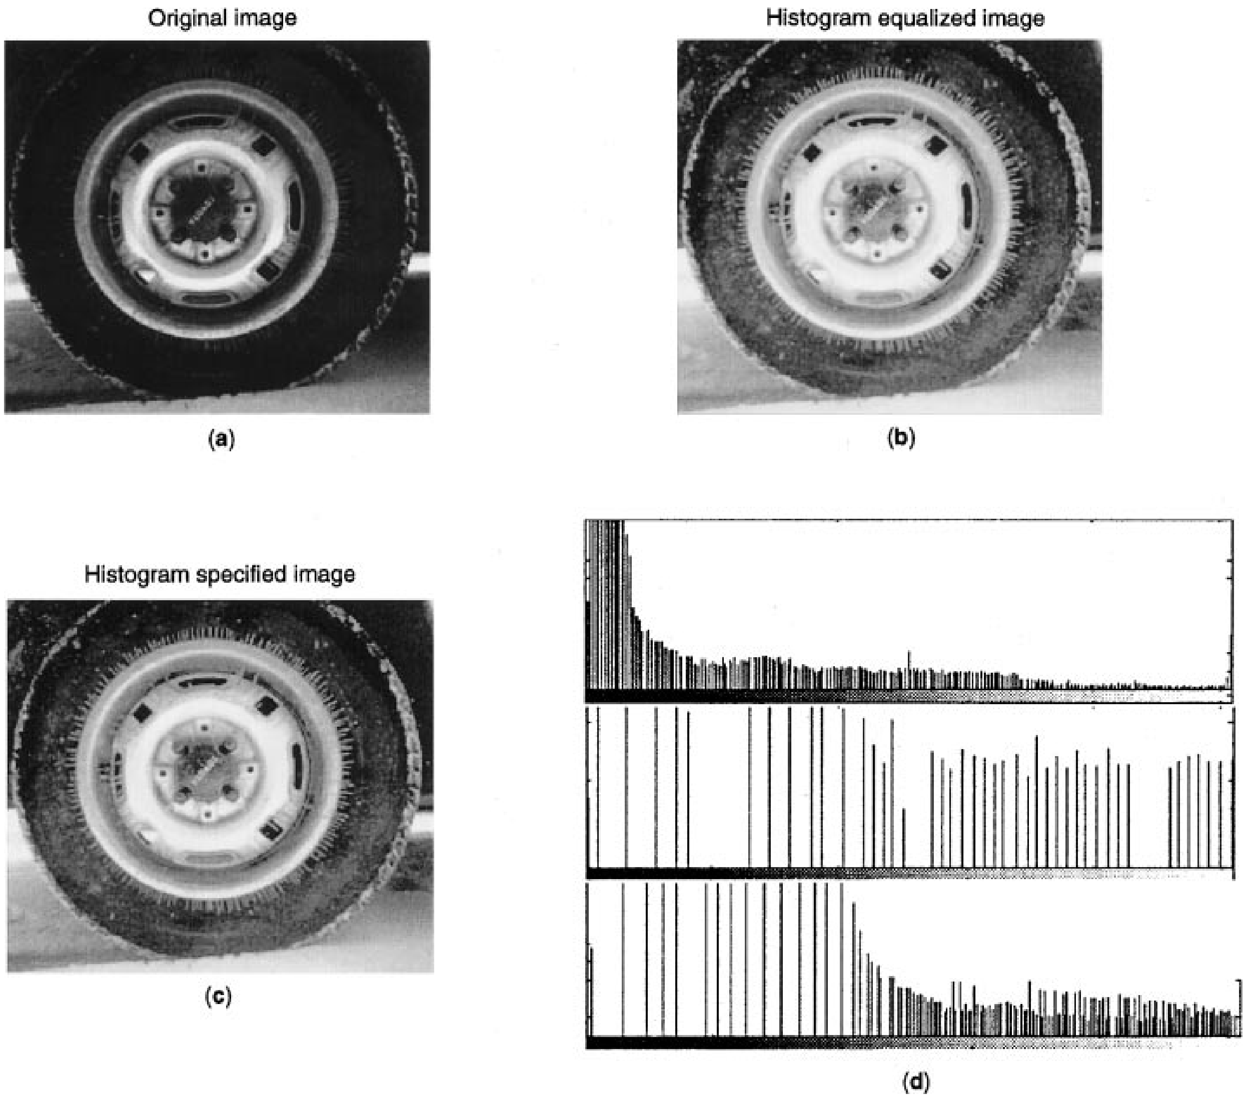

Fig. 6. Histogram specification.

In practice, the inverse transform from s_k to d_k is often not single-valued. This situation arises when there are unfilled levels in the specified histogram. Two solutions to this problem are as follows: The first is to specify a particular probability density function (such as the Gaussian density) and then form a histogram by digitizing the given function. The second approach consists of specifying a histogram shape by means of a graphic device (e.g., an interactive screen), whose output is fed into the processor executing the histogram specification algorithm.

Consider the image of a tire in Fig. 6(a) whose histogram is equalized in Fig. 6(b), and then specified in Fig. 6(c). The equalized image has an enhanced contrast but the rim and road regions are quite bright. The histogram specification in Fig. 6(c) uses the Gaussian function $e^{-x^2/2}\sqrt{2\pi}$ and shows a better contrast at the

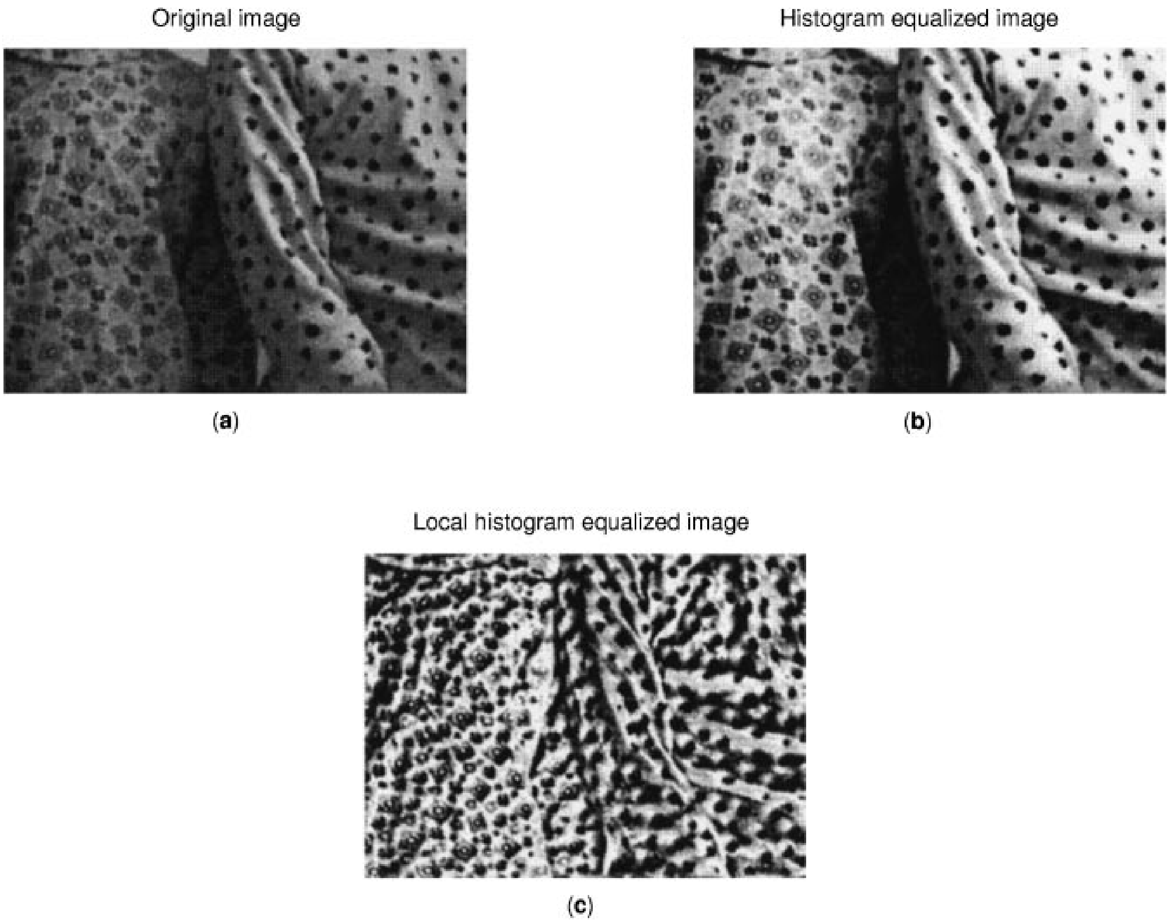

Fig. 7. Local histogram equalization.

rim and the road regions, giving the image a much more balanced appearance. The original, equalized, and specified histograms are also shown [Fig. 6(d)].

Local Histogram Modification. Local histogram modification includes several techniques under the topics of adaptive histogram modification and methods using structure. The two histogram modification methods discussed before are global methods, that is, the histogram is modified based on the gray level distribution over the entire image. Although these methods are useful for overall contrast enhancement, it is often necessary to enhance details over small areas of the image. The number of pixels in these areas may have negligible effects on the histogram of the entire image. Hence, the local methods devise transformation functions based on the gray level distribution and other properties in the neighborhood of every pixel in the image.

One method of local histogram enhancement consists of defining a rectangular neighborhood and moving the center of this neighborhood from pixel to pixel. At each location, the histogram of the points in the neighborhood is computed and a global histogram modification method (e.g., histogram equalization) is applied. Since only one new row or column of the neighborhood changes during a pixel-to-pixel translation of the neighborhood, the histogram can be updated by using an adaptive algorithm. Variations of the adaptive method, which limits the contrast enhancement that can be obtained at every pixel position, are also available.

Figure 7(a) below shows the image of texture which is enhanced by adaptive local histogram equalization method. A 15×15 neighborhood is moved from pixel to pixel and the enhanced image is shown in Fig. 7(c). Figure 7(b) shows the result of global histogram equalization which illustrates considerable enhancement of the contrast. Note, however, that the local histogram equalization method reveals finer details in the texture that are not apparent in Fig. 7(b). The grainy appearance in Fig. 7(c) is due to the fact that the 15×15 neighborhood was placed at every fourth row and column pixel instead of every pixel, for purposes of speed.

High-Pass Filtering and Unsharp Masking. High-pass filtering emphasizes the high-frequency components of the image while reducing the low-frequency components. Because edges or fine details of an image are the primary high-pass components of an image, high-pass filtering often increases the local contrast and sharpens the image.

Unsharp masking, a well-known photographic technique, is closely related to high-pass filtering. It first blurs (unsharpens) the original image, and then a fraction of the unsharp image is subtracted from (or masked from) the original. In order to demonstrate the relation between unsharp masking and high-pass filtering, consider a digital image I, which is unsharpened to obtain an image $LP(I)$, where $LP(\cdot)$ is a low-pass operation that causes the blurring or unsharpening of the original image I. Let c (≤ 1) be a positive scalar, and let U be the processed image. Then, U can be given by

$$U = I - cLP(I) \tag{4}$$

A low-pass-filtered image $LP(I)$ can be computed as the difference between the original image I and a high-pass-filtered version $HP(I)$ as $LP(I) = I - HP(I)$. Substituting this in the above equation, one obtains

$$U = (1 - c)I + cHP(I) \tag{5}$$

When $c = 1$, the standard high-pass result is obtained. When $c < 1$, part of the original image I is added back to the high-pass result $HP(I)$, which restores partially the low-frequency components lost in the high-pass filtering operation. The result is that the image processed by unsharp masking looks more like the original image, with a relative degree of edge enhancement that depends on the value of c. This method is one of the basic tools for contrast enhancement in the printing and publication industry.

Figure 8 shows an image whose contrast is enhanced by unsharp masking. Although unsharp masking enhances the local contrast of the image, it also enhances the background noise which typically has high-frequency components. A comparison of the background regions of Figs. 8(a) and 8(b) shows that the unsharp mask filtered image appears more noisy than the original.

One way to solve the graininess and the accentuation of the background noise caused by unsharp masking is to use an edge-strength mask which is an approximation of the gradient magnitude field. Then use the unsharp masking filter only near the edges. This method is applied to the image in Fig. 8 and the Sobel kernel (4) is used to find the binary edge mask. This method is often called *spatially varying unsharp masking.*

Homomorphic Processing. In some applications it is necessary to reduce the dynamic range of the image and increase the local contrast. One such application is for the printing industry, where an image with a large dynamic range (e.g., a scene on a sunny day) is transferred to a film or paper with a small dynamic range. The homomorphic filtering approach enhances the image contrast by simultaneously reducing the dynamic range and increasing the local contrast before the image is recorded on film or paper. Let $i(x, y)$ represent the illumination model and $r(x, y)$ represent the reflectance model for an image $f(x, y)$. Then, one simple model for image formation is

$$f(x, y) = i(x, y)r(x, y) \tag{6}$$

Original image

(a)

Unsharp masked image

(b)

Fig. 8. Unsharp masking.

In developing a homomorphic system for contrast enhancement, the illumination component $i(x, y)$ is assumed to be the primary contributor to the dynamic range and has low-frequency components, while the reflectance component $r(x, y)$ is the primary contributor to the details and local contrast and has high-frequency components. In order to reduce dynamic range and enhance local contrast, it is necessary to reduce $i(x, y)$ and increase $r(x, y)$. To decouple the two components, take the logarithm of $f(x, y)$ and obtain:

$$\log f(x, y) = \log i(x, y) + \log r(x, y) \qquad (7)$$

By low-pass filtering $\log f(x, y)$, $\log i(x, y)$ is obtained, and by high-pass filtering $\log f(x, y)$, $\log r(x, y)$ is obtained. Then deemphasize $\log i(x, y)$ to reduce dynamic range, and emphasize $\log r(x, y)$ to increase local contrast.

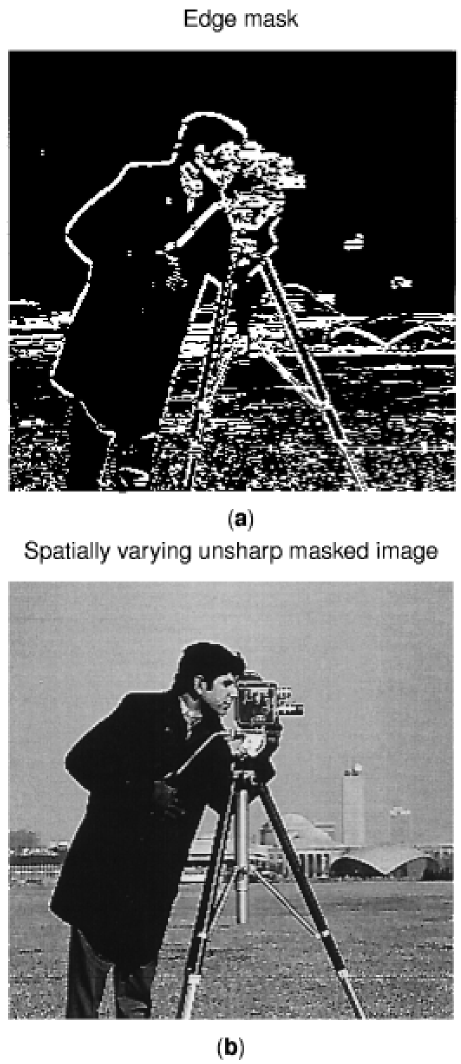

Fig. 9. Spatially varying unsharp masking.

The two components are then combined, and the result is exponentiated to get back the processed image. The algorithm can be written as:

(1) Decouple the illumination and reflectance components of the image as additive terms by taking the logarithm of the image $f(x, y)$, as in Eq. (7).
(2) Obtain the low-pass component $\log i(x, y)$ by low-pass filtering the logarithm image.
(3) Obtain the high-pass component $\log r(x, y)$ by high-pass filtering the logarithm image.
(4) Combine a weighted version of the two components as:

$$\log f'(x, y) = a \log i(x, y) + b \log r(x, y) \tag{8}$$

where $a < 1$ and $b > 1$.
(5) Take the exponent of $\log f(x, y)$ to obtain the final processed image.

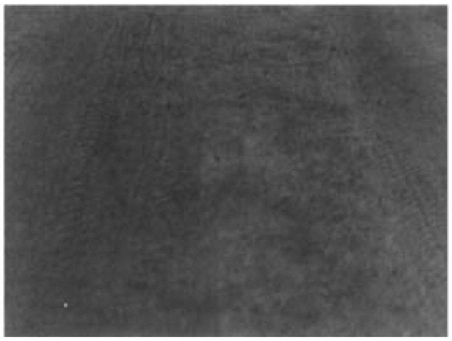


Fig. 10. Homomorphic filtering.

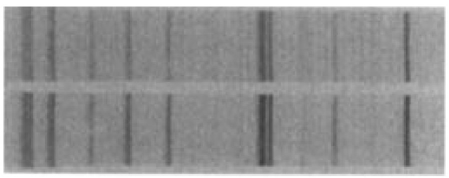
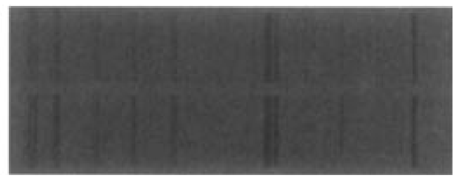

Fig. 11. Adaptive local contrast-modification.

Figure 10 shows an image processed by homomorphic filtering. The processed image reveals several features not visible in the original.

Adaptive Modification of Local Contrast. Some applications require that the local contrast be modified by an adaptive method, as the local contrast in the image varies. One such application is an image of a highly reflecting object with subtle surface features. In such an image, regions of the image have increased local reflectance due to the reflection of light from the shiny surface, and decreased local contrast due to poor signal from the surface features. One approach to contrast enhancement is to increase the local contrast and decrease the local luminance whenever shiny regions are detected.

Let I denote the original image. First regions of high reflectance are obtained by using a low-pass filter on I to obtain image LP. The image HP which denotes the local contrast is obtained by subtracting LP from I. The local contrast is modified by multiplying each pixel $HP(x, y)$ with scalar $a(x, y)$, where $a(x, y)$ is a function of $LP(x, y)$ as:

$$HP_{\mathrm{new}}(x, y) = a[LP(x, y)]HP(x, y)$$

When local contrast needs to be enhanced as in Fig. 11, we choose $a(x, y) > 1$. The modified local contrast image HP_{new} is added to the low-pass image LP to obtain the final processed image.

Discussion. Histogram equalization is powerful in "normalizing" and enhancing the image by spreading widely the histogram peaks. A disadvantage of this method is shown in Fig. 12. This method may also enhance the noise in the image, obscure necessary features, and reveal unnecessary features. Ideally one would like to minimize the spread in width of the individual peaks, while pulling the peak center positions so that the contrast between different regions in the image are enhanced. The histogram specification methods can, however, achieve this result, and have wider applications than contrast enhancement. However, histogram specification requires a lot more computation than histogram equalization.

The adaptive and structure dependent methods are local operations. They are suitable for enhancing subtle details at the expense of principal features which may be lost in the process. A major drawback of

Fig. 12. Problems with histogram equalization.

adaptive methods such as local histogram modification "is that it carries a heavy computational burden", and "practically prohibits the use of these techniques on general purpose computers in online or real-time applications" (24).

Analog-to-Digital (Hybrid) Contrast Enhancement

Next, a hybrid contrast enhancement method using the analog to digital converter (ADC) is discussed, for use in high-speed industrial applications. In the previous section, contrast enhancement methods, which enhance the image in digital domain in a pixel-by-pixel operation, were discussed. In comparison, the hybrid method processes a small region of the digital image to compute the low and high offsets of the ADC. These offsets are then used to acquire a new enhanced image. In particular, the method allows the user to select two regions of the image for enhancement, and select two *desired gray values* for these regions, respectively. The method then automatically generates images of this desired contrast. Applications of the method include efficient on-line contrast enhancement for manual identification of features on parts in manufacturing lines.

Some key features of the hybrid method are:

(1) Efficient and Automatic Enhancement The method is computationally efficient due to the signal processing in the analog domain by the ADC. The method requires two image acquisitions. In the first, ADC offsets are computed from a small probing region of the image. In the second, an enhanced image is acquired by controlling the ADC offsets. Thus, if acquisition is done in 33 ms, the entire enhancement is performed in under 100 ms. Processing time for half-resolution images are further reduced. In comparison, an experimental study with digital domain methods shows that they require much longer processing time in simple imaging systems. Furthermore, the speed of these methods depends on the sizes of the regions that are enhanced. The hybrid method, on the other hand, enhances the entire image, and its speed is independent of the region size. Moreover, the method automates the process of generating images of similar contrast.

(2) Economy of Implementation The hybrid method can be implemented with no custom hardware. A processor with an ADC and a frame buffer is sufficient for its implementation.

(3) Desired Image Generation The method allows the user to specify the desired contrast level. In applications where images of varying contrast are obtained, this method can be used to automatically generate images of a user-specified contrast. An example is discussed in a later section.

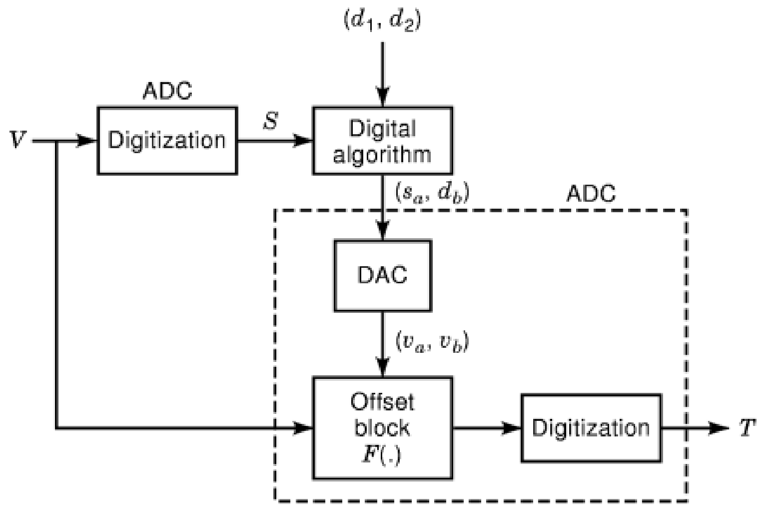

Fig. 13. Hybrid method to obtain enhanced image from analog input.

In summary, the method consists of two steps: (1) The user selects two gray values d_1 and d_2 for two regions of the image that need to be enhanced. *Note that the exact boundaries of the regions are not necessary; just their approximate positions are required.* (2) The method acquires an image and automatically estimates the ADC offsets. (3) It then acquires a new image with these offsets, which is the enhanced image. Note that in the above process, the user selects the approximate positions of the two regions and their desired gray values once for all parts. Thus, once these parameters are determined, the enhancement process is automatic for each part.

Enhancement Model. The enhancement model used to transform the input analog signal to obtain an enhanced gray scale image is described next. Let V be an input analog image signal, and S be a digitized gray scale image obtained by digitizing V. Let $[v_{\min}, v_{\max}]$ be the full dynamic range of input V, and let $[s_{\min}, s_{\max}]$ be the dynamic range of the intensity in S. Let v_a be the low offset and v_b be the high offset of the ADC, that are automatically computed by the hybrid method.

The user selects two gray values d_1 and d_2 for two regions in the image S to be enhanced. Note that the exact boundaries of these regions are not required, just their approximate positions. The algorithm automatically generates digital offsets (s_a, s_b). These offsets are transformed by the digital to analog converter (*DAC*) to analog offsets (v_a, v_b). Note that the DAC is included in almost all ADCs. The ADC then uses (v_a, v_b) to transform input V by an affine function F described below. The transformed input is digitized by the ADC to obtain the enhanced image T. Thus, the ADC performs two functions: (1) an affine transform of V by a function F with offsets (v_a, v_b) (see Fig. 13), and (2) digitization of the transformed input $F(V)$ to obtain enhanced image T.

In modeling the ADC, it can be viewed as a linear transform followed by the quantization process. Ignoring quantization, the linear transform can be expressed as: $s - s_{\min} = c(v - v_{\min})$, where c is a scalar constant, v is the input signal, and s its digitized value. The effect of quantization and nonlinearities are discussed later. The offset block F in Fig. 13 can be described by the following affine transform:

$$F(v) = \left(\frac{v_{\max} - v_{\min}}{v_b - v_a}\right)(v - v_a) + v_{\min} \qquad \text{for} \quad v_a \le v \le v_b \quad (9)$$

Gray value t belonging to enhanced image T is obtained by digitizing of $F(v)$ as shown below:

$$t = g(s - s_a) + s_{\min} \qquad \text{for } s_a \leq s \leq s_b \tag{10}$$

where s is the *initial contrast* of the image, and g is the *gain* described as:

$$g = \frac{s_{\max} - s_{\min}}{s_b - s_a} \tag{11}$$

Quantization, nonidealities, and nonlinearities are commonly found in real ADCs (25,26,27). Conforming to (4) these errors shall be represented by an affine transform: $s - s_{\min} = c(v - v_{\min}) + e$ instead of the ideal linear transform above. Here e is a *nonlinear* error due to the quantization and nonidealities of common ADCs. Considering this affine transform, one obtains (11) the following nonlinear expression for the enhanced image T instead of (2):

$$t = gs + k_1 g s_a + k_2 s_a + k_3 s + k_4 \tag{12}$$

In Eq. (12), (k_1, k_2, k_3, k_4) are unknown real-valued model constants. The ideal values are: $k_1 = -1$, $k_2 = k_3 = 0$. Constant k_4 represents any bias in the system due to an affine ADC. Ideally for an unbiased system, $k_4 = s_{\min}$, which is usually 0. Gain g is as described in Eq. (11).

In formulating Eq. (12) for the estimation of g and s_a, note the following two constraints:

(1) In most practical ADCs, the linearity assumption in Eqs. (10) and (12) is valid for $g < g_{\max}$, where $g_{\max}$ is an upper bound determined experimentally [see (10) for details]. Thus, g must lie within a range $R(g)$ defined as:

$$R(g) : 1 \leq g \leq g_{\max} \tag{13}$$

(2) Similarly, low offset s_a must lie within a range $R(s_a)$ defined as:

$$R(s_a) : s_{\min} \leq s_a \leq s_{\max} - \frac{s_{\max} - s_{\min}}{g} \tag{14}$$

High offset s_b is obtained from:

$$s_b = s_a + \frac{s_{\max} - s_{\min}}{g} \tag{15}$$

Enhancement Method. The enhancement method consists of two steps: (1) an *off-line* estimation of model constants (k_1, k_2, k_3, k_4), and (2) the *on-line* estimation of the ADC offsets (s_a, s_b). The enhanced image T is obtained by acquiring an image with offsets (s_a, s_b).

Off-Line Estimation of Model Constants (k_1, k_2, k_3, k_4). The model constants (k_1, k_2, k_3, k_4) are estimated by the following algorithm in an off-line process:

(1) Choose p offsets (s_{aj}, s_{bj}), $j = 1, \ldots, p$, satisfying Eq. (14), and compute gain g_j from Eq. (11). Here p (>4) is an arbitrary number of offsets (say 10) that are selected by the user, such that the least-squares Eq. (16) below is overdetermined.
(2) Choose two regions of the image that need to be enhanced. For each (s_{aj}, s_{bj}) above, acquire an image, and estimate gray values t_{1j} and t_{2j} by averaging pixels within the two regions, respectively. A small probing window may be used within each region to reduce computation.
(3) For each region i, Eq. (12) can be written as a linear equation below with unknown parameter vector $[k_3 s_i + k_4 \quad k_2 \quad k_1 \quad s_i]$ for $i = 1, 2$:

$$[1 \quad s_{aj} \quad g_j s_{aj} \quad g_j]\begin{bmatrix} k_3 s_i + k_4 \\ k_2 \\ k_1 \\ s_i \end{bmatrix} = [t_{ij}] \qquad \text{for } j = 1, \ldots, p \quad (16)$$

The above equation can be solved by the least-squares method.
(4) Final estimates of k_1 and k_2 are obtained by averaging estimates from the two regions.
(5) From the estimates of s_i and $(k_3 s_i + k_4)$ for $i = 1, 2$, obtained from Eq. (16), compute k_3 and k_4.

On-line Generation of Enhanced Image from Known Constants. Next generate the enhanced image, in an on-line process, from: (1) known model constants (k_1, k_2, k_3, k_4) computed above in an off-line process, and (2) user-specified gray value d_i for each region of the image to be enhanced. From Eq. (12), the following expression for d_i is obtained:

$$d_i = g s_i + k_1 g s_a + k_2 s_a + k_3 s_i + k_4 \qquad \text{for } i = 1, \ldots, m \quad (17)$$

Here m is the number of regions that are enhanced by the algorithm. Before discussing the estimation of (s_a, s_b), analyze the maximum number of regions (m) that can be simultaneously enhanced by this method. The following theorem summarizes the results.

Considering m regions of an image to be enhanced by Eq. (17), the necessary and sufficient condition for which solutions for gain g and low offset s_a exist is:

$$\frac{d_i - d_j}{s_i - s_j} = \frac{d_k - d_r}{s_k - s_r} \qquad \text{for all } i, j, k, r = 1, \ldots, m \quad \text{such that } i \neq j \text{ and } k \neq r \quad (18)$$

Proof:. For any pair of equations indexed by i and j, the solutions to Eq. (17) are:

$$\alpha_{ij} = g + k_3 = \frac{d_i - d_j}{s_i - s_j} \quad \text{and} \quad \beta_{ij} = (k_1 g + k_2)s_a + k_4 = d_i - s_i\left(\frac{d_i - d_j}{s_i - s_j}\right) \quad (19)$$

If Eq. (17) has a solution then $\alpha_{ij} = \alpha_{kr}$ and $\beta_{ij} = \beta_{kr}$. Considering only α_{ij} and α_{kr} one obtains Eq. (18). Next show that if Eq. (18) is true, then Eq. (17) has a solution. In other words, show that $\beta_{ij} = \beta_{kr}$ and $\alpha_{ij} = \alpha_{kr}$ if Eq. (18) is true. Clearly, from Eqs. (18) and (19) $\alpha_{ij} = \alpha_{kr}$. From Eq. (19) $\beta_{ij} - \beta_{kr} = (d_i - d_k) - c(s_i - s_k)$, where $c =$

$(d_i - d_j)/(s_i - s_j) = (d_k - d_r)/(s_k - s_r) = (d_i - d_k)/(s_i - s_k)$. From this equation one may conclude: $\beta_{ij} - \beta_{kr} = 0$, and hence, $\beta_{ij} = \beta_{kr}$.

One can choose up to two independent values of d_i for an image, the rest of which are fixed by (18). One can, therefore, simultaneously enhance up to two regions of an image by this method.

The proof of this corollary is a direct outcome of the theorem. The result is intuitive, because there are only two offsets (s_a, s_b) to modify input V, giving two degrees of freedom. In light of this discussion, two regions for the rest of this analysis will be considered. Given user-specified gray values d_1 and d_2 for the two regions, respectively, from Eq. (17) g and s_a are obtained as below:

$$g = \frac{d_1 - d_2}{s_1 - s_2} - k_3 \quad \text{and}$$
$$s_a = \left(\frac{d_2 s_1 - d_1 s_2}{s_1 - s_2} - k_4\right) \Big/ (k_1 g + k_2) \tag{20}$$

Given desired gray values d_1 and d_2 for the two regions, respectively, the algorithm for enhancement is as follows:

(1) Acquire an image with $g = 1$ and $s_a = s_{\min}$ and compute the initial contrasts s_1 and s_2 for the two regions, respectively, from the equation below [derived from Eq. (12)]:

$$s_i = \frac{t_i - (k_1 + k_2)s_{\min} - k_4}{1 + k_3} \qquad \text{for } i = 1, 2 \tag{21}$$

(2) Compute low offset s_a and gain g from Eq. (20). Compute high offset s_b from Eq. (15).

(3) Acquire a new image with offsets (s_a, s_b).

In order to gain computational efficiency in this on-line algorithm, it is necessary to reduce computation in the digital domain. Gray values t_1 and t_2 in Eq. (21) are the averages of gray values within the two regions, respectively. Instead of processing all pixels within each region, a *small probing window* is used to gather a sample of pixels for each region. This is justified if the two regions have a uniform gray value distribution. If, however, the gray values are nonuniform, the probing window will be placed near a feature that needs to be enhanced, such as the characters in the images in Fig. 15.

Experimental Results. In the authors' experiments, laser etched characters on industrial parts were enhanced. Due to variations in the surface quality of the parts, a wide range of contrast was obtained (see Fig. 15). In order to segment the parts and characters, the hybrid method was used to obtain a set of images with *same* contrast.

First $g_{\max}$ was estimated from the parts, and $g_{\max} = 3.2$ was obtained. The details of this process is described in (10). The model constants (k_1, k_2, k_3, k_4) are estimated by the off-line algorithm shown earlier. Table 1 shows the results for ten different parts.

From Table 1 the following are observed: (1) estimates of (k_1, k_2, k_3) are close to their ideal values of $k_1 = -1$, $k_2 = k_3 = 0$, and (2) the estimates of bias constant k_4 are different from their ideal value of 0. This deviation shows that the ADC is an affine transform and not an ideal linear one. Thus the nonideal model [Eq. (12)] is a more accurate representation of the ADC than the ideal model [Eq. (10)].

First, the performance of the hybrid enhancement method is tested by computing the error between the desired and the actual enhanced gray values in the image. Let t_1 and t_2 be the average enhanced gray values for the part and background, respectively, and d_1 and d_2 be the desired gray values selected by the user. The desired gray value d_1 was changed from 160 to 250, and d_2 was changed from 10 to 60 in increments of 10, and

Table 1. Estimates of Model Constants (k_1, k_2, k_3, k_4) from Ten Parts (Ideal k_1 5 21, k_2 5 k_3 5 k_4 5 0)

Part #	k_1	k_2	k_3	k_4
1	21.033	20.036	0.094	28.917
2	20.992	20.074	0.138	29.148
3	21.045	20.000	0.120	29.205
4	21.021	20.034	0.128	29.424
5	21.026	20.034	0.124	29.419
6	21.035	20.045	0.096	28.986
7	21.028	20.037	0.107	28.971
8	21.033	20.014	0.122	29.539
9	21.041	0.010	0.115	29.121
10	21.049	0.009	0.105	28.863

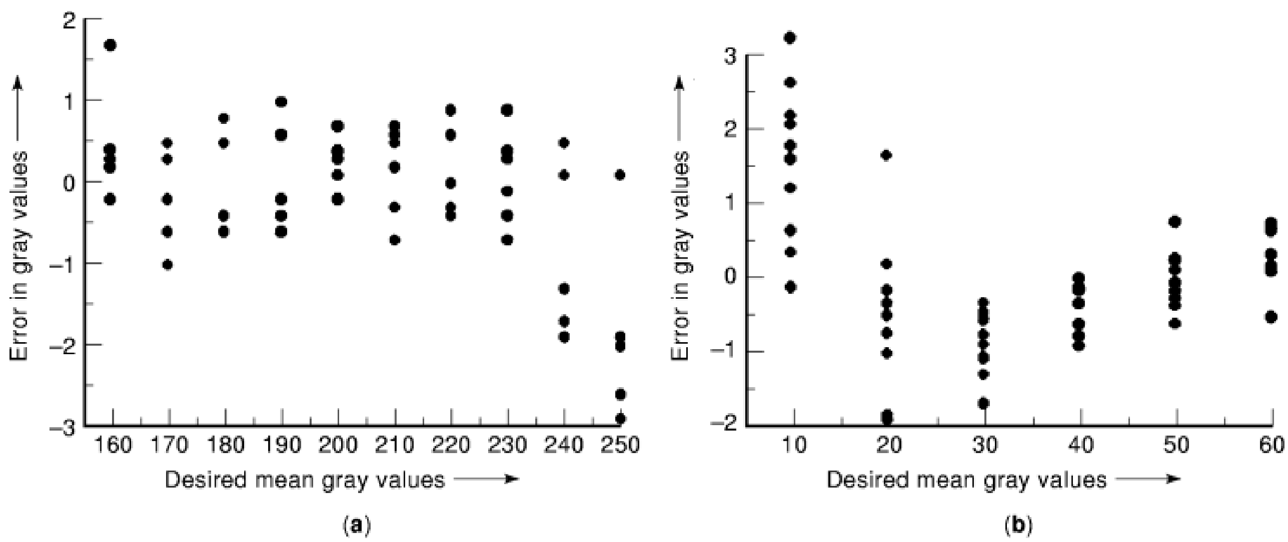

Fig. 14. Errors in enhanced image for (a) Part, and (b) Background.

reported errors $|t_1 - d_1|$ for the part and $|t_2 - d_2|$ for the background. Figure 14 shows the errors for different choices of d_1 and d_2. Here one can observe that the enhanced gray values are very close (within ± 3) to the desired values for both the part and background.

In the next experiment, the part was enhanced from the background. Enhancement results are shown in Fig. 15. The parts used in this experiment have varying initial contrast. However, all these parts are enhanced to the *same final gray value of* $d_1 = 160$, and the background has $d_2 = 0$. Clearly all parts are enhanced such that their contrast from the background are much higher. All images are now segmented by a single algorithm.

Comparison with Digital Domain Linear Contrast Enhancement. Figure 16 shows the parts in Fig. 15 enhanced by the digital domain linear contrast enhancement method to the same desired gray values, that is, $d_1 = 160$ and $d_2 = 0$ as used previously. As expected, the results are similar. However, for the simple processor under consideration, the processing time is much larger. An imaging system (6) containing a 68000 processor with a 10 MHz clock rate was used. The image resolution is 512×480 pixels and the gray values are 8bits/pixel (i.e., 256 gray levels).

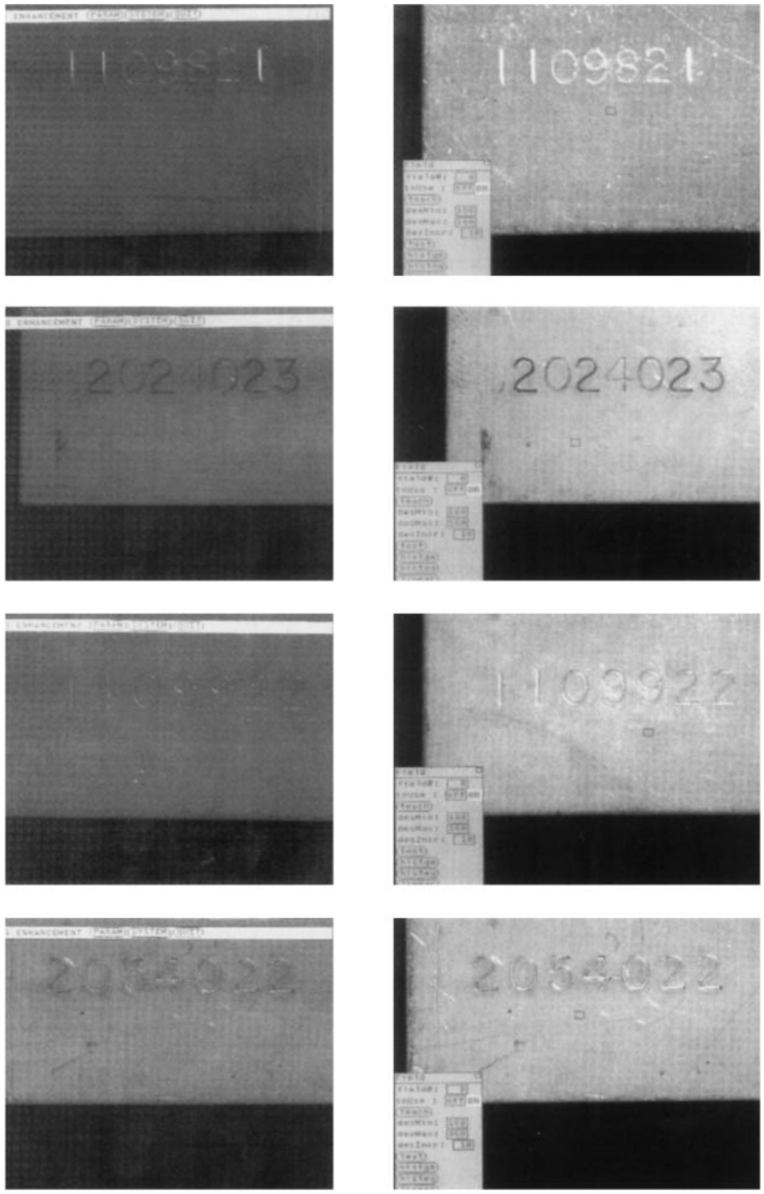

a b

Fig. 15. (a) Original images of steel samples with laser-etched characters on them, (b) Images enhanced by the hybrid method for $d_1 = 160$ and $d_2 = 0$.

The digital domain method applies the following transform f for each pixel with gray value z: $f(z) = 255 (z - z_{\rm low})/(z_{\rm hi} - z_{\rm low})$ for $z_{\rm low} \le z \le z_{\rm hi}$, $f(z) = 255$ for $z > z_{\rm hi}$, and $f(z) = 0$ for $z < z_{\rm low}$. The same probing windows as before are used for the computation of $z_{\rm low}$ and $z_{\rm hi}$. Although the computation for $(z_{\rm low}, z_{\rm hi})$ is the same as the hybrid method, the processing speed for the enhancement step is much higher, and also depends on the size of the enhancement region. Table 2 shows an example of the processing time for different region sizes.

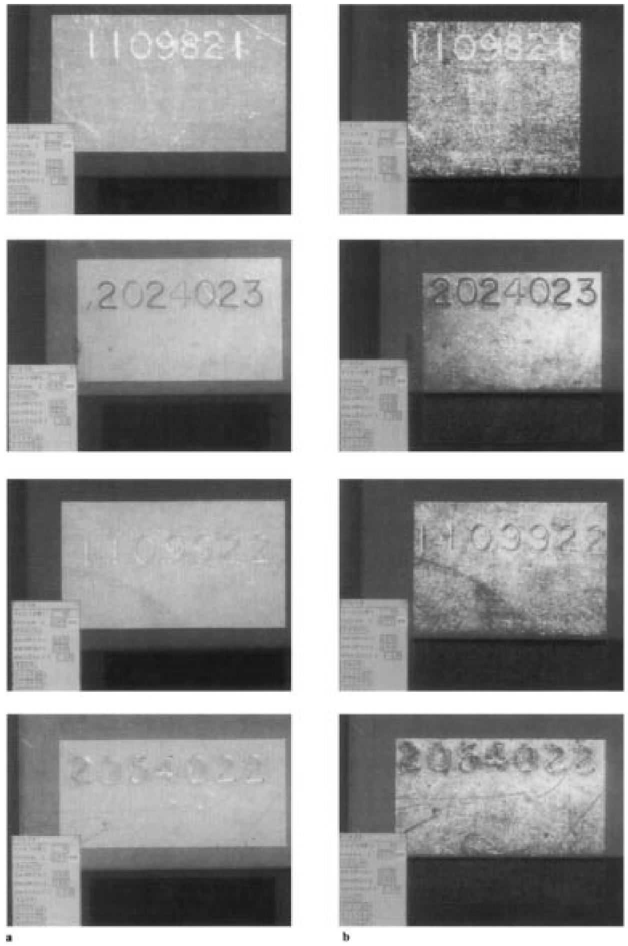

Fig. 16. (a) Images enhanced by digital domain linear contrast enhancement, (b) Images enhanced by histogram equalization.

In contrast, the hybrid method enhanced the *entire image* in 100 *ms.* Similar results are obtained for other imaging systems (7,8). Thus, for simple imaging systems, the hybrid method has a notable speed advantage. However, note that the above performance in digital domain can be vastly improved by incorporating hardware components. Thus the above speeds, although not typical, serve as an example for an inexpensive vision system.

Comparison with Global Histogram Equalization. Although histogram modification methods are powerful enhancement tools, in this application, these methods can have a poorer performance, as shown in Fig. 15. Here histogram equalization was used on the parts in Fig. 15. Figure 16 shows that histogram equalization magnifies noise drastically by spreading widely the histogram peaks due to noise in the background. Small specular reflections from the background and unnecessary surface features on the part are enhanced, and the effective contrast between the part and background is diminished. Furthermore, the processing time is comparable to the linear contrast enhancement method in Table 2. Ideally one would like to minimize the spread in the width of the individual peaks in the histogram, while pulling the peak center positions so that the contrast between the part and background are enhanced. The histogram specification methods can obtain this result, but in doing so they become similar, in speed and function, to the linear contrast enhancement

Table 2. Processing Time for the Linear Contrast Enhancement Method in Digital Domain for Different Sizes of the Two Regions Enhanced

Size of Region (pixels)	Processing Time (s)
50 3 80	12.72
75 3 80	18.79
100 3 80	24.85
125 3 80	31.16
150 3 80	37.39
250 3 80	62.16

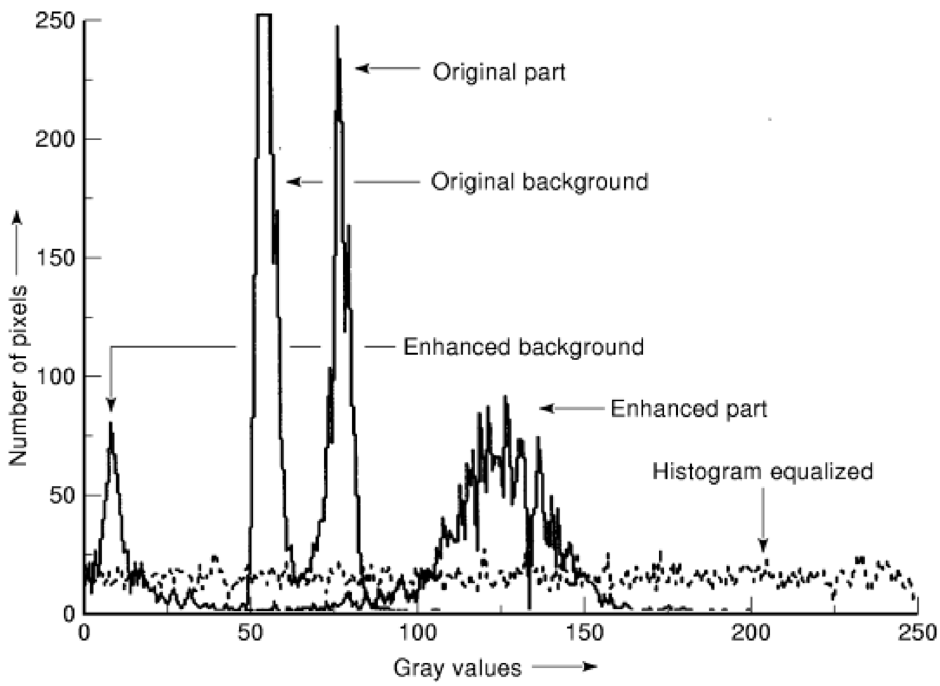

Fig. 17. Histograms for original, enhanced and equalized images.

methods discussed above. Figure 17 shows the histograms of the original, enhanced (by the hybrid algorithm) and the histogram equalized images. Clearly, the peaks for the part and background are pulled apart leading to a higher contrast image.

Generalization of the Hybrid Method. The hybrid method discussed above can be generalized to enhance multiple regions in an image. The details of this method is given in (10). Here one not only obtains a user-defined desired mean d_i for each region in the image, but one also obtains a user-defined desired standard deviation $\sigma(d_i)$ for each region. Moreover, this method can be used for multiple (≥ 2) regions in an image for enhancement.

The enhancement method was applied on six different steel parts. Figs. 18(a, c, e) show the six parts with very different surface textures, and also different contrast of the etched characters. Rectangular regions around some of these characters are identified for enhancement. For example, in Fig. 18(b) the rectangular

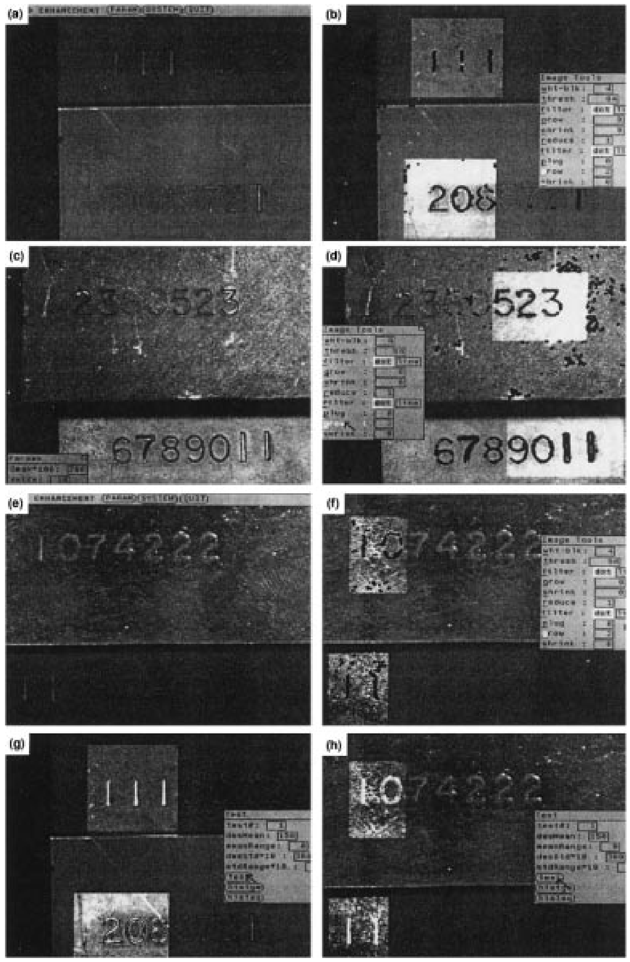

Fig. 18. (a), (c) and (e)—Camera images of steel parts with varying contrast and laser-etched characters on them; (b), (d) and (f)—Enhanced images with highlighted rectangular region in each part. The same segmentation algorithm is used in all parts to segment the characters. All characters in the enhanced regions are clearly segmented; (g) and (h)—Gray-scale enhanced images for (a) and (e), respectively.

regions enclose the characters "111" and "208" in the top and bottom parts, respectively. In Fig. 18(d), the rectangular regions enclose the characters "523" and "9011" which are enhanced. In Fig. 18(f), the rectangular regions enclose the characters "10" and "11" which are enhanced. The remaining characters in all images are not enhanced.

By using the enhancement method on each region separately, an enhanced image of characters in Fig. 18(b, d, f), respectively, was obtained. All images were enhanced to the desired mean $d_i = 150$ and desired standard deviation $\sigma(d_i) = 30$ for all regions. A segmentation algorithm was then applied to the gray scale enhanced images. The *same segmentation algorithm* was used for all images. The segmentation algorithm consists of gray scale closing, thresholding, reduction, filtering, and a binary closing in this sequence. The results show that the characters in the enhanced regions are clearly segmented.

Concluding Remarks

Three approaches to contrast enhancement were presented. Out of these, the digital domain techniques are the most popular and have wider applicability. However, a new technique based on a hybrid scheme is also discussed. The hybrid technique, although not as versatile as the digital domain techniques, is fast and can operate in near-real-time speeds. The hybrid method presented here is applicable only to enhance two regions in an image to a desired gray value distribution. This method can be generalized to enhance multiple regions and has been presented in detail in Ref. 10.

Acknowledgments

The authors thank GDE Systems Inc. of San Diego, CA, and The Math Works Inc., of Natick, MA, for providing some of the images used in this article.

BIBLIOGRAPHY

1. R. C. Gonzalez B. A. Fittes Gray Level Transformations of Interactive Image Enhancement, *Mechanism Mach. Theory*, **12**: 111–122, 1977.
2. R. C. Gonzalez P. Wintz *Digital Image Processing*, Reading, MA: Addison-Wesley, 1987.
3. R. C. Gonzalez R. E. Woods *Digital Image Processing*, Reading, MA: Addison-Wesley, 1992.
4. A. Rosenfeld A. C. Kak *Digital Picture Processing*, vol. 1, New York: Academic Press, 1982.
5. J. C. Russ *The Image Processing Handbook*, Boca Raton, FL: CRC Press, 1992.
6. AIS-3500 Vision Processor, Applied Intelligent Systems Inc., Ann Arbor, MI, 1990.
7. CVIM2, Configurable Vision Integration Module 2, Allen Bradley Corp., Milwaukee, WI, 1993.
8. Cognex-3000 Vision Processor, Cognex Corporation, Needham, MA, 1990.
9. C. Chatterjee V. P. Roychowdhury Efficient image processing algorithms for enhanced desired gray scale images, *Proc. SPIE Int. Symp. Opt., Imaging, Instrum.*, San Diego, CA, 1993.
10. C. Chatterjee V. P. Roychowdhury Models and algorithms for a real-time hybrid image enhancement methodology, *Pattern Recognition*, **29** (9): 1531–1542, 1996.
11. C. Chatterjee V. P. Roychowdhury A hybrid contrast enhancement method using off-the-shelf imaging systems, *Mach. Vision Appl.*, **9** (3): 97–105, 1996.
12. W. B. Green *Digital Image Processing: A Systems Approach*, New York: Van Nostrand Reinhold, 1983.
13. R. A. Hummel Image enhancement by histogram transformation, *Comput. Vision Graph. Image Process.*, **6**: 184–195, 1977.
14. W. Frei Image Enhancement by Histogram Hyperbolization, *Comput. Graph. Image Process.*, **6**: 286–294, 1977.
15. A. K. Jain *Fundamentals of Digital Image Procesing*, Englewood Cliffs, NJ: Prentice-Hall, 1989.
16. J. L. Harris, Jr. Constant variance enhancement: A digital processing technique, *Appl. Opt.*, **16**: 1268–1271, 1977.
17. E. Alparslan F. Ince Image enhancement by local histogram stretching, *IEEE Trans. Syst. Man, Cybern.*, **SMC-11**: 376–385, 1981.
18. R. A. Hummel Histogram modification techniques, *Comput. Vision, Graph. Image Process.*, **4**: 209–224, 1975.
19. R. Gordon Enhancement of mammographic features by optimal neighborhood image processing, *IEEE Trans. Med. Imaging*, **5**: 8–15, 1986.
20. K. Rehm *et al.* Design and testing of artifact-suppressed adaptive histogram equalization: A contrast-enhancement technique for the display of digital chest radiographs, *J. Thoracic Imag.*, **5** (1): 85–91, 1990.
21. S. M. Pizer *et al.* Adaptive histogram equalization and its variations, *Comput. Vision, Graph. Image Process.*, **39**: 355–368, 1987.
22. R. Cromartie S. M. Pizer Structure-sensitive adaptive contrast enhancement methods and their evaluation, *Image Vision Comput.*, **11** (8): 460–467, 1993.

23. R. Dale-Jones T. Tjahjadi Four algorithms for enhancing images with large peaks in their histogram, *Image Vision Comput.*, **10** (7): 495–507, 1992.
24. K. W. Leszczynski S. Shalev A robust algorithm for contrast enhancement by local histogram modification, *Image Vision Comput.*, **7** (3): 205–209, 1989.
25. ADS-927, Low Power Sampling A/D Converter, 1 MHz, 14-bit, *Datel Inc. International*, Tokyo, Japan.
26. Bt253, 18MSPS Monolithic CMOS Triple-Chan. 8-bit Image Digitizer, Brooktree Corp, San Diego, CA.
27. MP7684A, Monolithic CMOS 8-Bit Analog-to-Digital Flash Converter, Micro Power Systems Inc., Santa Clara, CA.

CHANCHAL CHATTERJEE
GDE Systems Inc.
VWANI P. ROYCHOWDHURY
UCLA