Display nonlinearity in digital image processing for visual communications

Eli Peli, MEMBER SPIE
Harvard Medical School
Schepens Eye Research Institute
20 Staniford Street
Boston, Massachusetts 02114

Abstract. The luminance emitted from a cathode ray rube (CRT) display is a nonlinear function (the gamma function) of the input video signal voltage. In most analog video systems, compensation for this nonlinear transfer function is implemented in the camera amplifiers. When CRT displays are used to present psychophysical stimuli in vision research, the specific display nonlinearity usually is measured and accounted for to ensure that the luminance of each pixel in the synthetic image properly represents the intended value. However, when using digital image processing, the linear analog-to-digital converters store a digital image that is nonlinearly related to the displayed or recorded image. The effect of this nonlinear transformation on a variety of image-processing applications used in visual communications is described.

Subject terms: cathode ray tube; displays; image enhancement; vision; halftones; color image coding.

Optical Engineering 31(11), 2374-2382 (November 1992).

1 Introduction

The luminance emitted from a cathode ray tube (CRT) display L is a nonlinear function (the gamma function) of the input video signal voltage v. The nonlinearity usually is approximated by a power law

$$L(v)=(kv)^{\gamma} \ , \tag{1}$$

where k and γ are constants. The value of the exponent γ for most displays ranges from 2.2 to 2.5. (A gamma value of 2.2 was assumed as a standard for the NTSC system.) This nonlinear transformation results in compression of the dark luminance levels and expansion of the bright levels. In most analog video systems [Fig. 1(a)], compensation for this nonlinear transfer function is implemented in the camera amplifiers.[1,2] Thus, the video signal applied to the CRT produces light distribution that is approximately linearly related to the light distribution of the original scene.

Many of the images commonly used in image processing (i.e., images distributed by the University of Southern California,[3] including the pictures of Lena and the Baboon, etc.) were digitized using a gamma-compensated camera [Fig. 1(b)]. The linear analog-to-digital converters (ADC) resulted in storage of a digital image with data values nonlinearly related to the luminances in the original scene (the gamma-compensated signal). Because the images are displayed on nonlinear displays, the stored image also is nonlinearly related to the output image being viewed. Image-processing computations are performed on the stored nonlinear data rather than on the input image. This paper describes the effects of this nonlinear transformation on a variety of image-processing applications used in visual communications. The emphasis is on image communication since the effect is meaningless and not relevant when processing nonvisual (radar, acoustic, or IR) data or when the processing remains machine related, i.e., when the processed image need not be used by an observer. I show that these effects can be of significant magnitude and therefore should be incorporated, or accounted for, in most visual communication applications of image processing, which is rarely the case.

2 Luminance Measurement and Calibration

The luminance-voltage characteristics of each display device can vary. Therefore, when CRT displays are used in precision applications (e.g., to present psychophysical stimuli in vision research), the specific display nonlinearity is measured and accounted for to ensure that the luminance of each pixel in the image properly represents the intended value.[4,5] For such critical applications, the approximation represented by the simple power function may poorly fit the display response.[2] Therefore, Pelli and Zhang[5] have proposed a more complete description of the nonlinearity form,

$$L(v)=\begin{Bmatrix}\alpha+(\beta+kv)^{\gamma} & \text{if } \beta+kv\geq 0\\ \alpha & \text{otherwise}\end{Bmatrix} \ , \tag{2}$$

where α is the minimum luminance and k and β represent approximately the effects of the monitor's contrast and brightness knobs, respectively. These authors indicated that even this function is not a good fit for some monitors and suggested that a polynomial function can provide a better fit. The fitted function can be inverted to obtain a correcting function. A curve fitting is not the only way to calibrate a monitor and correct for its nonlinearity. A table of measured luminance values can be interpolated and used to compute an inverse correction look-up table to be interposed between the frame buffer and the video digital-to-analog converter (DAC) circuitry [Fig. 1(c)].

Luminance values measured from our monochrome monitor (U.S. Pixel) are shown in Fig. 2 together with a sche-

Paper 04121 received Dec. 9, 1991; accepted for publication March 21, 1992. This paper is a revision of a paper presented at the SPIE conference Visual Communications and Image Processing, November 1991, Boston, Massachusetts. The paper presented there appears (unrefereed) in SPIE Proceedings Vol. 1606.

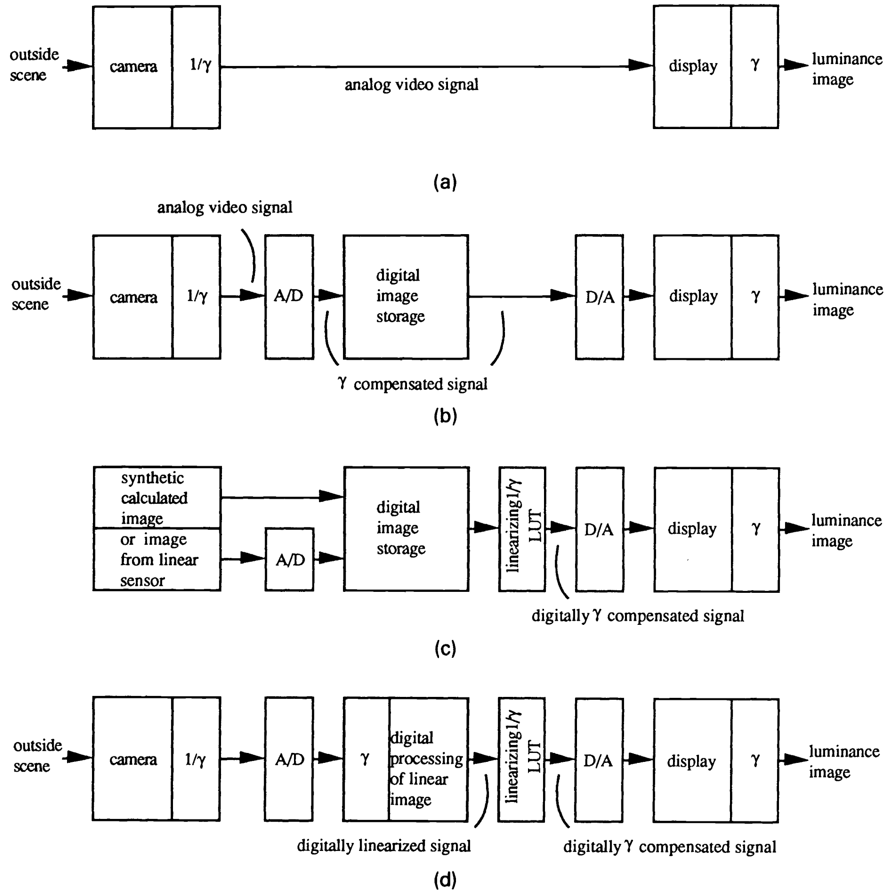

Fig. 1 Schematic illustration of the various types of systems and signals used in digital processing of video images. (a) The standard analog video system includes compensation within the camera for the nonlinear response of the display. (b) The standard digital image-processing system, where the analog video signal is digitized, resulting in storage and processing of a gamma-compensated signal rather than the luminance signal. (c) In critical applications, the synthetically generated image is displayed after digital look-up table compensation for the display's gamma nonlinear transformation. (d) In visual communication applications, where the camera and display image should remain as in (a), a linearization correction can be applied digitally before processing and then reversed using the look-up table before transfer to the display. This is the method used here to obtain a linearly displayed image of the gamma-compensated original stored image. Note the identification of various signals as used in the text.

matic illustration of the method of computing the correcting look-up table and the luminance measurement verifying the linear outcome of the procedure. With the use of a correcting look-up table, the display luminance becomes a linear function of the pixel values stored in the frame buffer [Fig. 1(c)]. Since we use an Adage display system with 10-bit DACs, we enter 10-bit words in the look-up table for our 8-bit-per-pixel frame buffer,[4] which permits a more precise presentation of low-contrast stimuli.

Luminance calibration using an accurate photometer is necessary for human vision studies. For other applications, i.e., medical displays, brightness linearization also may be desirable.[2,6] Calibrations can be achieved without a photometer using a visual adjustment task.[2,7]

One method of calibrating the display luminance involves displaying alternating raster lines of maximum (1.0) and minimum (0.0) normalized luminance values on half of the screen and a uniform adjustable luminance on the other half. The subject task is to adjust the apparent brightness of the uniform half by changing the look-up table until the two halves appear identical. The DAC look-up table value determined this way is the one corresponding to normalized luminance of 0.5. Note that if the display is characterized, for example by $k=1$ and $\gamma=2$ in Eq. (1), then the linear look-up table value that corresponds to a luminance value of 0.5 is found to be

$$v=(0.5)^{1/2}=0.707 \ . \quad (3)$$

Other values can be found by recursively repeating this procedure. For example, by replacing the bright lines with lines of 0.5 normalized luminance (i.e., 0.707) normalized pixel value, as determined in the first experiment, the uniform half of the screen can be adjusted to 0.25 normalized luminance and the DAC value for that luminance can be found. Horizontal raster lines should be used to reduce the effects of spatial nonlinearities on the luminance linearization. Various effects of spatial nonlinearity on the display luminance are addressed in detail in Naiman's[2] recent paper.

Using a linearized screen permits image processing and enhancement to be carried out without concern for the specific display to be used.[6] In Sec. 3, the consequences of displaying the result of image enhancements on a nonlinear screen are described.

3 Image Enhancement

In a linear system, image-enhancement procedures can be described and analyzed either in the frequency or the space domain, because the two are interchangeable. However, the nonlinear relationship of light distribution to the digitized image leads to an incorrect mathematical description of the actual process applied in either domain. Thus, even when an algorithm, applied without consideration of the display nonlinearity, proves to be effective, its explanation may need to be modified. The correct explanation is essential for proper application in varying conditions.

Look-up table manipulations, similar to the one described above for linearization, have been used as a simple method of enhancement. For example, an inversion of the pixel values (creating a negative image) has been found to be a useful enhancement in some cases. This process is mostly effective in increasing the visibility of image features that were dark in the original image. Such features will be transferred from the compressed end of the luminance range to the expanded end, and thus become more visible (on the nonlinear display).[8] If the same inversion is applied to an image displayed on a linearized screen, dark image features become less visible because their local physical contrast is lower when presented on a bright mean luminance background.[9] Therefore, on a linearized display, inversion should be more effective in enhancing features that appear on a bright background in the original image. In fact, look-up table modification aimed at compensating for the display nonlinear characteristics has been proposed by Lim[10] as a method of enhancement.

Histogram modification techniques such as histogram equalization or hyperbolization also can be implemented as look-up table modification. However, in these cases the look-up table structure is dependent on the original input image. Histogram equalization transforms the image to one with a more uniform histogram of gray levels.[8] The histogram and its modification are calculated from the gamma-corrected digitally stored representation of the image. Therefore, the histogram of the displayed image is shifted toward the dark levels (Fig. 3). In fact, the displayed histogram will be similar to the one calculated using the histogram hyperbolization technique (Fig. 4). However, even with this relative shift to low-luminance values, the equalized histogram displayed image still contains more values at the brighter end than the displayed original image. Thus, the enhancement achieved with histogram equalization may be, in part, the result of shifting important details from the dark, compressed display range of pixel values toward the brighter part of the display, resulting in increased contrast.

The rationale for the histogram hyperbolization was based on the compressive nature of visual brightness perception in response to illumination.[11] The logarithmic relation between brightness and displayed luminance has led to the hyperbolized histogram solution. When the method is applied, as is commonly the case, to the stored digital image and is displayed on a nonlinear display, the relation between the perceived brightness B and the pixel value remains logarithmic, if a simple power function is assumed as in Eq. (1):

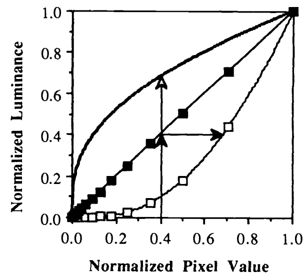

Fig. 2 Normalized measured luminance as a function of the normalized pixel values is illustrated for the nonlinear display (open squares) and following the linearization (filled squares). The look-up table used for the correction was calculated as illustrated by the arrows. For each pixel value, the corresponding linear luminance value (solid-head arrow pointing up) determines the level of luminance required (horizontal arrow) and the nearest pixel value that provides this level of luminance is obtained. The open-head arrow pointing up (equal in length to horizontal arrow) defines the value to be used in the linearization look-up table (thick line).

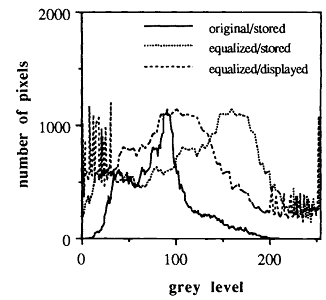

Fig. 3 Effects of the display nonlinearity on histogram equalization. The gray level histogram of the luminance values displayed on the screen is compressed toward the low values in comparison with the equalized histogram of the digital image stored in the frame buffer.

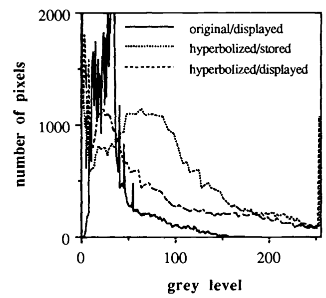

Fig. 4 Effects of the gamma transformation on histogram hyperbolization. The histograms of both the original image and the hyperbolized image are shifted toward low levels.

$$B = \log(L) = \log(kv)^{\gamma} = \gamma \log(kv) \ . \tag{4}$$

However, because the histogram hyperbolization is a nonlinear operation, the result will be modified differently than just being scaled by γ. The differences between the calculated hyperbolized histogram and the displayed one are illustrated in Fig. 4.

It is interesting to note that equalization of the stored (γ-compensated) image or of the displayed (digitally linearized) image gives essentially the same result. Displaying this equalized image on a linearized screen usually results in a too bright, washed-out image, whereas displaying the hyperbolized image on a linearized screen gives a similar ap-

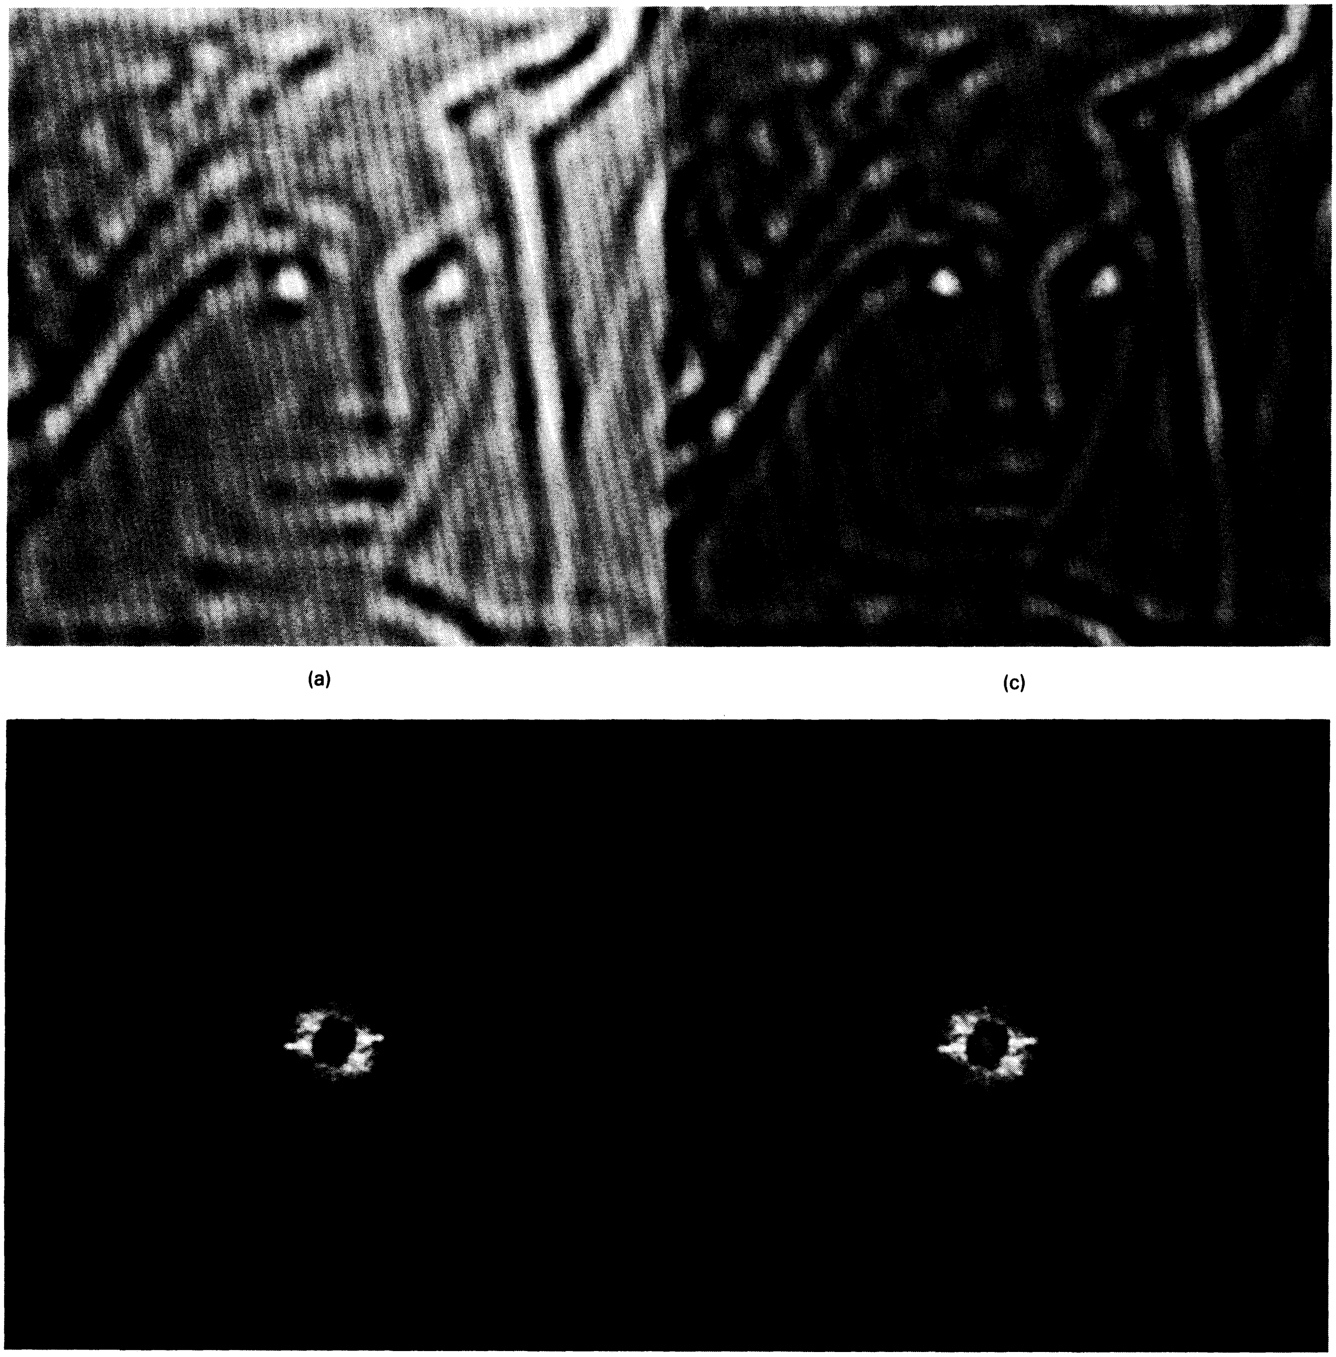

(a) (c)

(b) (d)

Fig. 5 Change in the spectral content of the bandpass filtered image. (a) The original bandpass filtered image, as stored in the frame buffer (displayed on the linearized screen). (b) The spectrum of the image in (a). (c) The appearance of the image when displayed on a nonlinearized display. (d) The spectrum of the luminance image presented in (c) contains low and high frequencies outside the range of the bandpass filter.

pearance to that of the equalized image on a nonlinear screen. This suggests that histogram equalization is effective frequently only because the gamma function changes the histogram of the displayed image to an approximately hyperbolized one.

3.1 *Spatial Filtering*

Because of the nonlinear relationship between the screen luminance and the stored image, all nominally linear operations lose their convenient linearity. Thus, even simple averaging of pixel values (used to reduce noise) does not result in averaging of the corresponding luminances.

High-pass or bandpass filtering frequently is used to enhance images. When the results of such processing are displayed on a nonlinear CRT there are two effects. The first and most obvious is that the displayed image is darker than an image properly filtered in the luminance domain (Fig. 5). This effect is the result of the necessary rescaling of the

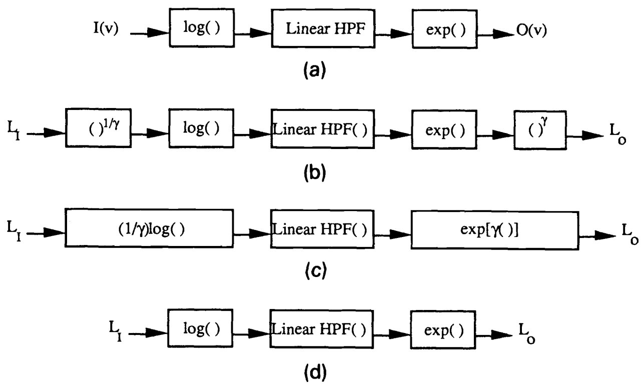

Fig. 6 Combined effect of homomorphic filtering and the display nonlinearity. (a) The calculated homomorphic filtering as applied to the stored, digital image. (b) The actual relation between the displayed image luminance before and after homomorphic filtering. (c) Combining the nonlinear operations. (d) The scaling by a factor of $1/\gamma$ before linear filtering and by γ after filtering is cancelled out to yield homomorphic filtering of the luminance displayed image.

filtered image, which is bipolar, with zero mean, into the display range of positive pixel values. The rescaling usually translates the zero value in the filtered image to 128 or close to the 128 (mid value) in the frame buffer, which results in an image with mean luminance much darker than the middle of the luminance range. The second effect, which is not as apparent when comparing the images, is that the spectrum of the displayed image contains energy at frequencies outside the passband of the filter (Fig. 5). Such ''distortion products'' have been shown to contain low-frequency spectral information corresponding to the original unfiltered image.[12] In many cases, this spurious information in the displayed filtered image may have little consequence on the effectiveness of the enhancement. However, when filtration is used in studies of the effects of various frequency bands on the recognition of images,[13–15] care should be exercised not to allow the introduction of spurious frequencies. Such nonlinearities may result from the use of nonlinearized displays or from the use of photographic slides taken off the screen, where the gamma function of the film causes further distortion.[16,17]

I have shown that high-pass-filtered images, transformed by a square power law, result in low-frequency information that is opposite in phase to the original prefiltered image.[12] Hayes and Ross[18] have shown that adding low-frequency information to a line drawing reduces recognition when the image is presented in negative form. They concluded that low-frequency information presented in inverse phase may present difficulties to the visual system. Thus, the effect of the gamma transform may reduce the enhancement of visibility obtained through the filtering. The nonlinearity also may express itself in a variety of artifacts when spatially varying filters are applied to the stored, gamma-compensated image for the purpose of noise reductions or correcting geometric distortions.[7]

3.2 *Homomorphic Filtering*

In homomorphic filtering,[19] the image is compressed by a logarithmic transformation, high-pass filtered, and then antilogged before presentation (Fig. 6a). It has been suggested that the logarithmic transformation is useful because it transforms the image, modeled as a product of high-frequency reflectance function and low-frequency illumination, to an additive function, where the high-pass filtering can be implemented easily.[19] It also was argued that since the human visual brightness response appears to present a logarithmic response to luminance intensity, the processing of images in the density (log) domain may be more appropriate (even though it is computationally expensive) and may not perform better in every application.[10,20] When the results of homomorphic filtering are presented on a nonlinear display device, the processed image first was compressed by the inverse gamma response of the camera, followed by digitally applied homomorphic filtering, and then expanded again by the gamma function of the display. However (if power law gamma function is assumed), the implementation of homomorphic filtering on the stored, gamma-corrected image is equivalent to applying the homomorphic filtering to the displayed luminance image, as illustrated in Fig. 6. Because only linear filtering is applied to the logged image, the scaling of the log transform by the gamma or inverse gamma functions remains unchanged and is cancelled out.

Schreiber[21] argued against Stockham's explanation of the effect of the homomorphic filter on the low-frequency luminance component and suggested that a significant effect occurs only with unusual images, such as the boiler room image [Fig. 7(a)]. I believe that even for these extreme cases some of the dramatic effects of image improvement obtained with homomorphic filtering are not necessarily a result of the proposed processing; instead much of the effect can be achieved simply by shifting the working point up on the nonlinear gamma curve. To illustrate this point, I processed the image by simply increasing the luminance at the dark end and keeping the bright end fixed [Fig. 7(d)]. This look-up table transformation reduces the contrast at every level of the stored image, and more significantly at the low-luminance levels. Yet, as can be seen in the image displayed on the nonlinear screen, the details at the low-luminance areas are highly enhanced [Fig. 7(d)]. Similar transformations of the low-luminance mean were included explicitly in the adaptive enhancement of this image[22] and in the original processing by homomorphic filtering.[19]

To appreciate the effect of the gamma transformation on the visibility of details in the original boiler image, one can compare the bandpass filtered versions of the stored [gamma-compensated image, Fig. 8(a)] to that of the digitally linearized image [Fig. 8(e)]. These images demonstrate that the bandpass-filtered amplitudes inside the dark room are lower for the displayed image than for the stored (gamma-compensated) one. However, since amplitude does not directly correspond to contrast, which is the visually important parameter,[9] we also have to consider the changes of local mean luminance that occur under the gamma transformation. Therefore, I calculated the band-limited local physical contrast[9] for the same image using first the digitally stored image [Fig. 8(b)] and then the digitally linearized image, representing the displayed luminance [Fig. 8(f)]. The calculated band-limited physical contrast is higher for the displayed image [Fig. 8(g)] than for the stored, linearly displayed image [Fig. 8(c)], because contrast is expressed as the ratio of amplitude to local luminance mean.[9] Only if the observer's reduced contrast sensitivity at low luminances is considered,[23] the calculated perceived contrast of the dis-

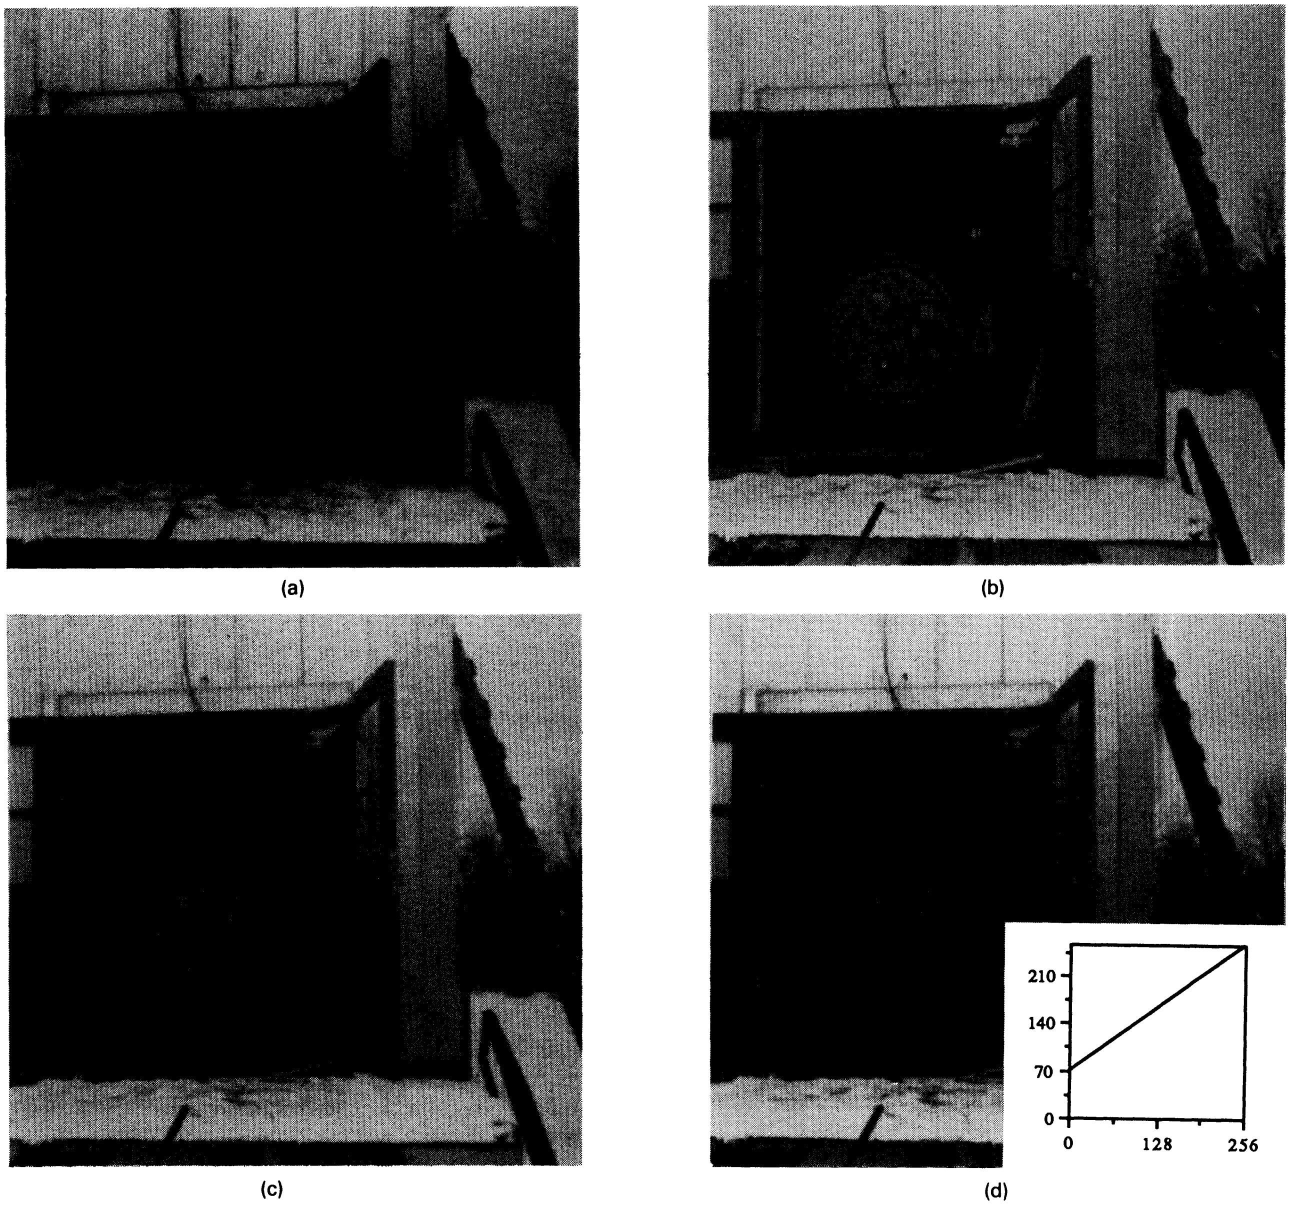

Fig. 7 Role of display nonlinearity in the effectiveness of the homomorphic filtering for the boiler image enhancement. (a) The original image as displayed on a nonlinear CRT display. (b) The image in (a) after homomorphic filtering displayed in the same way as the original. (c) The image in (a) displayed on a linearized display, resulting in increased visibility of details in dark sections of the image. (d) The image in (a) displayed on a nonlinearized display through a look-up table that shifts only the low pixel values to higher ones (inset). The details in dark sections of the display become clearly visible. Reproduction of these images in print introduces an additional nonlinear transformation; however, on the display, details in the dark area that are visible in (b) are visible in (c) and (d) as well.

played image [Fig. 8(h)] is lowered substantially, resulting again in a lower perceived contrast for the nonlinearly displayed image than for the linearly displayed stored image [Fig. 8(d)].

Note that, although much of the detail enhancement in the dark areas can be attributed to the shift in local luminance mean, all these enhancement algorithms include other factors that further improve the visibility of details at all luminance levels. These further improvements occur in the moderate- and high-luminance areas in the images.[22,24]

4 Digital Halftoning

When digital images displayed on a nonlinear CRT are transferred via digital halftoning to be printed, they commonly appear flatter and lower in contrast than the image on the CRT, because most halftoning techniques are designed to produce a fairly linear reflectance as a function of the input image.[25,26] In many cases, the reflectance is not strictly linear or even a monotonic function of the input pixel value because of the varying spreading interactions of

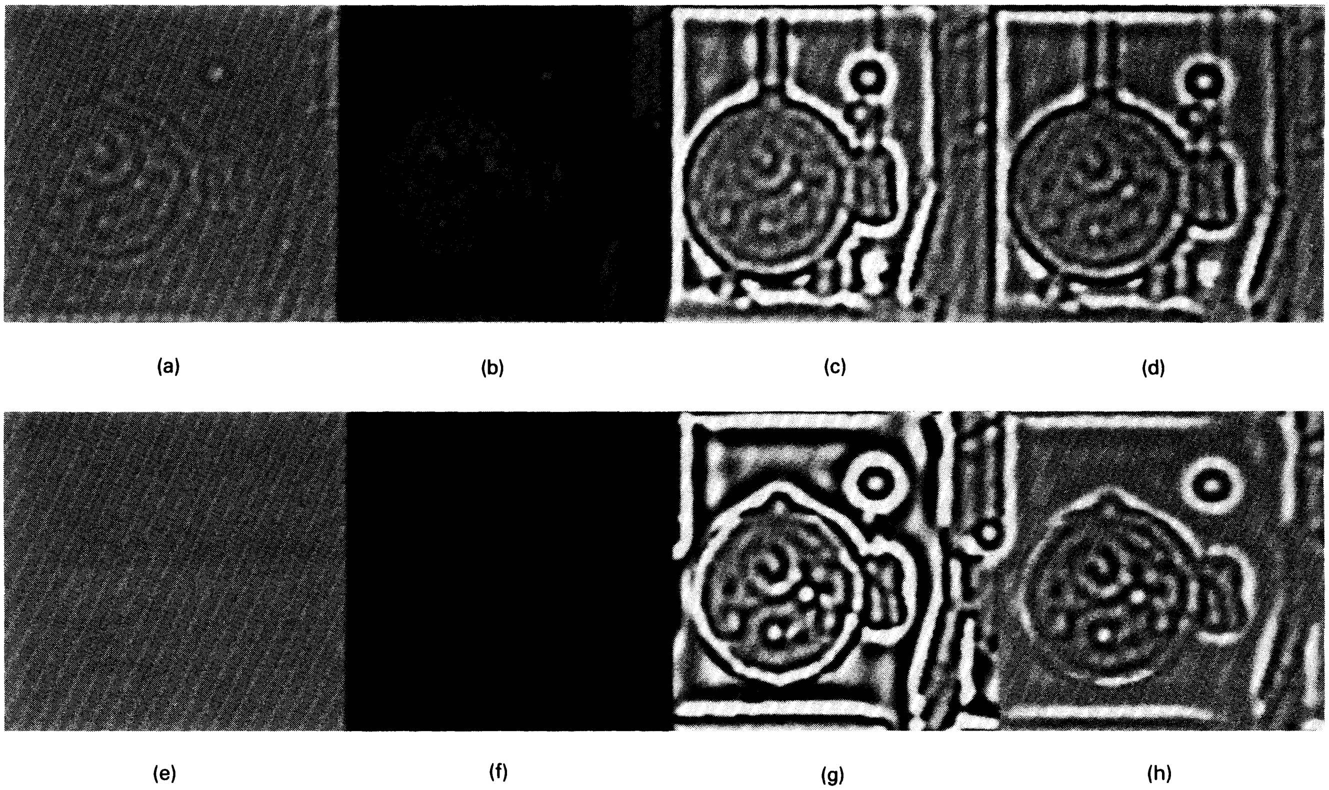

Fig. 8 Interactions of the effects of the display nonlinearity and observers contrast perception at low luminance in determining the visibility of the boiler room image. (a) and (e) The bandpass filtered versions of the gamma-compensated, stored image (b) and the digitally linearized luminance image (f), respectively, both displayed on linearized display. (c) and (g) The calculated physical contrasts. (d) and (h) The corresponding calculated perceived contrasts accounting for observer's reduced sensitivity to contrasts at low luminances. See text for details.

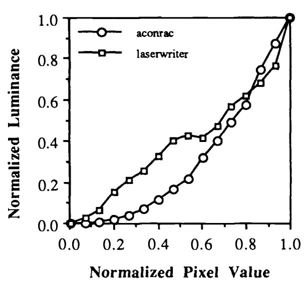

Fig. 9 Gamma characteristic of a Conrac CRT monitor is compared with the reflectance of the halftone patterns generated, on white paper, with the Apple LaserWriter for the same pixel values. Note that the halftone reflectance is less regular than the CRT but closer to linear response, especially for low pixel values.

ink on paper (Fig. 9). However, even in such cases the function is closer to linear than the typical CRT response function. The halftone pixel-to-reflectance function does not show the strong compression at the dark values, and therefore, halftone-printed images appear much lighter and usually flatter than the corresponding CRT image (Fig. 10). This effect can be quite large and requires preprocessing of the image or modification of the halftoning algorithm. Many users are satisfied with arbitrary enhancement of the image before printing, usually by application of a power function with a user-selectable exponent. For any critical application where faithful, accurate reproduction of the displayed image is desirable, calibration of both the CRT and the laser printer is required to provide the correcting look-up table. The required correction depends on the response characteristics of both the display and the printer used. Although images printed with such corrections usually appear sharper and more similar to the displayed image (Fig. 10), one often gets disturbing false contour artifacts in areas of low gradients over low luminance in the image.

5 Color Images

An analysis of the gamma artifacts associated with color image processing was presented by Miller.[7] Briefly, when using a conventional RGB system, the same gamma compensation discussed above is applied to each of the three color channels at the camera. Processing in most image-processing systems is carried in this gamma-compensated RGB domain. The luminance Y, hue U, and saturation V are expressed as linear transformations of the RGB values:

$$\begin{aligned} Y &= 0.299\mathrm{R} + 0.587\mathrm{G} + 0.114\mathrm{B} \ , \\ U &= 0.5\mathrm{R} - 0.419\mathrm{G} - 0.081\mathrm{B} \ , \\ V &= -0.169\mathrm{R} - 0.331\mathrm{G} + 0.5\mathrm{B} \ . \end{aligned} \tag{5}$$

However, because image processing by the computer actually operates on $\mathrm{R}^{1/\gamma}$, $\mathrm{G}^{1/\gamma}$, and $\mathrm{B}^{1/\gamma}$ instead of R, G, B, respectively, the result is that the Y component loses its

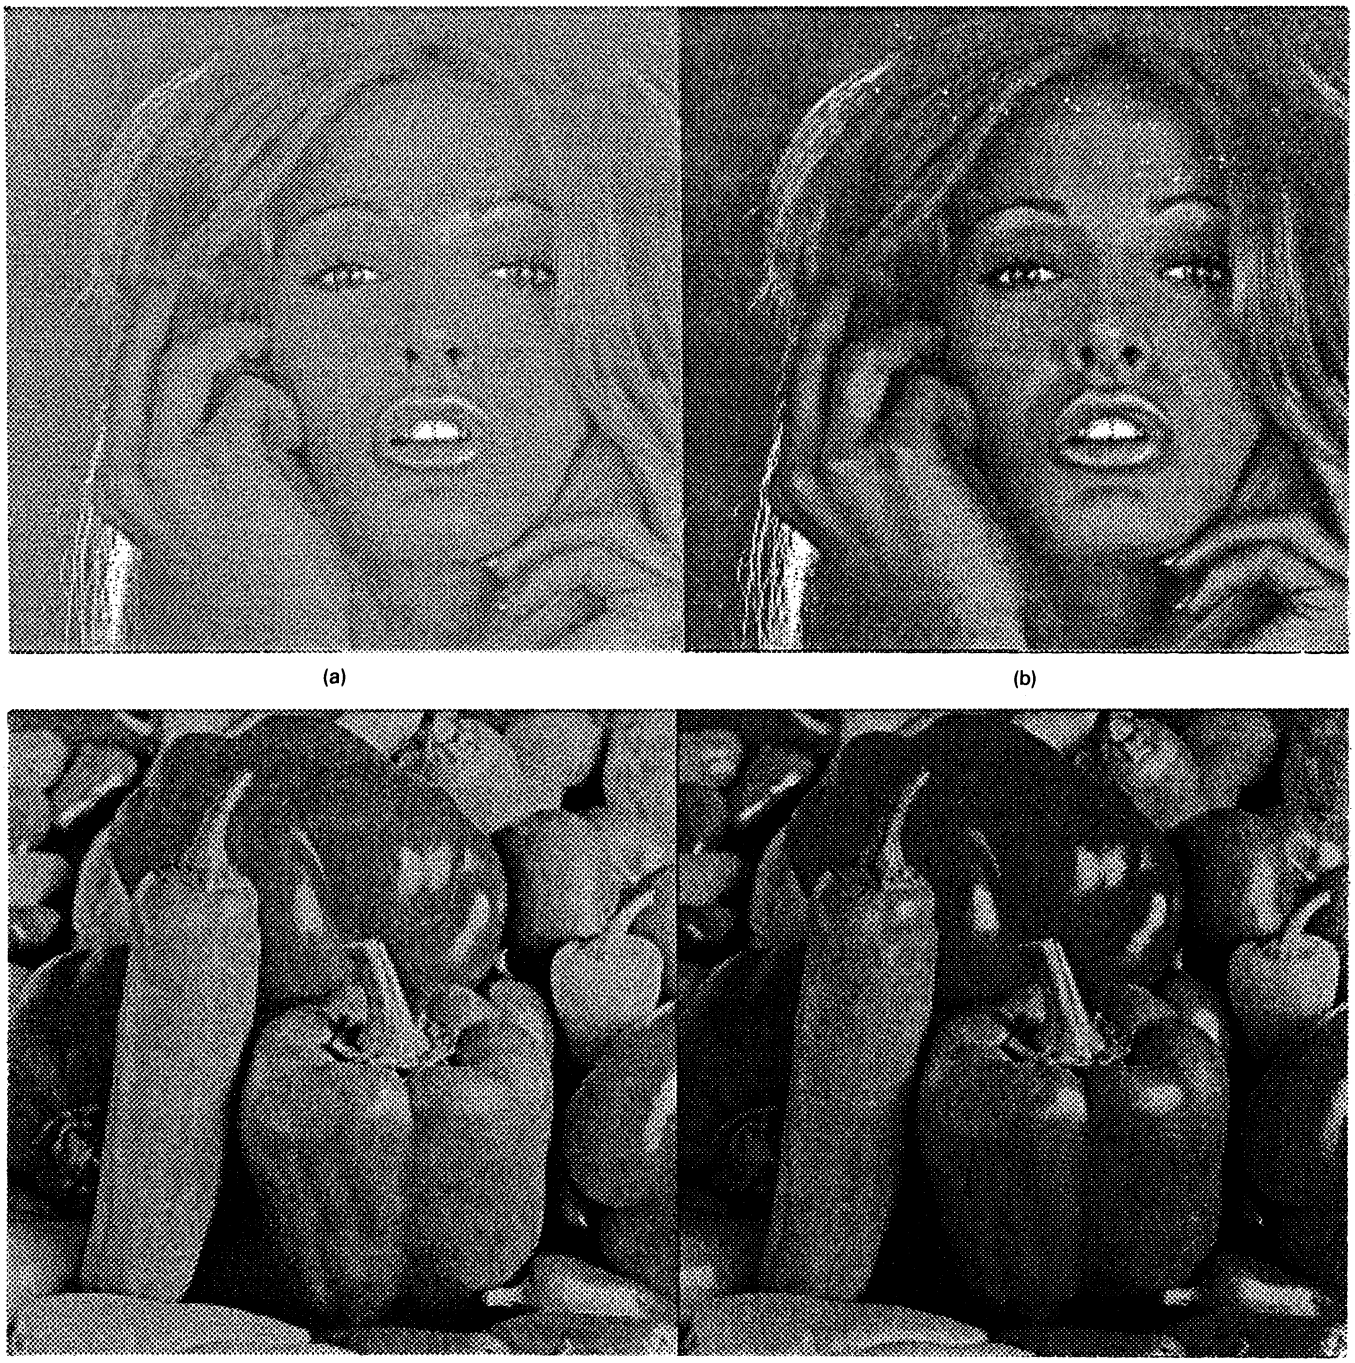

(a) (b)

(c) (d)

Fig. 10 Comparing images printed in halftone with and without correction for the gamma compensation applied in digitization. (a) and (c) Images printed without correction appear flat, of lower contrast, and brighter than the CRT displayed image. (b) and (d) Images printed after being transformed by a gamma power law corresponding to the measured CRT response.

pure correspondence with physical luminance, and similarly, the U and V coordinates do not correspond to pure hue and saturation. Thus, for example, an equal change in R, G, and B that should represent only a change in luminance in fact "leaks" some changes into the U and V components.

When processing is done in the "YUV domain,"[27,28] it is in fact carried in the $Y'U'V'$ domain where, for example,

$$Y' = 0.299\mathrm{R}^{1/\gamma} + 0.587\mathrm{G}^{1/\gamma} + 0.114\mathrm{B}^{1/\gamma} , \tag{6}$$

which is quite different from the luminance channel. This becomes critical in many cases when the bandwidth used to transmit the chroma information U and V is much lower than the bandwidth used to transmit luminance. The result of subsampling U' and V', as is commonly done[27,28] is, therefore, to reduce the resolution of the portions of the Y signals that have leaked into the primed chroma channels.

6 Conclusion

In vision research and other critical medical and scientific applications, the best solution for the problems discussed in this paper is to use a linear acquisition camera and correct for the display nonlinearity using a display-dependent look-up table [Fig. 1(c)]. However, in applications related to mass-media video signal processing, it is important to maintain the normal nonlinear relationship between the video voltages and the displayed images. In such applications, or whenever the images are already available in the standard gamma-compensated format, digital images can be converted by a gamma transformation before processing [Fig. 1(d)], allowing the processing to be applied in an approximately linear luminance domain. Following the processing, the image should be displayed on a linear display or transformed back by the gamma-compensating function before being presented on the display [Fig. 1(d)].

These transformations are not without problems. When applied in a fixed point (8-bit) system, they result in reduced resolution. Depending on the specific nonlinearity, the number of usable levels may be reduced by as much as 30% (i.e., 700 out of 1024 for a 10-bit system).[4] In some cases, noticeable transitions causing false contours can result. The commonly used, 8-bit resolution is, in fact, inadequate for quantization of the linear luminance domain, especially at the dark end of the range. Thus, images we transformed to be displayed on the linearized display frequently have false contours appearing in dark sections of the image. The loss of resolution can also be noted in Fig. 8, where Fig. 8(f) is the digitally linearized version of Fig. 8(b). These problems can be avoided or reduced if the standard analog video signal (or video signal that was digitized at higher than 8-bit resolution) is transferred to the luminance linear domain using an analog converter before digitization.

Accounting for the gamma transform of digital images in video communications systems also enables proper incorporation of the human visual model and permits image compression by quantization that is based on actual image amplitude (contrast), rather than a nonlinearly transformed version of the image.

Acknowledgments

Supported in part by grant #EY05957 from the National Institutes of Health, by the James and Alice Teubert Charitable Trust, and by the Ford Motor Company Fund. I thank George Young for technical help and Larry Arend for valuable discussions.

References

1. A. Benson, *Television Engineering Handbook,* pp. 2.26–2.29, McGraw-Hill, New York (1985).
2. A. C. Naiman, "CRT spatial non-linearities and luminance linearization," in *Raster Imaging and Digital Typography,* R. Morris, Ed., pp. 43–53, Cambridge University Press, Cambridge, England (1991).
3. A. G. Weber, "Image data base," USC Signal and Image Processing Institute, Report 101, Los Angeles, University of Southern California (1988).
4. A. B. Watson, K. R. K. Nielsen, A. Poirson, A. Fitzhugh, A. Bilson, K. Ngugan, and A. J. Ahumada, "Use of a raster framebuffer in vision research," *Behavior Res. Meth., Instr. Comput.* **18,** 587–594 (1986).
5. D. G. Pelli and L. Zhang, "Accurate control of contrast on microcomputer displays," *Vision Res.* **31,** 1337–1350 (1991).
6. S. M. Pizer, J. B. Zimmerman, and R. E. Johnston, "Contrast transmission in medical display," in *Proc. First Int. Symp. on Medical Imaging, Image Interpretation ISM III,* IEEE, pp. 2–9 (1982).
7. R. L. Miller, "The effects of gamma on digital image processing," Unpublished manuscript (1985).
8. W. K. Pratt, "Image enhancement and restoration," in *Digital Image Processing,* pp. 307–407, Wiley & Sons, New York (1978).
9. E. Peli, "Contrast in complex images," *J. Opt. Soc. Am. A* **7,** 2032–2040 (1990).
10. J. S. Lim, "Image enhancement," in *Digital Image Processing Techniques,* M. P. Ekstrom, Ed., pp. 1–51, Academic Press, New York (1984).
11. W. Frei, "Image enhancement by histogram hyperbolization," *Comput. Graph. Image Process.* **6,** 286–294 (1977).
12. E. Peli, "The perception and interpretation of highpass filtered images," *Opt. Eng.* **31**(1), 74–81 (1992).
13. A. Fiorentini, L. Maffei, and G. Sandini, "The role of high spatial frequencies in face perception," *Perception* **12,** 195–201 (1983).
14. R. A. Schuchard and G. S. Rubin, "Face identification of band pass filtered faces by low vision observers," *Invest. Ophthalmol. Vis. Sci.* **30**(Suppl.), 396 (1989).
15. D. H. Parish and G. Sperling, "Object spatial frequencies, retinal spatial frequencies, noise and the efficiency of letter discrimination," Vision Res. **31,** 1399–1415 (1991).
16. T. Hayes, M. C. Morrone, and D. C. Burr, "Recognition of positive and negative bandpass-filtered images," *Perception* **15,** 595–602 (1986).
17. J. Norman and S. Ehrlich, "Spatial frequency filtering and target identification," *Vision Res.* **27,** 87–96 (1991).
18. T. Hayes and J. Ross, "The visual properties of lines in line drawings," *Perception* **17,** 357 (1988).
19. A. V. Oppenheim, T. G. Schafer, and T. G. Stockham, "Nonlinear filtering of multiplied and convolved signals," *Proc. IEEE* **56,** 1264–1291 (1968).
20. T. G. Stockham, Jr., "The importance of nonlinear enhancement in visual prosthesis design," in *Visual Prosthesis in the Interdisciplinary Dialog,* T. D. Sterling, E. A. Bering, S. V. Pollack, and H. G. Vaughan, Eds., pp. 173–192, Academic Press, New York (1971).
21. S. W. Schreiber, "Image processing for quality improvement," *Proc. IEEE* **66,** 1640–1651 (1978).
22. T. Peli and J. S. Lim, "Adaptive filtering for image enhancement," *Opt. Eng.* **21,** 108–112 (1982).
23. E. Peli, J. Yang, R. Goldstein, and A. Reeves, "The effect of luminance on superthreshold contrast perception," *J. Opt. Soc. Am.* A**8,** 1352–1359 (1991).
24. V. T. Tom and G. J. Wolfe, "Adaptive histogram equalization and its applications," in *Application of Digital Image Processing IV, Proc. SPIE* **354,** 204–209 (1982).
25. H. S. Hou, *Digital Halftoning and Shading in Digital Document Processing,* pp. 83–115, Wiley & Sons, New York (1983).
26. R. Ulichney, *Digital Halftoning,* The MIT Press, Cambridge, Mass. (1987).
27. A. Leger, T. Omachi, and G. K. Wallace, "JPEG still picture compression algorithm," *Opt. Eng.* **30,** 947–954 (1991).
28. K. N. Ngan, H. C. Koh, and W. C. Wong, "Hybrid image coding scheme incorporating human visual system characteristics," *Opt. Eng.* **30,** 940–946 (1991).

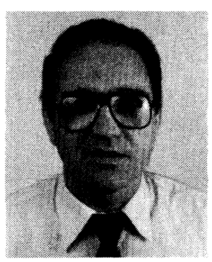

Eli Peli received a BSc degree, cum laude, and MSc degree, both in electrical engineering, from the Technion-Israel Institute of Technology in 1976 and 1979, respectively, and a doctorate from New England College of Optometry in 1983. He is a senior scientist and head of the Physiological Optics Unit at the Schepens Eye Research Institute and is on the faculty of Tufts University School of Medicine, Harvard Medical School Department of Ophthalmology, and is director of the Low Vision Service at New England Medical Center Hospitals in Boston. Dr. Peli's principal research interests are image processing in relation to visual function and clinical psychophysics. He also maintains an interest in oculomotor control and binocular vision. He is a fellow of the American Academy of Optometry and a member of the Association for Research in Vision and Ophthalmology, IEEE, SID, and SPIE.