PROCEEDINGS OF THE
SOCIETY FOR INFORMATION DISPLAY

SID

REPRINTED FROM
VOLUME 26 NO. 2 ● 1985
PAGES 101–107

FUNDAMENTAL REQUIREMENTS FOR PICTURE PRESENTATION

M. F. Cowlishaw
IBM UK Scientific Centre
Winchester, UK

Abstract—"How many bits per picture element does a display need?" For the display resolutions now commonly in use, only 4 bits per picture element (pel) are needed for the display of monochrome images. A total of 8 bits per pel are needed for color images. These conclusions are based both on experiment and on the theoretical model of the visual system in which the detectors in the eye are treated as simple photon-detectors. The results are applicable to "natural" images and also to computer-generated images. A particular 8-bit color encoding scheme is described that also has the advantage that natural images are displayable on monochrome displays.

FUNDAMENTAL REQUIREMENTS FOR PICTURE PRESENTATION

M. F. Cowlishaw
IBM UK Scientific Centre
Winchester, UK

Abstract—"How many bits per picture element does a display need?" For the display resolutions now commonly in use, only 4 bits per picture element (pel) are needed for the display of monochrome images. A total of 8 bits per pel are needed for color images. These conclusions are based both on experiment and on the theoretical model of the visual system in which the detectors in the eye are treated as simple photon-detectors. The results are applicable to "natural" images and also to computer-generated images. A particular 8-bit color encoding scheme is described that also has the advantage that natural images are displayable on monochrome displays.

I. INTRODUCTION

"How many bits per picture element does a display need?" is a common question asked by display designers and users. Answers to this question vary widely, but are usually based on research in which the sensitivity of the human visual system (HVS) is measured by observing various kinds of sinusoidal gratings. This research suggests that about 7 bits are needed for a monochrome picture, and it is often assumed that therefore three times as many are needed for a color picture.

My research has shown that these are overestimates. For the display resolutions now commonly in use, only 4 bits per picture element (pel) are needed for the display of monochrome images. A total of 8 bits per pel are required for color images. These conclusions are based both on experiment, and also on the theory of the visual system in which the detectors in the eye are modelled as simple photon detectors. The results are applicable both to "natural" images (from photographs and other natural sources of images) and to computer-generated images. A particular 8-bit color-encoding scheme is described that has the advantage that natural images are displayable on monochrome displays.

Although this paper mainly refers to the presentation of pictures on electronic display devices, the general conclusions and theory are equally applicable to hardcopy output.

II. THE FREQUENCY MODEL

For many years researchers have described and measured the limitations of the HVS in terms of a graph of contrast sensitivity plotted against spatial frequency. Many workers have presented results in this area (Mannos and Sakrison[1] present several results on a single graph), and Campbell's elegant demonstration of the curve[2] is a standard illustration in textbooks. Detailed investigation has shown that this graph can probably be described as the envelope of a number of bandpass filters in the HVS, but for the purposes of this paper we may use the curve as it stands.

For display researchers, an especially useful presentation of the curve is that in which the number of gray levels discernible is plotted against spatial frequency. The number of gray levels discernible is directly related to contrast sensitivity, which is the reciprocal of the contrast (contrast is defined here as the difference in intensity between an object and its background, divided by the intensity of the background). Such a curve (after Robson[3] and Mannos and Sakrison[1]) is shown in Fig. 1. The vertical axis is calibrated in the number of levels discernible, and also in the number of bits required to represent those levels. The horizontal axis shows spatial frequency in cycles per degree, and the equivalent in picture elements (pels) per millimeter at normal viewing distance (400 mm). Variations in gray levels or detail outside the shaded area are not detectable to the average observer; any combination inside the shaded area will normally be detected under suitable viewing conditions.

This curve is usually measured using sine-wave gratings whose frequency and contrast are varied. We may therefore read from the graph that, for example, observers will (on average) not be able to detect any sinusoidal grating which varies by less than one level in 190, or which has a spatial frequency greater than 60 cycles/deg. A typical image display with a raster of 4 pels/mm (14 cycles/deg at 400 mm) and 256 gray levels will therefore

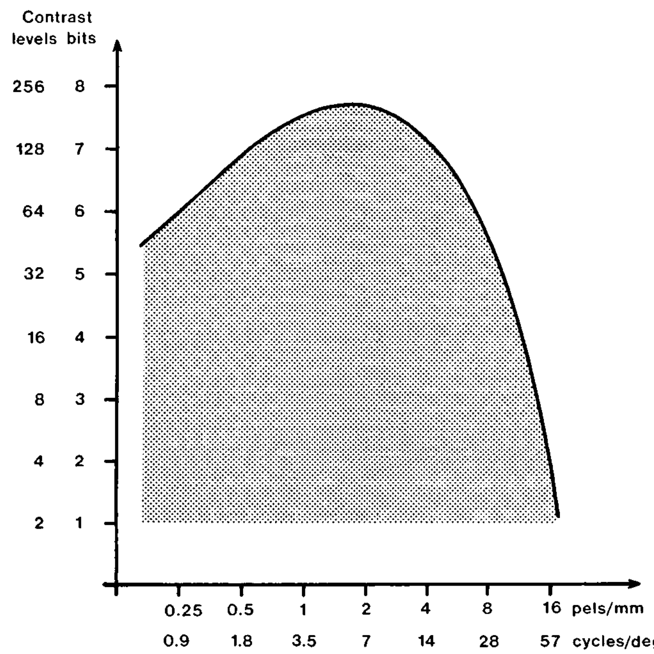

FIG. 1. Contrast sensitivity of the visual system, as a function of spatial resolution (after Robson[10] and Mannos and Sakrison[8]).

exceed the limitations of the HVS for gray levels, yet does not provide the spatial detail that the observer could resolve.

This graph is also useful for presenting other information, as we shall see.

III. THE TARGET MODEL

The curve in Fig. 1 is derived from measurement of the ability of the eye to detect linear features (gratings). It is instructive to also consider how well the HVS will perform at detecting a target (here defined as a small area with equal horizontal and vertical dimensions, and differing only in intensity from its background).

A forced choice experiment was conducted using 16 observers to measure this relationship, and I was surprised at the time to find that detection of a target of given contrast was proportional to the linear size of the target instead of to its area.

After some study, it was found that this is predicted by what I shall call the target model of the HVS. Several workers, notably Rose,[4] Schnitzler,[5] and Sturm and Morgan,[6] have considered the performance of the eye as modelled by an ideal photon detector. Blackwell's experiments in the 1940's[7] confirmed that over a significant range this model is appropriate, and the current research has shown that it is certainly applicable for CRT displays over at least the normal conditions of use as a computer output device. It is almost certainly applicable to all forms of visual presentation, including paper-based technologies.

The target model may be used to derive a simplified version of the formula that relates the various parameters that affect the detection of a small target against a constant background:

$$CA = \frac{kS}{D\sqrt{NTQ}}$$

where

- C = the contrast of the target when it is just detectable
- A = the angular size of the target
- k = constant that depends solely upon the units of the other terms
- S = the signal-to-noise ratio needed for reliable detection
- D = the diameter of the collection aperture (the pupil of the eye)
- N = the number of incident photons per unit area in unit time
- T = the integration time of the detector
- Q = the quantum efficiency of the detector.

Using this formula we find that for conditions of constant background intensity, quantum efficiency of the detector, aperture size, etc., then contrast multiplied by the angular (linear) size of a target should be constant. Using the approximate values suggested by Schnitzler and others for the terms in the formula, it was found that contrast multiplied by target size should equal approximately 16 minutes of arc under the viewing conditions of the experiment. The figure previously calculated from my experimental results was 12.5 minutes – a remarkably close result in view of the large approximations and ranges of the terms in the formula, and well within the limits of experimental error and observer variation.

There is thus both theoretical justification and experimental evidence that the limitation on detection of a target of a given contrast is proportional to its angular size, rather than to its area. Modelling the receptors of the eye as simple photon detectors is a valid method for describing their ability to detect targets, especially when the model is calibrated by experimental results. The limits of the HVS as derived from this target model may be plotted on the same graph as the frequency model. The two curves differ significantly, and the next section investigates why this might be so.

IV. TWO CONTRADICTORY MODELS?

If we plot the limits suggested by the target model (as determined by experiment) on the same graph as Fig. 1 (the limits, also determined by appropriate experiments, as suggested by the frequency model), we get the combined graph shown in Fig. 2.

The striking feature of the combined graph is that there is a large part of the area under the original curve that is above the limit found for target detection. If we look at the portion of the graph at 4 pels/mm, we see that from the frequency experiments (using gratings) we should be able to detect features differing by about 1 part in 128 (7 bits of gray level), yet from the target model curve and from experiment we know that this is not always so.

The explanation for this is of course that the two curves were measured in different ways: one measures the detection of symmetric small patches (targets), and the other measures the detection of gratings (targets greatly

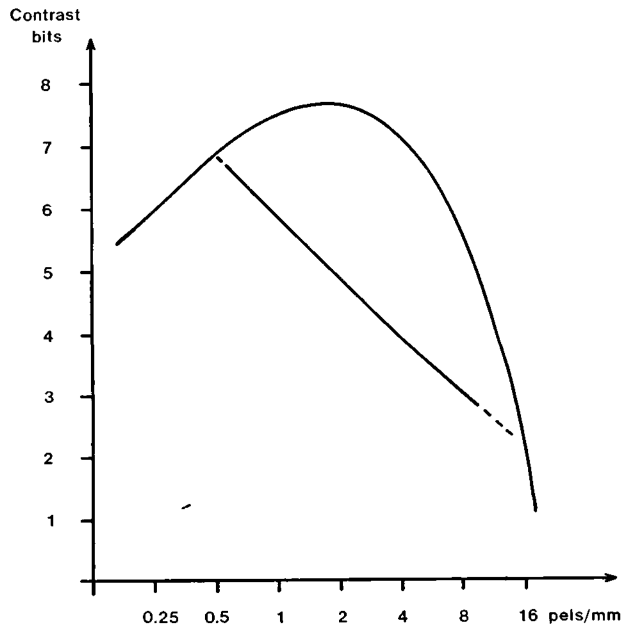

FIG. 2. As Fig. 1, also showing the target detection curve.

extended in one dimension). There are two possible reasons for gratings being more detectable than targets: it might be that the regular pattern evokes some kind of resonant response in the HVS; or it might simply be that the long dimension of each bar in the grating makes it more visible despite its small width.

To test which of these is the case, an informal experiment was carried out. In this experiment the visibility of a single bar from a grating (with frequencies up to 15 cycles/deg) was compared with that of the grating as a whole. It proved to be equally visible, indeed if anything more visible, which indicates that no resonant effect caused by the multiple bars is improving their visibility. (Relative visibility was simply measured by changing the viewing distance to find the point at which the object being observed merged with its background. This point is quite abrupt and repeatable.)

The single bar was then gradually reduced in length: it remained equally visible until its length became about five times its width, at which point its visibility deteriorated down to the point at which it became the same size (and visibility, of course) of the target of the appropriate size. Little difference was observed in the results for horizontal and vertical bars and gratings.

One conclusion that may be drawn from these observations is that the frequency-based (grating) model and curve will in fact describe the limitations of the HVS for line-like features such as single bars or wires (but not edges, unless they are of very low contrast), where the length of the feature is at least five times its width. The target model curve will describe the limitations of the HVS for the more regular type of feature whose width and height are similar. Objects between these two descriptions would fall in the area between the curves shown in Fig. 2.

These observations explain why half-toning (with two or more gray levels) works so well. Suppose we reduce a 4-pels/mm 256-gray-level picture to the 16 gray levels suggested by the target detection curve, and some area in the original picture is at a gray level midway between two of the output possibilities. By representing the intermediate level by a pattern in which 50% of the pels are set to the level above the desired level and the remainder are set to the level below (preferably randomly distributed) then the eye will not be able to detect the individual pels (which admirably fit the definition of a "target"), and so the area will appear to have a smooth gray appearance at the desired (intermediate) level.

A halftoning method which produces relatively few artificial linear features (such as Floyd and Steinberg's Error Diffusion algorithm[8]) will therefore look better than one (such as Judice, Jarvis, and Ninke's Ordered Dither algorithm[9]) that tends to produce linear features which by their very nature are most easily detected by the eye.

Another conclusion that may be drawn from these observations is that a small target can be made more visible if its size is increased in just one dimension, up to the point where one dimension is five times the other. A feature of this shape is detectable, even though it is apparently not sufficiently wide for detection. This would seem to confirm that the detectors in the eye are not independent, but instead have the capability of integrating events locally and can therefore act as a larger and more sensitive (or reliable) detector.

A. *Practical Observations*

The graph shown in Fig. 2, with both curves plotted, provides valuable insight into several observations. On a 4-pel/mm display, pictures displayed with 4 bits per pel (used fairly optimally by applying a halftoning algorithm such as error diffusion) are almost indistinguishable from the same picture displayed with 8 bits per pel. (If a slice near the center of an 8-bit-per-pel image is replaced with the same data error diffused to just 4 bits per pel, it is usually impossible to locate the slice, except with close inspection.) This result is contrary to that which would be predicted from looking at the upper curve in the graph, which indicates that at least 7 bits need to be used to reach the HVS limit. I suggest that in fact features in real pictures are of generally high contrast (or if of low contrast are rarely linear) and that almost invariably they will fall below the lower curve. Some pictures can be presented with just simple thresholding to 16 levels (that is, by using the four most significant bits), but a good halftoning algorithm allows any picture – including computer generated pictures – to be treated as though they consisted of just regular ("target") features. For most practical purposes we may therefore use the lower target detection curve to design our displays rather than the more demanding (and expensive) upper curve.

Certain applications – such as radiography – do require that low-contrast linear features be displayable, and for simplicity it might be wiser to use the upper curve as the guide for specialist research displays. In many cases, though, it will make more sense to process the image to bring the dynamic range of the image within the lower curve, hence increasing the probability of detection of all types of feature. As a general rule, enhancement by image processing should always aim to bring the dynamic range of the features to be detected within that defined by the lower curve, so that they will be detectable by the observer whatever their shape.

Variation between observers may also be taken into account, though in practice it has not been found to be a major factor. In my experiment, the contrast sensitivity of the observers varied by up to one-half of 1 bit above or below the line shown in Fig. 2. This variation can be explained by differences in visual acuity and pupil size of the observers.

B. *Using the Contrast Detection Formula*

In the discussion above, it has been implied that if a target requires a contrast of 1/16 (one part in 16) to be detected, then only 16 gray levels are required in a display to depict that target at various luminances. For practical purposes this is a good approximation, but it is possible to use the formula given above to determine more precisely the number of levels required, and what the luminance of each level should be.

To use the formula accurately, it is helpful if the various terms are expressed in units familiar to the display scientist. It is possible to derive from basic principles that for light with a wavelength of 555 nm:

$$C = \frac{kSV}{PD\sqrt{LTQ}}$$

where

- C = the contrast of the target when it is just detectable
- k = a constant that depends upon the units of the other terms and the wavelength of the light, and is equal to 1.93×10^{-8} in this formula
- S = the signal-to-noise ratio needed for reliable detection of the target
- V = the viewing distance (m)
- P = the diameter of the pupil (m)
- D = the diameter of the circular target (m)
- L = the luminance of the target (cd/m^2)
- T = the integration time of the detector (sec)
- Q = the quantum efficiency of the detector

For this formula to be accurate for all contrast levels, the contrast (C) must here be defined as the difference between the luminance of the target and the luminance of the background, divided by the luminance of the target (not the background).

This formula is a complete description of the fundamental information transmission system formed by the image displayed and the receptors of the eye. Typical figures for the human elements of this system are (again for a wavelength of 555 nm): S = 5, V = 0.4 m, P = 0.003 m, T = 0.1 sec, and Q = 0.14.

These figures may be used in a trial calculation in which the luminance of the screen (L) is 50 cd/m^2, and the pel size (D) is 0.25 mm. In this case the result is that C = 0.0615, very close to the value of 1 part in 16 described earlier.

Given the contrast required at 50 cd/m^2, we may therefore determine the luminance of the background for a pel of that luminance to be just visible (in this case 46.9 cd/m^2), and this value may in turn be used to calculate the next step. This process may be continued until the minimum (background) luminance of the screen is reached.

V. COLOR PICTURES

The preceding sections describe the performance of the HVS under optimal conditions, where the picture being observed is monochromatic with a hue to which the eye is most sensitive (i.e., green). The same results apply so long as the green component of the color of the picture is at least as large as any other component, as with white, yellow (amber), or green displays.

Unfortunately, so far as I have been able to determine, no experimenter has directly measured how contrast and spatial frequency sensitivity vary simply with the wavelength of light (though Campbell and Durden[10] do present results for the variation of vernier visual acuity with wavelength, and Thorell[11] refers to some related work). We do have, however, the well-known curve of total eye sensitivity as a function of wavelength (see Judd's modification of the CIE 1931 curve,[12] and Vos[13]). The formula that models the performance of the eye receptors as photon detectors shows that the contrast required for a given target to be detectable is proportional to the square root of the efficiency of the detector. By re-ordering and simplifying the formula given earlier in this paper we can show that K, the contrast sensitivity (the reciprocal of contrast) is given by

$$K = Zd\sqrt{Q}$$

where

- Z = effectively a constant for the eye system over the range of normal luminance of displays (it depends upon the terms for luminance, pupil diameter, signal-to-noise ratio, etc.)
- d = the diameter of the target
- Q = the efficiency of the detector

From the current experiment, a typical value for Z is 64 if d is the size of the target (expressed in mm and viewed at 400 mm), and Q is expressed as relative efficiency with a value of 1 at green (555 nm).

If we take the values from the eye sensitivity curve at wavelengths of 450 and 660 nm (blue and red) we find that the eye is approximately one-sixteenth as sensitive to these colors as it is to green. If Q is one-sixteenth of the value at green, then K must, in turn, be one-quarter of its value at green (the square root of the factor for Q). For a target to be detectable in these colors, it must therefore have four times the contrast of a green target of the same size and power.

There is much published evidence to support this result. For example, several experimenters (e.g., Martin et al.[14] and Kaiser et al.[15]) have measured the number of steps of color that can be discriminated between fully saturated hues and white. There is a pronounced minimum at green, as would be expected, since the steps of saturation available are determined here by the number of steps of luminance that can be discriminated in red and blue. (Since the amount of green stays constant, the desaturating of the green in an RGB coordinate system takes place by increasing the luminance of the other two colors.) The number of steps that could be discriminated at the minimum is approximately one-fourth of the number that could be discriminated at red and at blue, as predicted by the photon-detector model.

We may therefore draw an extremely important conclusion: the contrast required for the detection of a feature is four times higher for red and for blue (at 660 and 450 nm, respectively) than it is for green (at 555 nm). The red and blue signals in an RGB representation of a picture will

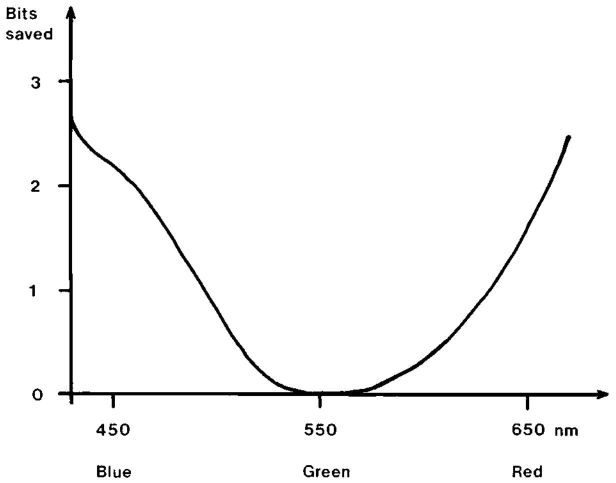

FIG. 3. Possible bit savings for image display as a function of wavelength.

therefore need 2 fewer bits (base 2 logarithm of 4) than the green signal. This conclusion is *independent* of the resolution of the device.

The dominant wavelength of the RGB primaries used has a significant effect on the number of bits required. If (as in the case for many red phosphors) the dominant wavelength of the red primary is less than 660 nm, then more bits are required for red than for blue. Figure 3 shows this effect by plotting the number of bits required relative to green as a function of wavelength, derived from the CIE curve (with Judd's 1951 modification). A designer may simply read off the bits to be saved for the primaries available.

A. More Practical Observations

If we consider a device with a resolution of 4 pels/mm (and viewed at 400 mm) we can say that under ideal viewing conditions (and viewing the most detectable image gratings) approximately 7 bits are needed for the green component of an RGB picture, but only just over 5 bits are needed for red and for blue. We will never require more bits than this, whatever the viewing conditions or the image being viewed. For a lower or higher resolution display fewer bits per pel are required.

These figures suggest that about 17 bits are needed for the optimal display of color pictures, but the requirement for color display can be reduced still further if we use the conclusions of the earlier sections of this paper. In practice, for a 4 pel/mm display, we can achieve equivalent results with just 4 bits per pel for the green plane. As just described, we may assign 2 bits less, just 2 bits, for each of red and blue. This makes the convenient total of 8 bits for high-quality practical color display.

For completeness, it must be mentioned that the graph in Fig. 3 (which is derived simply from the figures for detector efficiency) does not allow for the variation of the number of photons in light as a function of wavelength, nor for variations due to the color-matching functions of the eye (see Wyszecki and Stiles[16]). The color-matching functions may be used to determine the relative powers (and hence photon levels) of the three colors in use when white is the perceived result of the mixture.

Allowing for these factors indicates that slightly more information is required in red than would be read from Fig. 3. This difference is about 0.5 bits at a wavelength of 660 nm. The 8-bit encoding scheme described in the last paragraph is therefore slightly deficient in red. Similarly, the scheme is over-generous for blue by about the same amount. If a 12-bit encoding scheme were to be used, an almost ideal assignment of the bits (assuming we wished to keep an integral number for each color) would be 4 for red, 5 for green, and 3 for blue.

Workers in this field have already shown that by analysis of a color picture (e.g., using a Peano scan,[17] or by color space partitioning[18]) it is possible to produce a good-quality picture using 8 or fewer bits and an appropriate look-up table. This section has shown why this is so, and also that it is possible for 8 bits in a general way that always uses the same look-up table and does not require analysis of the color distribution in each image. This scheme has been tested on a wide variety of pictures using our image-processing system, with the results predicted. An example is shown in Fig. 4, though the variations due to the reproduction process must make this illustration less convincing than that on a real display.

It is important that the bits available are used intelligently. Figure 4 was generated using the error diffusion algorithm for each of the color planes. If, instead, we make no attempt to reduce errors and just use the high-order bits of each color (thresholding), then the inferior results shown in Fig. 5 are obtained.

My experiments have shown that a color picture encoded to 8 bits by simple error diffusion at this resolution provides excellent results, with no significant contouring or pel structure being visible at the standard viewing distance. Even with computer-generated images, only very slight contours are noticeable (due to deficiences in the error diffusion algorithms), and these may be eliminated by randomizing the algorithm appropriately. It has also been possible to confirm that the distribution of bits for color is at least approximately correct. If either red or blue is given the 4 bits instead of green, then the picture is noticeably inferior and pel structure becomes visible.

VI. CONCLUSIONS – WHAT THIS MEANS FOR REAL PICTURES

The preceding sections detail two important models for describing the limitations of the HVS. The frequency model describes the ability of the HVS to detect gratings or bars against a background. The target model describes the ability of the HVS to detect more regular (e.g., circular) targets against a background.

For a given display (or other output device), there is little point in providing image capability which exceeds that of

FIG. 4. Color image processed to 8 bits/pel, using 4 bits for green and only 2 bits each for red and for blue.

FIG. 5. As Fig. 4, using thresholding instead of error diffusion for each color.

the observer. The criterion that has usually been used for determining how many gray levels to provide has been the limitations of the eye as described by the frequency model (grating measurements). Yet, in real photographs, details tend to be relatively symmetric: it is very rarely necessary to detect bar-like features that are also low contrast. This almost certainly explains why nearly all images thresholded to 64 or even fewer gray levels look as good as the originals (sampled at 256 gray levels), and more accurate processing of the information allows the same pictures to be displayed with only 16 gray levels.

For transferring information to the viewer, the "target model" curve must be used as specification if all features are to be detectable regardless of their shape. In other words, for features of a given dimension, the contrast of the feature must be sufficiently high (either naturally, or after enhancement) so as to keep it below the lower of the two curves shown in Fig. 2. This again implies that, in turn, the presentation of pictures only need be to this same specification.

For color (RGB) displays, only the green signal need be to the accuracy just described, and the red and blue signals typically require 2 fewer bits. The actual savings may be determined approximately from Fig. 3 in which the number of bits to be saved is plotted as a function of the dominant wavelength of the color used for the primary.

The following new "rules of thumb" are suggested for the design of output devices used for the display of images:

• A cost-effective general-purpose display may be designed to perform no better than the target model curve (see Fig. 2). If it is to be used for basic research, or if the detection of low-contrast linear features is likely to be a significant application area, then using the frequency model curve as a design limit may be appropriate.

• For color images, fewer bits are needed for each of the red and blue planes than for the green plane. No quality improvement will be gained by using more than a total of 18 bits of intensity and color information for each pel of a color image, provided that the bits are assigned to each color appropriately. The number of bits that may be saved for red and blue is dependent on the dominant wavelength of the color used to represent these primaries, and may be read off the graph shown in Fig. 3. For example, if the dominant wavelengths are 660 and 450 nm, then 2 bits may be saved for each color. Choosing phosphors with more extreme dominant wavelengths could save even more bits (and allow a greater range of colors to be represented), but with the phosphors currently available this possibility will usually require more power and may disturb the color balance of existing images.

• A suitable number of bits for the green plane of a color image may be determined from the lower curve of Fig. 2 for a given output resolution. The bits for the other two primaries may then be deduced by subtracting the savings derived from Fig. 3. As an example, for a 4 pel/mm display we would use 4 bits for green and 2 each for red and blue. These may be conveniently assigned to a single byte.

• If the 2:4:2 scheme for red, green, and blue is used, it is recommended that green be placed in the most significant 4 bits so that images of real scenes may usually be viewed satisfactorily on a monochrome display. Red and blue, in that order, would be placed in the 4 least-significant bits. Thus,

MSB LSB

G	G	G	G	R	R	B	B

MSB = Most Significant Bit
LSB = Least Significant Bit

• If 12 bits are available for the display, then the best coding scheme (allowing for all the factors involved, and using an integer number of bits for each color) will use 5 bits for green, 4 for red, and 3 for blue:

MSB LSB

G	G	G	G	G	R	R	R	R	B	B	B

REFERENCES

[1] J. L. Mannos and D. J. Sakrison, "The Effects of a Visual Fidelity Criterion on the Encoding of Images," IEEE Trans. Information Theory **IT-20(4)**, 525-536 (July 1974).
[2] F. W. Campbell, "The Human Eye as an Optical Filter," Proceedings of the IEEE **56(6)**, 1009-1014 (June 1968).
[3] J. G. Robson, "Spatial and Temporal Contrast-Sensivity Functions of the Visual System," J. Optical Soc. Am. **56(8)**, 1141-1142 (August 1966).
[4] A. Rose, "The Sensitivity Performance of the Human Eye on an Absolute Scale," J. Optical Soc. Am. **38(2)**, 196-208 (February 1948).
[5] A. D. Schnitzler, "Analysis of Noise-Required Contrast and Modulation in Image-Detecting and Display Systems," Perception of Displayed Information, edited by L. M. Biberman (Plenum Press, New York-London, (1973), pp. 119-166.
[6] R. E. Sturm, and R. H. Morgan, "Screen Intensification Systems and their Limitations," Amer. J. of Roentgenology **62(5)**, 617-634 (November 1949).
[7] H. R. Blackwell, "Contrast Thresholds of the Human Eye," J. Optical Soc. Am. **36(11)**, 624-714 (November 1946).
[8] R. W. Floyd and L. Steinberg, "An Adaptive Algorithm for Spatial Greyscale," Proc. Soc. Inf. Display **17(2)**, 75-77 (2Q 1976).
[9] C. N. Judice, J. F. Jarvis, and W. H. Ninke, "Using Ordered Dither to Display Continuous Tone Pictures on an AC Plasma Panel," Proc. Soc. Inf. Display **15(4)**, 161-169 (4Q 1974).
[10] F. W. Campbell and K. Durden, "The Visual Display Terminal Issue: A Consideration of its Physiological, Psychological and Clinical Background," Ophthal. Physiol. Opt. **3(2)**, 175-192 (1983).
[11] L. G. Thorell, "Using Color on Displays: a Biological Perspective," Soc. Inf. Display Seminar Lecture Notes, 2, pp. 11.1-11.24 (May 1983).
[12] D. B. Judd, "Report of U.S. Secretariat Committee on Colorimetry and Artificial Daylight," CIE Proceedings, Vol. 1, Part 7, p. 11, Paris Bureau Central de la CIE, Stockholm (1951).
[13] J. J. Vos, "Colorimetric and Photometric Properties of a 2′ Fundamental Observer," Color Res. and Appl. **3(3)**, 125-128 (Fall 1978).
[14] L. C. Martin, F. L. Warburton, and W. J. Morgan, "Determination of the Sensitiveness of the Eye to Differences in the Saturation of Colours," Medical Res. Council, Special Rept. 188, London 1933.
[15] P. K. Kaiser, J. P. Comerford, and D. M. Bodinger, "Saturation of Spectral Lights," J. Optical Soc. Am. **66(8)**, 818-826 (August 1976).
[16] G. Wyszecki and W. S. Stiles, "Color Science" (John Wiley, New York, 1982).
[17] R. J. Stevens, A. F. Lehar, and F. H. Preston, "Data Ordering and Compression of Multispectral Images using the Peano Scan," IEE Int. Conf. Electronic Image Processing, IEE, London, pp. 209-213 (July 1982).
[18] Y. Tozawa, "Color Image Quantization – Population Equalization Algorithm," Proc. 29th Conference of Information Processing Society of Japan, pp. 1049-1050 (1984).