

# A Practical Approach to Calculating Luminance Contrast on a CRT\*

Blair MacIntyre and William B. Cowan

March 6, 1994

## Abstract

Luminance contrast is the basis of text legibility, and maintaining luminance contrast is essential for any colour selection algorithm. In principle, it can be calculated precisely on a sufficiently well calibrated display surface, but calibration is very expensive. Consequently, most current systems deal with contrast using heuristics. However, the usual CRT set-up puts the display surface into a state that is relatively predictable. Luminance values can be estimated based on this state, and these luminance values used to calculate contrast using the Michelson definition. This paper proposes a method for determining the contrast of coloured areas displayed on a CRT. It uses a contrast metric that is in wide use in visual psychophysics, and shows that the metric can be approximated reasonably without display measurement, as long as it is possible to assume that the CRT has been adjusted according to usual CRT set-up standards.

## 1 Introduction

*Luminance contrast* is a measure of the perceived lightness or brightness difference between two colours. For the perception of fine detail, such as in reading, contrast across edges is very important. It has been shown that legibility is well predicted by the Michelson definition:

$$C = \frac{L_{max} - L_{min}}{L_{max} + L_{min}}, \quad (1)$$

where  $L_{max}$  and  $L_{min}$  are the maximum and minimum luminances of the foreground and background colours. The contrast  $C$  ranges from 0 to 1.0 (Legge, Rubin and Luebker, 1987; Legge et al., 1990).

Contrast is important in systems that automate the selection of colours. Such systems must have an effective means for maintaining adequate contrast whenever legibility is needed. Systems in the literature tend to choose colours within a sufficiently limited scope so as to ensure legibility (Grosse, 1985; De Corte, 1986), or use heuristics that depend implicitly on the intuition of the system designer (Meier, 1988; Salomon and Chen, 1989). It would be desirable to be able to determine the contrast of two colours algorithmically for use in a more general legibility metric.

In principle, it should be possible to maintain contrast algorithmically by calibrating the display<sup>1</sup>. If the exact point-spread function of the display surface was known, it would be possible to determine the luminance of each pixel<sup>2</sup> and calculate the contrast. Unfortunately, detailed calibration is quite difficult and tends to be too expensive to be undertaken in any but research environments. Furthermore, inter-pixel interaction effects can be large (Lyons and Farrell, 1989), which has little effect on large area calibrations but is very important in the presentation of detailed material like text where contrast is very important. Thus, it is impossible to take inter-pixel non-linearities into account without very extensive calibration, and incorporating them into a practical display model is not currently possible. To avoid the impractical calibration overhead we may hope either that CRT manufacturers discover how to make display surfaces without inter-pixel non-linearities, or that simple methods for estimating them be developed.

---

\* Appeared in *ACM Transaction on Graphics* 11(4): 336-347, October, 1992.

<sup>1</sup> Calibration is the process of adjusting the display surface to a known state. When we say *calibration* in this paper, we mean calibration done to an exact colorimetric or radiometric standard. This almost always depends on measurement of the display surface, an expensive and laborious process (Cowan, 1989). What we call *set-up*, on the other hand, refers to calibration done to an imprecise visual standard (see Section 2.)

<sup>2</sup> By a pixel we denote a unit of writing to the display surface, where the writing model adds a distribution of light to whatever already exists on the surface, with the distribution centred at the pixel position. The display surface can be divided into a set of cells centred on pixels, where each cell consists of the points closer to the pixel position than to any other pixel position. Such a cell obtains most of its light intensity from the write operation to the pixel at its centre. These cells are sometimes called pixels, but should not be confused with the pixels described in this paper.

	Display A		Display B	
	Typical Light	Bright Light	Typical Light	Bright Light
Luminance of White	14.01	15.50	6.45	8.51
Luminance of Black	0.17	2.11	0.05	1.75
Relative Black Luminance	0.01	0.14	0.01	0.21
Contrast	0.98	0.76	0.98	0.66

**Table 1:** The luminance of large areas of black and white were measured on two typical displays without adjusting the brightness or contrast controls. Display A is a Sun colour monitor Model GDM-1662B. Display B is a DEC monochrome monitor Model VR 260-BB. The viewing environment is a computer lab equipped with both dim incandescent directional track lights and overhead fluorescent lights. The typical lighting condition occurs with only the dim incandescent lights turned on. The bright lighting condition occurs when the fluorescent lights are turned on as well. The typical low lighting condition is representative of many computer labs. The monitors were set up for the typical lighting condition using the standard set-up described in the main text. The contrast and brightness of the displays were not subsequently adjusted to account for the brighter light. The *relative black luminance* is the luminance of black relative to the luminance of white. The contrast was calculated using Equation 1.

This paper is based on a practical metric for determining luminance contrast which was developed while exploring a general approach to automatic colour selection for window systems (MacIntyre, 1991). The goal of the system was to choose colours for windows based on attributes of the window contents, such as the font to be used for displaying text in a terminal window and the relationship of the window to other windows on the display, without knowing the exact contents of the window. To this end, the metric was designed to provide a contrast measure for the average window contents. In this situation, a simple but imprecise metric combined with a fairly high *contrast threshold* (the contrast value used by the metric to indicate adequate contrast) produces results that are more useful than a complex, precise metric with a low threshold.

## 2 Calculating Luminance

For Equation 1 to yield useful results, a reasonable approximation to the luminance of the two colours is required. There are three distinct approximations to luminance that must be considered, corresponding to different assumptions about the model of the CRT and the viewer. First is the drive luminance ( $L_d$ ), which is calculated as a weighted sum of the gamma-corrected RGB values (Cowan, 1989), and is thus easy to control programmatically. It corresponds to a model which assumes a display surface that fits the set-up model exactly and ignores inter-pixel effects. Second is the emitted luminance ( $L_e$ ), which is the luminance of light emitted from the display surface. This model assumes an ideal viewing environment. Third is the perceived luminance ( $L_p$ ), which is the luminance of the total light that enters the eye.  $L_p$  is the quantity that matters for calculating contrast, but only  $L_d$  is easy to control from within an application. Controlling  $L_e$  requires careful calibration; controlling  $L_p$  requires attention to the visual environment and the state of the user.

While  $L_p$  is the relevant quantity, it is not possible to calculate it in practice. First, the lighting conditions in the area of the display significantly affect the perceived luminance and are unknown in typical environments. Second, most displays are equipped with two controls, a brightness control and a contrast control, which greatly affect  $L_e$  (and thus  $L_p$ ). These controls are intended to be used to set up the monitor to account for lighting conditions in the viewing area. Table 1 shows the effect of different lighting conditions on the luminance of large areas of black and white on two typical displays when the displays are not readjusted to account for a change in viewing conditions. As can be seen, increasing the ambient light increases the luminance of all colours on the display (by increasing the amount of light reflecting off the display). The result is a dramatic decrease in contrast, since both black and white are increased equally. However, when the lighting conditions change, the user is expected to readjust the brightness and contrast controls so as to maintain in the display surface the perceptual characteristics on which the brightness and contrast settings are based.

These characteristics are reflected in the standard CRT set-up used in this paper, which is performed as fol-

lows (Cowan, 1989; Benedikt, 1985). First, the focus is set by the *shrinking raster method*: a flat field is displayed and the focus adjusted to be as sharp as possible while ensuring that raster lines are invisible in the centre of the screen. Second, the *black level*, or background level, is set by adjusting the brightness knob. It is set to the brightest level consistent with zero input ( $L_d = 0$ ) producing perceived black ( $L_p = 0$ ) under the usual ambient illumination conditions. Finally, the contrast is set to the highest level that produces no blooming when the CRT is used at full input intensity. Blooming occurs when the electron guns supply too much energy to the display surface, resulting in wide, blurry pixels and desaturated colours (Cowan, 1989).

While the exact  $L_e$  values can not be determined, the relative luminance values change very little under different lighting conditions, assuming the CRT is properly set up each time the lighting conditions change. For example, the perceived luminance of black on a properly set-up CRT is roughly 1-2% of the perceived luminance of white, regardless of the actual luminances. (In fact, if the CRT is unable to provide enough intensity to produce this ratio for high ambient illumination it tends to look washed out and users attempt to find better viewing conditions.) Since contrast is a function of *relative* differences, it is therefore reasonable to ignore the effects of lighting conditions and display adjustment for the purpose of calculating contrast<sup>3</sup>.

This simplification is important, but not sufficient. It still remains to determine  $L_e$ , which can normally only be done through calibration. Fortunately, simple scaling arguments allow us to get a rough connection between  $L_e$  and  $L_d$  without doing any measurements. Recall that black, the darkest colour, corresponds to  $L_d = 0$ , and white, the brightest colour, corresponds to the maximum value of  $L_d$ . Between them, as long as reasonable gamma correction is performed,  $L_e$  scales linearly with  $L_d$  for large, homogeneous patches measured at a single location on the display. Without knowing something about the adjustment of the CRT controls, the effective gamma needed to adjust the RGB values can be in the range 1.4 to 4. But, assuming a properly set-up monitor, the gamma used to adjust the RGB values can be assigned a typical value, around 2.3.  $L_d$  is then approximated by the Y coordinate of the adjusted colour value expressed in the CIE XYZ colour space. Once the black level is taken into account, contrast calculated using  $L_d$  is exactly the same as the contrast calculated using  $L_e$ , for large areas of colour. The next section shows how to account for the black level.

The above calculations do not take into account the inter-pixel effects which are also part of  $L_e$ . There are two types of inter-pixel effects: display non-linearities and pixel broadening (bleed). Lyons and Farrell show that the non-linearities that exist in CRTs can be sizable and suggest possible approaches to account for them. Nevertheless, no standard theory exists which accounts for them and could be used in a rule of thumb calculation like the one suggested here, making this a topic for future research.

Pixel bleed implies that pixels are not sufficiently localized spatially for  $L_d$  (adjusted only by the black level) to be a good approximation for images with fine detail like text. Section 2.2 shows how to correct for the spatial extent of an average pixel. For high precision applications, this procedure requires expensive and difficult measurement, but the calculations below show that a reasonable approximation can be produced based on monitor set-up conditions. Thus,  $L_d$  with both effects included, which we call  $L'_d$ , scales linearly with  $L_e$  and can be used in place of  $L_e$  when calculating Michelson contrast.

## 2.1 Correction for Non-Zero Black Level

The black level, or background level, of the display represents the minimum luminance that can be produced on the display. The definition of black internal to the graphics system, a zero input to the CRT, does not imply zero luminance on the screen: scattered and reflected light prevents it, not to mention residual beam current that is present even at zero input.<sup>4</sup> Fortunately, the luminance of black on a well-adjusted CRT can be estimated from the set-up conditions. As mentioned above, on a properly adjusted display the luminance of black is 1-2% of the luminance of white, as is the case with both of the displays in Table 1 when properly adjusted to account for the low ambient lighting conditions.

However, when the  $L_d$  value of black is calculated, by taking the zero input to the monitor and gamma correcting it, it will be 0. To correct this, a luminance of 1-2% of the  $L_d$  value of white should be added to all  $L_d$  values before

<sup>3</sup>There are known but poorly measured effects by which the state of adaptation of the eye (which is a function of the prevailing lighting conditions) changes the luminance scaling of the response of the eye. When the CRT set-up is changed to account for different ambient lighting, changes to the contrast and brightness controls change the CRT characteristics to take this into account. But this effect is small and varies from CRT to CRT so it is ignored in this paper.

<sup>4</sup>In fact, if zero luminance black were obtainable, any colour whatsoever, even the darkest gray, would have a contrast of 1 compared to it. Then all colours would produce equally legible results when used with black, which experience shows is obviously incorrect.

they are used in the contrast calculation. Thus,  $L_d$  with the black level correction, which we shall call  $L_{db}$  is:

$$L_{db} = Y_d + B \cdot L_{Wd}, \quad (2)$$

where  $B$  is the fraction of white used to approximate the black level of the display and  $L_{Wd}$  is the  $L_d$  value of white.

## 2.2 Pixel Bleed

When a single pixel is displayed on the screen of a CRT it has a spatial profile that extends beyond the borders of the cell in which it is conceptually localized. This ‘leakage’ turns out to be the limiting factor in the ability of the CRT to deliver contrast on its screen. The discussion below is based on the shrinking raster set-up, and should be close to standard for all CRTs that are focussed using it.

The luminance contribution of any pixel to its neighbours varies from display to display, but is significant in all cases. The spatial profile of a stationary electron beam can be expressed as a Gaussian function centred at the pixel location  $(x_0, y_0)$  by<sup>5</sup>:

$$\Phi(x, y) \propto \frac{1}{2\pi\sigma^2} \exp\left(\frac{-((x - x_0)^2 + (y - y_0)^2)}{2\sigma^2}\right). \quad (3)$$

The amount a pixel influences others around it depends on the value of the constant  $\sigma$ , which represents the amount of pixel bleed. This function is normalized so that an infinite series of pixels will have unit integrated weight in a unit interval.

This equation is a good approximation to the vertical profile of a moving electron beam, but less so to the horizontal (ie. scanning direction) profile, which is affected by the performance of the video amplifiers (Mulligan and Stone, 1989). The degree of interaction between adjacent horizontal pixels varies from CRT to CRT, and involves substantial non-linearities on some CRTs (Lyons and Farrell, 1989; Naiman and Makous, 1992). But Naiman and Makous’ study, in particular, shows that while there are substantial non-linear inter-pixel interactions, the horizontal profile of an individual pixel remains qualitatively Gaussian. Thus, although the effect is very important for optimal quality font display, it seems reasonable to assume that the horizontal profile is comparable in size to the vertical one, as Naiman and Makous’ measurements seem to show for their CRT<sup>6</sup>. Among other things, any CRT that severely violates this assumption would display undesirable visual anisotropy.

The  $\sigma$  value in Equation 3 varies from display to display. Using a light meter capable of measuring the emitted luminance of individual pixels, a reasonable approximation of this value can be determined for a given display, but this procedure is both time-consuming and expensive. Instead, consider the effect on the appearance of the display of varying  $\sigma$ . Imagine displaying a flat field of equal intensity pixels. If the  $\sigma$  value is small enough so that there is no noticeable pixel bleed, it is possible to see the grid pattern of the pixels clearly on the display surface of a properly set up CRT because the spatial profile of the pixel is limited to its imaginary area, resulting in the pixels edges being noticeably darker than their centres. As  $\sigma$  increases, the pixels increase the emitted luminance at their edges, as well as contributing to the emitted luminance of their neighbours. At the optimal level of pixel bleed, the emitted luminance of the pixel edges in a uniform field is close enough to the emitted luminance of the pixel centres that the difference is imperceptible, resulting in the impression of an area of uniform emitted luminance. Increasing the pixel bleed beyond this optimal level causes the area to become even smoother (though not perceptibly), but the display loses its sharpness (Klassen, 1989). Choosing the minimal value of  $\sigma$  for which no further improvement is possible is the essence of the shrinking raster set-up condition, and it defines a specific value for  $\sigma$ .

Thus, while measuring individual displays may be impractical, it is possible to get a reasonable estimate for  $\sigma$  for a well set up CRT. Klassen suggests that a value of  $\sigma/a = .51 \pm .01$  is optimal, where  $a$  is the distance between pixel centres. By using  $\sigma = a/2$ , Equation 3 becomes:

$$\Phi(x, y) \propto \frac{2}{\pi a^2} \exp\left(\frac{-2((x - x_0)^2 + (y - y_0)^2)}{a^2}\right). \quad (4)$$

Using Equation 4, it is possible to calculate the contribution that a pixel’s neighbours make to its emitted luminance. Consider a pixel **P** with all its neighbours the same intensity. Figure 1 shows the pixel neighbours and the amount each contributes to the intensity of pixel **P**. 47% of the luminance of the area is contributed by the pixel itself. The

<sup>5</sup>This formula is actually a simplification. It is complicated for colour displays by the existence of a shadow mask.

<sup>6</sup>This does not address the non-linear spatial interactions between neighbouring horizontal pixels. As mentioned on page 2, while they may be large, they are not sufficiently well understood and are therefore ignored in this paper.

E	D	C	D	E
D	B	A	B	D
C	A	<b>P</b>	A	C
D	B	A	B	D
E	D	C	D	E

**Figure 1:** The centres of the immediate horizontal and vertical neighbours of pixel **P** (labeled A) are a distance of  $a$  from the centre of **P** and each contributes 10.7% of the luminance of **P**. The centres of the immediate diagonal neighbours (labeled B) are a distance of  $a\sqrt{2}$  away and contribute 2.5% each. The centres of the more distant neighbours are a distance of  $2a$  (labeled C),  $a\sqrt{5}$  (labeled D) and  $2a\sqrt{2}$  away (labeled E). Pixels C through E together contribute less than 0.1% of the luminance of **P**.

remaining 53% is contributed almost completely by the immediate neighbours, neighbours C through E contributing a negligible amount. Given the insignificant contribution of the more distant neighbours, only the immediate neighbours need be considered to calculate the intensity of a pixel.

In principle, it is possible to calculate the intensity of all relevant pixels. Pappas and Neuhoff (1991) use such an approach when halftoning images for output on a laser printer. They model the interaction of halftone dots on the printed page and use this information to determine how dark any given spot will actually appear. However, their goal is to model accurately the display device (a laser printer) and exploit this model to increase the quality of halftoned images. The goal of this paper, on the other hand, is to create an approximate model of the output device (a CRT) and use this model to estimate the contrast of a typical image. There is no desire to accurately calculate exact contrast values for specific pixels on the display. Systems such as (MacIntyre, 1991), which generate colours automatically, must choose colours before they are used, so a contrast metric is required which provides a single number that estimates the performance of the whole screen. Therefore, instead of examining a particular screen image, we determine the average adjacency of foreground and background pixels (see Section 3) and then use these values to calculate average luminance values for foreground and background pixels.  $L'_d$  of the average background pixel is approximated by:

$$L'_{dbg} = \xi L_{dbg} + (1 - \xi)(L_{dbg} N_{bg} + L_{dbfg}(1 - N_{bg})), \quad (5)$$


where  $L_{dbfg}$  and  $L_{dbg}$  represent the  $L_{db}$  luminance (from Equation 2) of the foreground and background colours.  $N_{bg}$  represents the average number of background pixels adjacent to a background pixel. The contribution of a pixel to its own total luminance,  $\xi$ , would typically use the value of 0.47 calculated above.  $L'_d$  of the average foreground pixel can be calculated similarly:

$$L'_{dfg} = \xi L_{dfg} + (1 - \xi)(L_{dbg}(1 - N_{fg}) + L_{dfg} N_{fg}), \quad (6)$$

where  $N_{fg}$  represents the average number of foreground pixels adjacent to a foreground pixel.

Except in the rare case where the contents of the display are known ahead of time and remain fixed, the adjacency values are most reasonably approximated using averages such as those computed in the following section. Three important observations should be made. First,  $N_{fg}$  is not necessarily equal to  $N_{bg}$ . Second, when calculating the average adjacency, only pixels that are at a luminance edge should be considered because they are the ones that determine the contrast at that edge.

Thirdly, this section is primarily of concern for rendered objects whose image is only a few pixels thick, as pixel bleed does not have as significant an effect on thick objects, where more than 2 pixels is considered thick. Since only a pixels immediate neighbours contribute significantly to the emitted luminance of its cell, the maximum luminance of an object that is many pixels thick is controlled by pixels in its uniform central region. Thus, while the edges of thick objects are blurred by pixel bleed, their maximum luminance is essentially equal to their  $L_{db}$  luminance.



**Figure 2:** The average number of foreground (black) pixels around any foreground edge pixel in a font varies widely between fonts. Over 74% of the foreground edge pixels in font (a) have only one or two foreground neighbours and less than 5% have four or more neighbours, with an average of 2.3 neighbours. Conversely, over 60% of the foreground pixels in font (b) have five or more neighbours and less than 8% have two or less, with an average of 4.4 neighbours. In both cases, most background pixels that are adjacent to at least one foreground pixel have 3 or less foreground neighbours. Only 9% of the background pixels in font (a) have 4 or more foreground neighbours, with 2.4 on average. Similarly, 16% of the background pixels in font (b) have 4 or more foreground neighbours, with 2.6 on average.

### 3 Guidelines for Determining Pixel Adjacency

In this section, guidelines for determining the average pixel adjacency values  $N_{fg}$  and  $N_{bg}$  are discussed. The *average pixel adjacency* is the average number of background or foreground pixels which are immediate neighbours of a typical background or foreground pixel, respectively. The approaches suggested here are based on those developed in (MacIntyre, 1991) and produced acceptable results in that context. As discussed in the previous section, the more that is known about the context in which the colours will be used, the more accurate the estimate of average pixel adjacency that may be calculated. Text-based applications are considered first, followed by graphics-based applications.

#### 3.1 Text Based Applications

For text based applications, two observations can be made. First, the average neighbour of a foreground pixel is largely determined by the font being used. Second, the average neighbour of a background pixel is largely independent of the font being used.

Consider the two fonts shown in Figure 2, rendered in black text on a white background. The white background pixels bleed into any adjoining black foreground pixels, increasing their emitted luminance and therefore lowering contrast. The average foreground pixel of font (a) (the thinner font) has almost six background neighbours, whereas the average foreground pixel in font (b) (the thicker font) has just over three. As a result, the foreground pixels of font (a) have their emitted luminance increased by roughly twice as much as the foreground pixels of font (b), resulting in a lower contrast for font (a) than for font (b).

Conversely, consider what happens when these fonts are rendered in white text on a black background. The white foreground pixels bleed into the adjoining black background pixels, increasing their emitted luminance and therefore lowering contrast. The average foreground pixel of font (a) has two foreground neighbours, whereas the average foreground pixel of font (b) has approximately four. As a result, the foreground pixels of font (a) have their emitted luminance increased by roughly half as much as the foreground pixels of font (b), resulting in a lower contrast for font (a) than for font (b).

The effect in both cases is the same. Font (a) has less contrast than font (b) when displayed with the same colours. While it is conceivable that the average pixel adjacency relationships could be calculated for each font used, fonts can more usefully be divided into a few categories such as *thick* and *thin*. Thick fonts are those whose foreground pixels are adjacent to many other foreground pixels, similar to font (b) in Figure 2. Font (a), on the other hand, is representative of thin fonts; most of the pixels adjacent to its foreground pixels are background pixels.

Note that a font's "thickness" is dependant on its bitmap representation, and is not intended here to encompass its typographic characteristics, such as its *weight* or *spacing*, which may also affect a font's legibility (see (Rubinstein, 1988) for a discussion of digital typography). It should also be kept in mind that the effect discussed here is device-dependent. It would be natural, for example, for the scan conversion software driving a CRT to take the weight of a font

Pixel Adjacency	Colour Orientation	Black Level								
		.01			.025			.05		
		$L'_{d_{bg}}$	$L'_{d_{fg}}$	$C$	$L'_{d_{bg}}$	$L'_{d_{fg}}$	$C$	$L'_{d_{bg}}$	$L'_{d_{fg}}$	$C$
Thick Font	Black on White	0.84	0.25	0.54	0.85	0.26	0.53	0.88	0.29	0.50
	White on Black	0.18	0.77	0.61	0.20	0.79	0.59	0.22	0.81	0.57
Thin Font	Black on White	0.85	0.39	0.38	0.87	0.40	0.37	0.89	0.43	0.35
	White on Black	0.17	0.63	0.58	0.18	0.65	0.56	0.21	0.67	0.53
Single Pixel	Black on White	0.95	0.54	0.27	0.96	0.56	0.27	0.99	0.58	0.26
	White on Black	0.07	0.48	0.73	0.09	0.49	0.70	0.11	0.52	0.64
Large Area	Black on White	1.01	0.01	0.98	1.02	0.03	0.95	1.05	0.05	0.91
	White on Black	0.01	1.01	0.98	0.03	1.02	0.95	0.05	1.05	0.91

**Table 2:** The effects of different black levels and pixel adjacency relationships are apparent. Changing the black level is insignificant compared to the effect of changing the pixel adjacency relationship. The thick and thin fonts are those shown in Figure 2. The contrast was calculated using the values suggested in the text. The *single pixel* values represent the calculated contrast for a single foreground pixel completely surrounded by background pixels. The *large area* values represent the maximum possible contrast, obtained by ignoring the effects of pixel bleed for large areas of uniform colour.

into account so that larger characters do not appear to have a different luminance contrast than smaller ones. Similarly, these considerations are taken into account intuitively when a font is hand tuned for a CRT of a given resolution or when fonts or scan conversion algorithms are hand-tuned for printers of specific resolution.

### 3.2 Graphics-Based Applications

In a graphics based application, the interaction of the foreground and background is much harder to determine without significant work because the relatively simple pattern found with text is not as likely to exist. One approach when no pattern can be found is to use the worst case situation, in which eight pixels of the lighter colour surround each pixel of the darker colour, significantly lightening the darker colour. While this situation may be extremely unlikely, depending on the application, this assumption provides a lower bound on  $L'_d$  and, therefore, contrast. Additionally, it is likely to be the case in which diminished contrast is most visible.

As with text based applications, a more reasonable estimation is possible if something is known about the objects rendered. For example, consider the two fonts shown in Figure 2 as simple bi-level graphics. Font (a) corresponds to a drawing consisting of single pixel width objects. Slightly thicker objects would correspond to Font (b).

Table 2 shows the contrast calculated for black and white foregrounds and backgrounds with an optimal  $\xi$  value of 0.47. As can be seen, the effects of using different black levels is not very significant compared to the effect of using different pixel adjacency values.

## 4 Contrast Metric

By using the average adjacency relationship and the optimal  $\xi$  value, a reasonable metric for measuring luminance contrast can be determined by setting  $L_{max}$  and  $L_{min}$  in Equation 1 to the appropriate one of  $L'_{d_{bg}}$  and  $L'_{d_{fg}}$  from Equations 5 and 6, as follows

$$L_{max} = \max \left( L'_{d_{bg}}, L'_{d_{fg}} \right), \quad (7)$$

$$L_{min} = \min \left( L'_{d_{bg}}, L'_{d_{fg}} \right). \quad (8)$$

In practice, this metric could be used in two ways. First, given two colours and some information about how they will be used (ie. background and foreground colours for displaying text using a certain font) an estimate of the contrast between these colours could be produced. To accomplish this, first transform the colours into the CIE XYZ colour

space (Foley et al., 1990). Next, use the Y coordinates of the XYZ triples as the  $L_d$  values in Equation 2. Equations 1, 5, 6 and 8 can then be used to determine the contrast for the two colours.

A slightly different use of this metric is to determine, given a colour and a desired contrast threshold, the minimum (or maximum)  $L_d$  value a second colour would need so that the pair of colours, when used together, exceed the specified threshold. This is accomplished by rearranging the equations and solving for either  $L'_{d_{bg}}$  or  $L'_{d_{fg}}$ , as desired.

In (MacIntyre, 1991), a contrast of 30% was used as a legibility threshold. This number was arrived at by informal testing of a relatively small number of people, and represented a calculated contrast which allowed text to be read comfortably across a wide range of colours and fonts. The formal verification of this metric is a subject of future investigation.

## 5 Conclusion

By assuming that a CRT has been set up in a standard way, both the lighting conditions and the display controls can be effectively ignored since the display controls are adjusted specifically to account for the lighting conditions and put the display surface in a state that is relatively perceptually predictable. A set of simple equations are provided to calculate the contrast relationships based on this state. The equations are developed from simple heuristics that account for the most important factors affecting the perceived contrast on a CRT display, the black level and pixel bleed.

These equations provide a fast and reasonable estimate of luminance contrast which can be used when pairs of colours must be tested for legibility on the fly.

## 6 Acknowledgments

This work was supported by the Natural Science and Engineering Research Council of Canada through a post-graduate scholarship and grants to the University of Waterloo Computer Graphics Laboratory; the Ontario Information Technology Research Centre through grants to the laboratory; Digital Equipment Canada through equipment donations; Apple Computer Corporation through grants to the laboratory; Sun Microsystems through software support; Steven Feiner and the Dept. of Computer Science at Columbia University for support.

## REFERENCES

- Benedikt, F. (1985). A tutorial on color monitor alignment and television viewing conditions. Technical report, Canadian Broadcasting Corporation, Montreal. Development Report 6272.
- Cowan, W. B. (1989). Colorimetric Properties of Video Displays. Notes for Course 25 at the Annual Meeting of the Optical Society of America.
- De Corte, W. (1986). Finding appropriate colors for color displays. *Color Research and Application*, 11(1):56–61.
- Foley, J. D., van Dam, A., Feiner, S. K., and Hughes, J. F. (1990). *Computer Graphics: Principles and Practice*. The Systems Programming Series. Addison-Wesley Publishing Company, New York, 2nd edition.
- Grosse, E. (1985). Automatic Choice of Colors for Level Plots. Technical report, AT&T Bell Laboratories. Numerical Analysis Manuscript 85-1.
- Klassen, R. V. (1989). *Device Independent Image Construction for Computer Graphics*. PhD thesis, University of Waterloo. Available as Research Report CS-91-19 from the Department of Computer Science.
- Legge, G. E., Parish, D. H., Luebker, A., and Wurm, L. H. (1990). Psychophysics of reading. XI—Comparing color contrast and luminance contrast. *Journal of the Optical Society of America*, 7(10):2002–2010.
- Legge, G. E., Rubin, G. S., and Luebker, A. (1987). Psychophysics of reading—V. The role of contrast in normal vision. *Vision Research*, 27(7):1165–1177.



- Lyons, N. P. and Farrell, J. E. (1989). Linear systems analysis of CRT displays. *Society for Information Display 1989 Digest*, 20:220–223.
- MacIntyre, B. (1991). A constraint-based approach to dynamic colour management for windowing interfaces. Master's thesis, University of Waterloo. Available as Research Report CS-91-55 from the Department of Computer Science.
- Meier, B. J. (1988). ACE: A color expert system for user interface design. In *Proceedings of the ACM SIGGRAPH Symposium on User Interface Software*, pages 117–128.
- Mulligan, J. B. and Stone, L. S. (1989). Halftoning methods for the generation of motion stimuli. *Journal of the Optical Society of America A*, 6(8):1217–1227.
- Naiman, A. C. and Makous, W. (1992). Spatial non-linearities of grayscale CRT pixels. In *Proc. SPIE: Human Vision, Visual Processing, and Digital Display III*, volume 1666. In Press.
- Pappas, T. N. and Neuhoff, D. L. (1991). Model-based halftoning. *Proc. SPIE: Human Vision, Visual Processing, and Digital Display II*, 1453:244–255.
- Rubinstein, R. (1988). *Digital Typography: An Introduction to Type and Composition for Computer System Design*. Addison-Wesley.
- Salomon, G. B. and Chen, J. (1989). Using neural nets to aid color selection. *Society for Information Display 1989 Digest*, 20:326–329.