Spatial non-linearities of grayscale CRT pixels

Avi C. Naiman[†]
Walter Makous[†‡]

[†] Center for Visual Science
[‡] Psychology Department
University of Rochester
Rochester, New York, 14627

Abstract

We measured the overall luminance, the luminance profile of individual pixels, and the interactions among neighboring pixels of a grayscale (i.e., non-color) CRT, using a two-dimensional CCD camera. We find that the luminance of a pixel depends on the luminance of the two preceding pixels in the raster, and that the interaction results in superadditivity. Furthermore, the luminance of a pixel depends on the proportion of the screen illuminated. Two consequences of these non-linearities are that average luminance (as in a halftone image) cannot be predicted from the linear summation of individual pixel spread functions, and that inverting the polarity of a display does not simply invert the luminance profiles. These interactions must be taken into account wherever the luminance profile of displayed stimuli is important.

Introduction

To display stimuli on a raster-scan system, digital images must be quantized in space, time, and intensity. The ideal display would have no perceivable spatial grain or flicker, and a full gamut of colors. In reality, the designer is faced with cost-related tradeoffs and technical limitations in determining the amount of resolution to provide, the temporal frequency at which information is to be updated, and the number of discrete colors that should be available.

A digital *frame buffer* stores an $x \times y$ array of n-bit pixel values corresponding to discrete regions to be intensified on an optical display such as a *cathode ray tube* (CRT) (see [Conrac 85]). These values are sequentially output, in a raster-scan fashion, to a *digital-to-analog converter* (DAC), which generates a video signal to control the device's electron gun. A beam of electrons is directed, by magnetic or electrostatic deflection, across the surface of the display, bombarding a phosphorescent material. Some of the power is converted to light by the phosphor (the remaining becomes heat), the spectral power distribution being a function of the chemical composition of the phosphor. A particular region on the display is stimulated in proportion to the number of electrons emanating from the gun at the time the beam is aimed at that location.

Although there may be a direct relationship between the display signal coming to a DAC and the voltage applied to the CRT, the luminous output of the phosphor is not directly proportional to either of these values. In fact, luminance levels on the display are related in a monotonic, but non-linear, fashion to pixel values. Unless a device has been linearized in luminance, it will not have equal intensity changes across the available range; for example, an intended doubling of the luminance may actually result in an increase by a factor of 1.5 or 3. Because unit DAC values correspond to different-sized luminance steps, the relative increase in luminance for a unit increase in DAC value depends on the luminance at which the increase is made (see [Murch and Weiman 90]).

In the television industry, the practice is to correct for this non-linearity by compensating with the inverse function at the signal's source, namely the television camera. A convention has been adopted in computer graphics to approximate the non-linear relationship with a power function, $L=V^{\gamma}$, where L is luminance and V is voltage. Compensation for this non-linearity is often called *gamma correction*, and is achieved by taking the inverse of the power function. Unfortunately, many monitors are poorly characterized by even the best-fitting gamma function, with some values off by as much as a factor of two [Cowan 83].

The non-linear luminance relationship can be determined accurately by measuring the screen intensity for a representative sample of DAC values over its full range. Measurements can be made with any light meter the output of which is proportional to the power of the photons stimulating it. It is common for a light meter to be equipped with a broadband spectral filter that approximates the human photopic sensitivity function; the light meter is then called a photometer. A second- or third-order function can be fit to these data and inverted to generate a compensation table for the non-linearity. This compensation table can then be incorporated into the software used to generate frame buffer pixels, or stored in a *lookup table* (LUT) interposed between the frame buffer and the circuitry that drives the DAC, where the computed intensity is used as an index into the table and the stored value is used as the input to the DAC. In fact, it is common for manufacturers to include at least one such lookup table in hardware.

In addition to the *point* non-linearities discussed above, the luminance of CRTs may exhibit various *spatial* non-linearities as well. For example, since no constraints are placed on the relationship between adjacent frame buffer pixels, either in the x or y direction, the digital hardware can specify arbitrary waveforms. However, even if a signal from adjacent pixels at full extremes of luminance can be generated by the sample-and-hold DAC circuitry, it must subsequently pass through other physical components (amplifiers, encoders, cables) that have a limited slew rate. Therefore, there are time limitations on the change of the analog signal on the CRT display; often the signal is severely distorted when it changes abruptly between extremes [Lyons and Farrell 89]. This reflects inadequate video bandwidth, with the video circuitry acting as a spatial low-pass filter in the scanning direction [Mulligan and Stone 89].

Throughout computer graphics, models of display devices typically presume that each displayed pixel is square (or perhaps rectangular), and measure the accuracy of a displayed image with metrics based on pixel center sample points. In reality, displays have rounded, overlapping pixels, with the *intensity distribution* – or *point spread function* – decreasing monotonically from the center of the pixel. Treating a video image as a matrix of points and ignoring the influence of pixel size and shape, as is frequently done, is not always satisfactory. Pixels of standard size, from 20 to 40 pixels/cm, viewed from typical distances of 30 to 45 cm, subtend a visual angle from 1.9 to 5.7 arc minutes. The geometry of pixels of this size can significantly affect image appearance.

The purpose of these observations is to measure non-linear spatial interactions that affect the luminance and shape of the pixels of a representative grayscale display.

Equipment & Methods

We examined the luminance and spatial non-linearities of a representative grayscale CRT (i.e., *not* a color display), the Apple Two-Page Monochrome Monitor. (We chose this display system merely out of convenience; we suspect that the non-linearities reported below are present in most CRT display systems.) The display area of the monitor was 38 cm wide and 29 cm high, with 870 raster lines and 1152 pixels along each raster, resulting in approximately 3 pixels per mm both horizontally and vertically. The monitor's phosphor was EIA Type P4 (white), its video signal was RS-343 standard, and its refresh rate was 75 Hz. As per specifications, the monitor was warmed up for at least 30 minutes before calibration and measurement sessions. The brightness knob was set to the highest level that did not noticeably brighten a black screen; the contrast knob was set at maximum.

Measurements of the luminance of a circular area at the center of the screen were made with a Minolta Chroma Meter CS-100. The 1° measuring spot had a 25-mm diameter at the surface of the monitor. The luminance at the center of a full screen grew in proportion to the 1.45th power of the voltage of the electron gun (Figure 1), and measured 65 cd/m^2 when driven to its maximum. The luminance non-linearity (as derived from measurements at the center of a fully-illuminated screen) was inverted and loaded into the system's LUT, and used throughout the remainder of the measurements.

Spatial profiles were measured with a Photometrics Series 200 CCD Camera System, incorporating a two-dimensional Kodak KAF-1400 CCD array. The camera made measurements at -40°C. The central 256×256 CCD elements (out of a possible 1320×1035) were used, without binning; each CCD element measured 6.8 μm on a side. The system gain was unity, and 12-bit measurements were recorded from each CCD element. The camera head was mounted on a three-dimensional carrier, the lens was adjusted for closest focus, and the instrument was advanced to the position where the display was in best focus at the center of the screen. This procedure resulted in 74 CCD elements per CRT millimeter, or about 25 samples per pixel, both horizontally and vertically. Use of the CCD camera to observe an opaque edge attached to the face of the screen showed that any movement of the camera relative to the screen during 1 minute did not exceed 10% of the width of a CCD element. We compared measurements made with the CCD and the photometer under identical conditions to confirm that they yielded identical results.

Figure 2 shows the two-dimensional profile of a single pixel. However, the interactions among neighboring pixels that interest us here are more easily seen in one-dimensional cross-sections at appropriate orientations. Hence, we will, in general, show only one-dimensional profiles below. Unless otherwise noted, the profiles all pass through the centroid of the pixel or pixels under study.

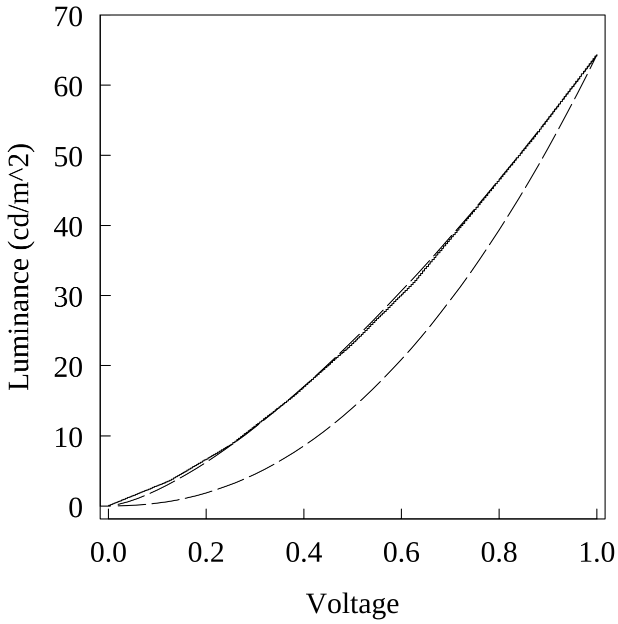

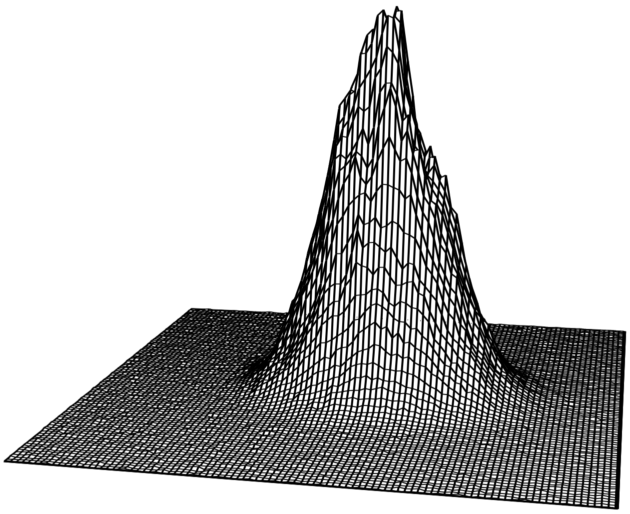

Figures 1: Luminance Non-Linearity (left) & 2: Two-Dimensional White Pixel (right). The solid curve on the left plots the luminance of the Apple Two-Page Monochrome Monitor as a function of voltage. A voltage of half maximum yielded a luminance of 22 cd/m^2, or 0.36 of maximum. The dashed curve straddling the solid curve is a power function with an exponent of 1.45. For comparison, the lower dashed curve plots a power function with the NTSC-recommended exponent of 2.2. On the right is the two-dimensional luminous-intensity profile of a white pixel on a black background. There are 96×96 CCD samples in this height field.

The exposure period of the CCD camera was not synchronized with the raster scan; consequently, most exposures began and ended while the raster scan was in mid-flight, rather than during vertical retrace. The resulting error is inversely proportional to the number of scans during the exposure, and so we used 0.25-second exposures, the longest that did not saturate the CCD elements. Thus, each exposure is the sum of 19 refreshes, and any error introduced by partial scans is less than 6%.

To check the stability of the display and the CCD camera, we repeated 14 identical measurements of a single pixel during a 5 minute period (using quarter-second exposures). The solid line in Figure 3 shows the means of the results, and the vertical lines represent ±1 standard deviation at each point. Note that a relative luminous intensity of 1.0 corresponds to the luminance generated by a full screen, but the ‘luminance’ of a single pixel is considerably less than that of a homogeneously-illuminated field, a topic taken up below.

Irregularities in the pixel profiles on a scale much smaller than the pixel are visible in these repeated observations. Both the amplitude of the irregularities and the variability of the observations are directly proportional to the signal. All observed profiles show such irregularities. Although the irregularities differ for different pixels and are greater than those observed on a homogeneous field supplied by an integrating sphere, we are not certain whether they are properties of the pixels themselves or of the measuring instrument.

Pixel Shape as a Function of Luminance

The profile in an undeflected electron beam is usually radially-symmetrical and nearly Gaussian. However, when the electron beam strikes the phosphor, the light output distribution differs. Furthermore, beam deflection, magnetic fields, power supply ripple, and vibration may all influence the shape of the light distribution. Although many of these factors can be eliminated to characterize the CRT on its own, they need to be considered as part of the entire display system when determining the overall quality of displayed images.

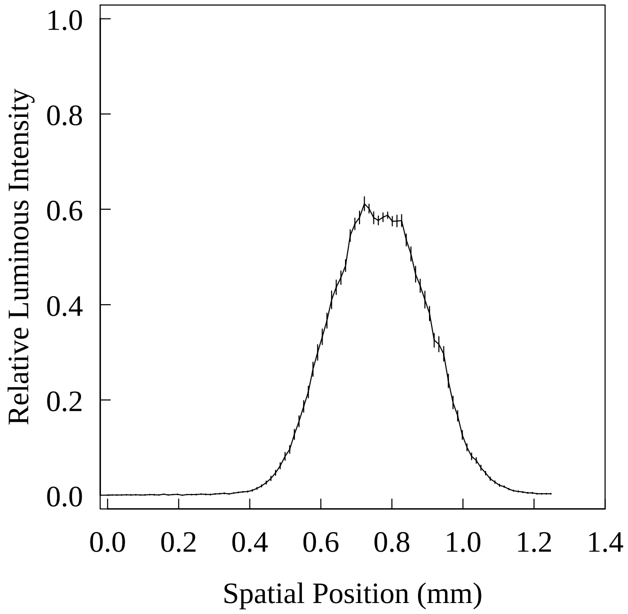

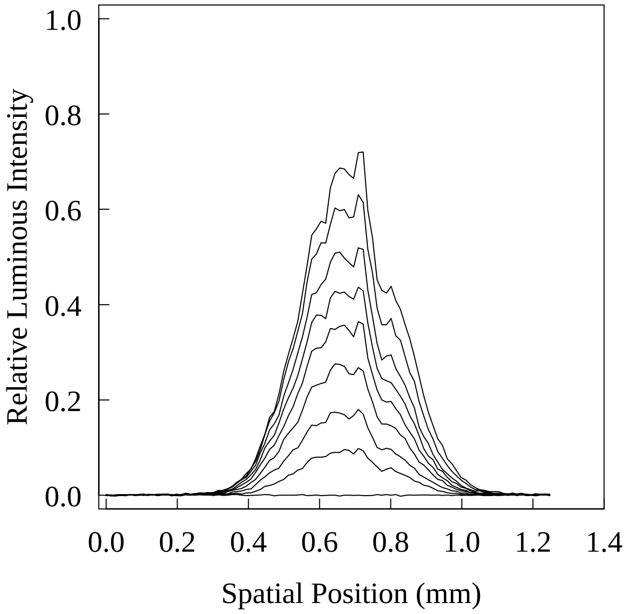

Figures 3: Measurement Variability (left) & 4: Horizontal Profiles of Light Pixels (right). On the left, the mean and ±1 standard deviations are shown for 14 measurements of a single pixel over a 5 minute period (using quarter-second exposures). Note that, in this case (and for most of the figures below), there are 96 CCD elements represented along the x-axis and the y-axis represents relative luminous intensity. The normalization factor used is the mean luminance of a full screen of white pixels. The curves on the right are the horizontal profiles through the centroid of a single pixel presented at 9 nominal luminances, ranging from 0 to 1 in steps of 1⁄8.

It is of interest whether varying the luminance of a pixel changes its shape. Figure 4 shows the horizontal profiles of light pixels of varying luminance against a black background, and Figure 5 shows that they have the same shape after being scaled by their nominal luminances, confirming what others have found (e.g., [Lyons and Farrell 89]). Figure 6 shows that the vertical profiles also have the same shape after being scaled by their nominal luminances. However, the pixels we measured are not radially-symmetric, as is demonstrated in Figure 7, which shows the horizontal and vertical profiles through a white pixel's centroid. Fortunately, though, the horizontal profiles are independent of their vertical position within the pixel (Figures 8 and 9).

Figure 10 shows the two-dimensional profile of a black pixel on a white background. The shape of black pixels is complicated by the fact that they constitute depressions embedded in bright ridges formed by the scan lines and dark furrows interposed between scan lines. The noisiness of the high luminances that represent the effective baselines for black pixels further obscures their structure. For the sake of clarity, we have clipped the top of the plot at a level just above the bottom of the furrows. To understand the complexities of this two-dimensional plot, it is helpful to examine various cross-sections of the black pixel.

Figure 11 shows vertical profiles along the horizontal centroid of dark pixels of varying luminance against a white background. It is evident that increasing the darkness of the pixel (starting from white) at first merely reduces the height of the white ridge; eventually, the darkest pixels entirely invert the raised ridge. However, the total inversion occurs only near the center of the pixel. Figure 12 shows, for a black pixel, various vertical profiles, beginning from the horizontal centroid of the pixel and moving out toward one side of the pixel. As one moves from the center of the pixel along the scan line, the luminous intensity changes from black to white, and the profile follows a course analogous to that shown in the preceding figure.

Figure 13 shows various horizontal profiles of a black pixel, beginning from the vertical centroid of the pixel and moving out toward one side of the pixel. These profiles cross one another because those passing through the blackest region at the center ride on the top of a base provided by a white ridge, and those near the edge of the scan line ride on a base nearer the dark furrows between scan lines.

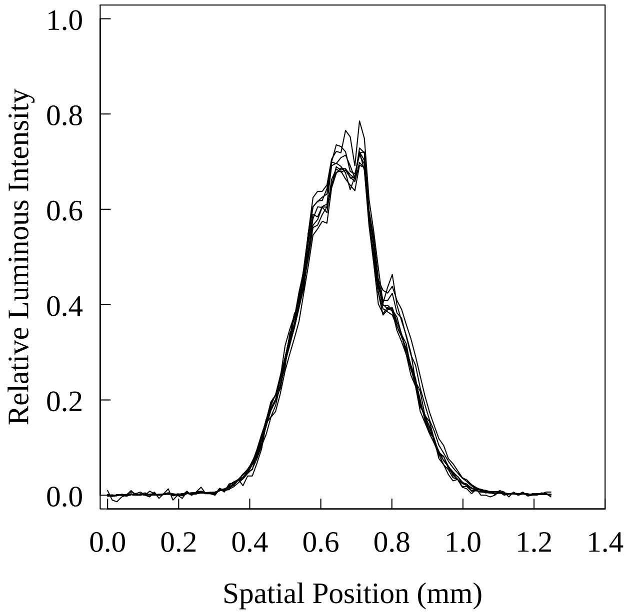

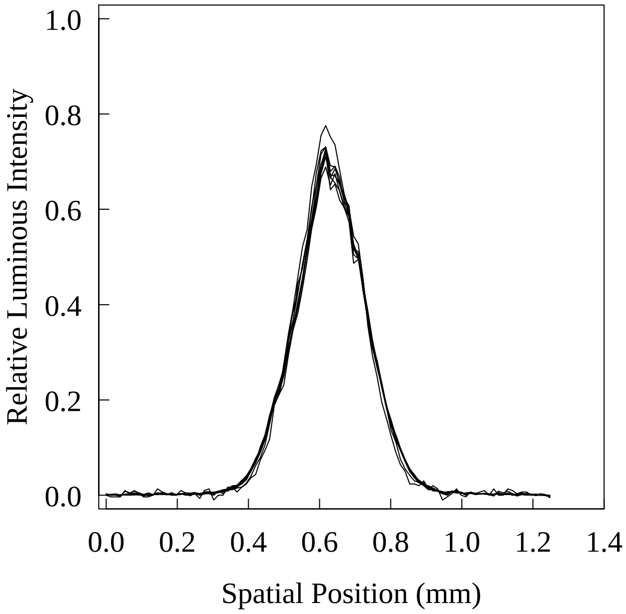

Figures 5: Normalized Horizontal Profiles of Light Pixels (left) & 6: Normalized Vertical Profiles of Light Pixels (right). On the left, the curves from the previous figure have been scaled by their nominal luminances, demonstrating that the profiles do not change shape at different luminance settings. Analogously, vertical profiles through the centroid of a single pixel presented at various nominal luminances have been normalized in the right figure, demonstrating that these profiles also do not change shape at different luminance settings.

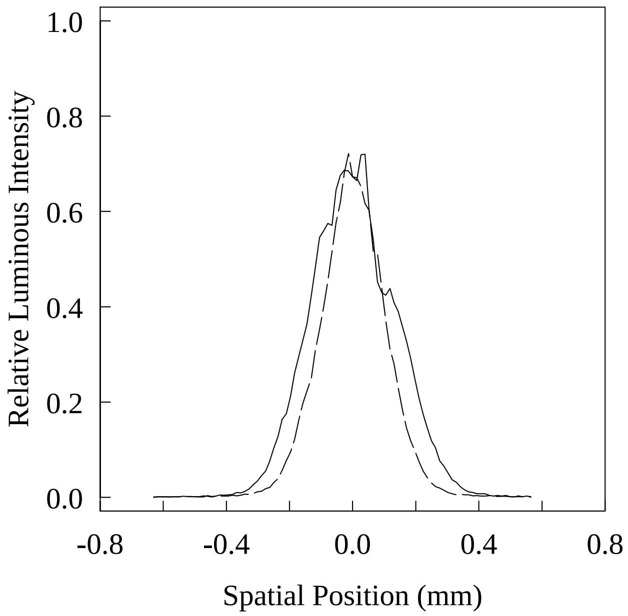

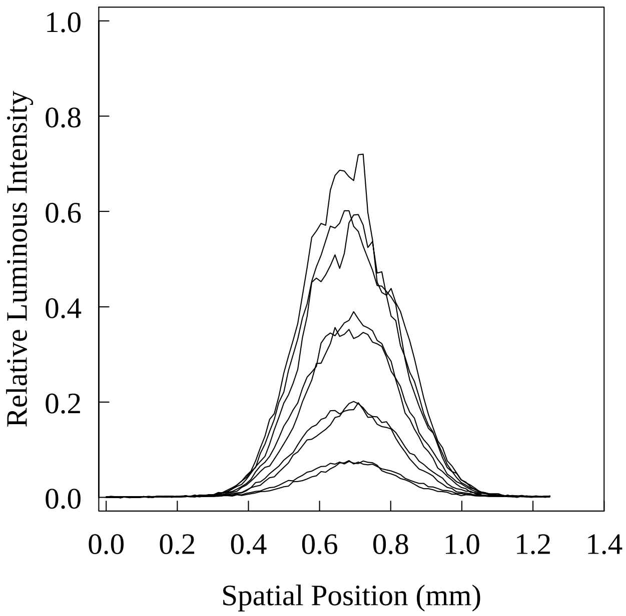

Figures 7: Horizontal vs. Vertical Profiles of White Pixel (left) & 8: Horizontal Profiles of a White Pixel (right). On the left, the horizontal (solid) and vertical (dashed) profiles through the centroid of a white pixel have been superimposed, demonstrating that the pixel's two-dimensional profile is not radially-symmetric. The top curve on the right is the horizontal profile of a white pixel taken through the vertical centroid of the pixel. The pairs of curves below that one represent equi-distant horizontal profiles above and below the vertical centroid at distances of 0.0525, 0.105, 0.1575, and 0.21 mm.

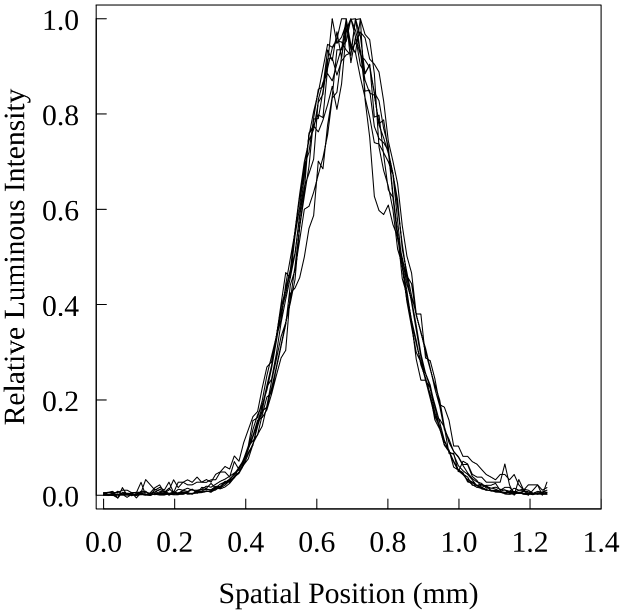

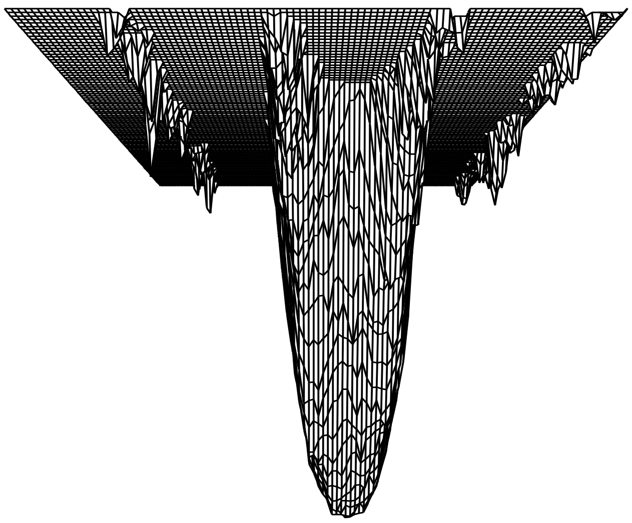

Figures 9: Normalized Horizontal Profiles of a White Pixel (left) & 10: Two-Dimensional Black Pixel (right). The profiles from the previous figure have been scaled on the left to demonstrate that a white pixel's horizontal profile is independent of its vertical position within the pixel. Note that, in this case, the curves have been normalized to the height of the horizontal profile through the pixel's vertical centroid (i.e., the highest curve in the previous figure). On the right is the two-dimensional luminous-intensity profile of a black pixel on a white background, peering down the scan lines of the display. When viewed as a surface, the black pixel serves as a depression in a field of white; for the sake of clarity, we have clipped the image just below the bottom of the white background. The two rows of downward-pointing spikes on the far left and far right of the image are the (relatively) dark furrows between bright scan lines. Similarly, there are two of these furrows flanking the black pixel. Note the horseshoe-shaped intersection between the pixel and the clipping plane, discussed in the text. There are 96×96 CCD samples in this height field.

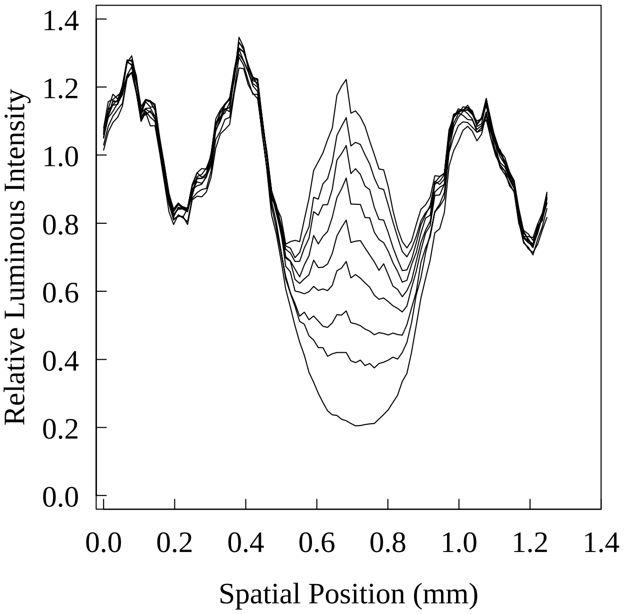

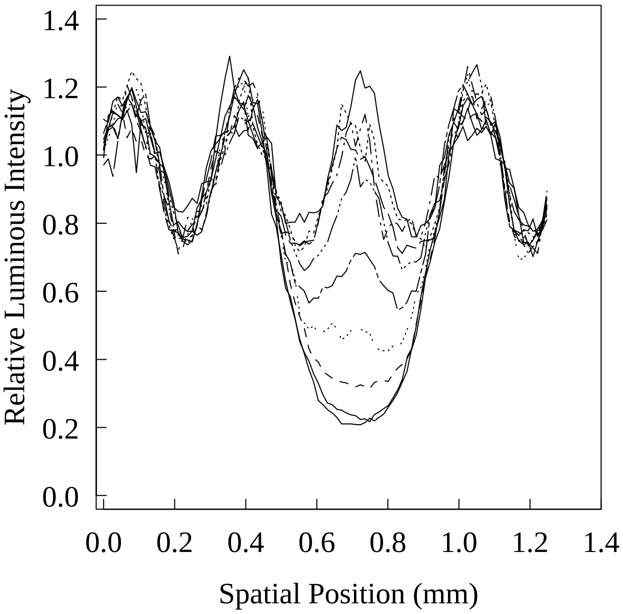

Figures 11: Vertical Profiles of Dark Pixels (left) & 12: Vertical Profiles of a Black Pixel (right). The curves on the left are the vertical profiles through the centroid of a single pixel presented at 9 nominal luminances, ranging from 0 to 1 in steps of 1⁄8, on a white background. Note the inversion of the white ridge as the pixel is darkened. The bottom (solid) curve on the right is the vertical profile of a black pixel taken through the horizontal centroid of the pixel. The other curves represent vertical profiles of the pixel at distances of 0.0525 to 0.4725 mm from the horizontal centroid, in steps of 0.0525 mm.

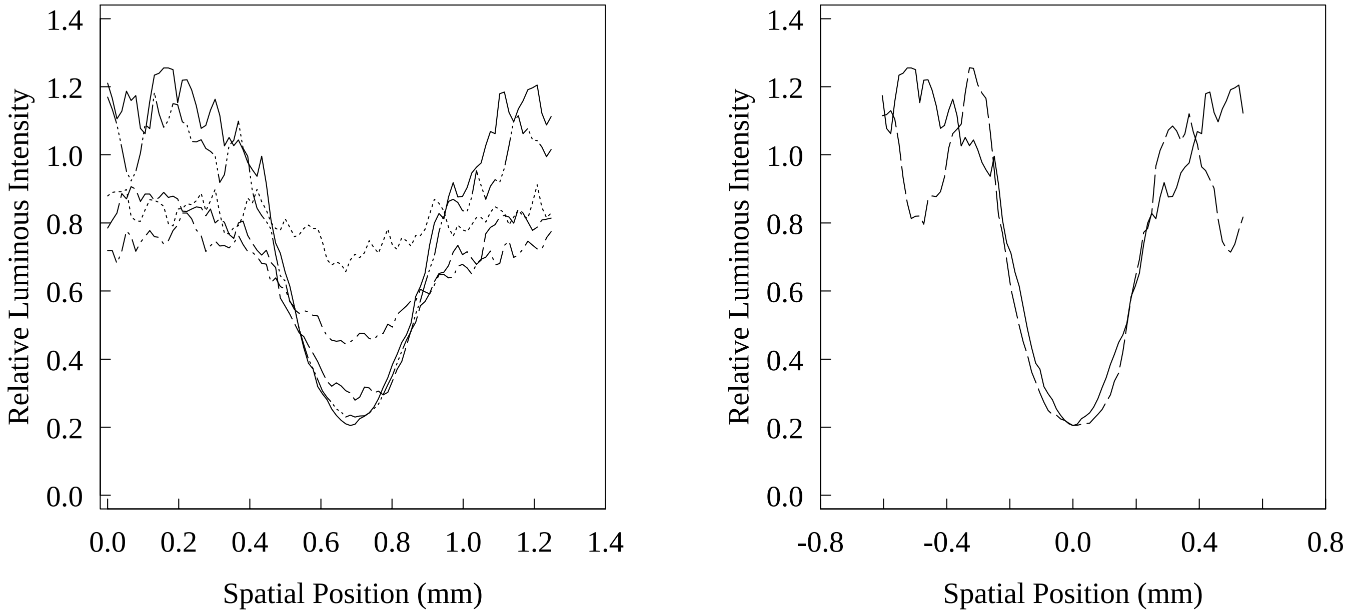

Figures 13: Horizontal Profiles of a Black Pixel (left) & 14: Horizontal vs. Vertical Profiles of Black Pixel (right). The bottom (solid) curve on the left is the horizontal profile of a black pixel taken through the vertical centroid of the pixel. The other curves represent horizontal profiles of the pixel at distances of 0.0525, 0.105, 0.1575, and 0.21 mm from the vertical centroid. On the right, the horizontal (solid) and vertical (dashed) profiles through the centroid of a black pixel have been superimposed. Note that the central portion of this figure corresponds to the pixel, while the outlying areas correspond to the surrounding white background, either along or perpendicular to the scan line.

All of this leads to the complexity of the two-dimensional profile show above. In view of this complexity, it may be surprising to find that horizontal and vertical profiles through the centroid of a black pixel are as similar as shown in the central part of Figure 14. This is because the complexities of black pixels lie near the diagonals. Owing to these complexities along the diagonals, black pixels lack radial symmetry, and, unlike a white pixel on a black background, a black pixel on a white background is not *x-y* independent.

Figure 15 shows horizontal profiles of dark pixels of varying luminance against a white background. Figure 16 shows the sum of each of the horizontal profiles of a dark pixel on a white background with the complementary profile of a light pixel on a black background. Evidently, the sums of each pixel pair do not have the same shape. This demonstrates that even along the middle of the scan line, where things are simplest, profiles of black and white pixels are not simple inverses of one another.

We will now examine this shape difference in two dimensions. If one added a white pixel to a black pixel that was its exact complement, it would simply form a white screen. However, a white screen – with its ridges, furrows, and noise – is a complicated and noisy background for inspection of any deviations. To remove this background, we subtracted from the image of a white screen the sum of black and white pixels (Figure 17), thereby forming a flat background against which any failures of cancellation between the pixels can be viewed. If the black pixel were exactly the complement of the white pixel, then the result would be a uniform flat surface. The systematic deviations from flatness, then, show the differences in the shapes of black and white pixels. The direction and shape of these differences demonstrate that the black pixel is larger than the white pixel in all respects.

Luminance

Note that the black pixel on a white background (cf. Figure 15) starts from a luminance higher than that achieved by an individual white pixel on a black background (cf. Figure 4), and reaches minimum about 20% above 0 luminance. Both effects (the higher maximum and lower minimum) are in part consequences of the size of the local area that is white. This can be seen in Figure 18, which shows the horizontal profiles of vertical stripes of varying width (note the change of scale on the

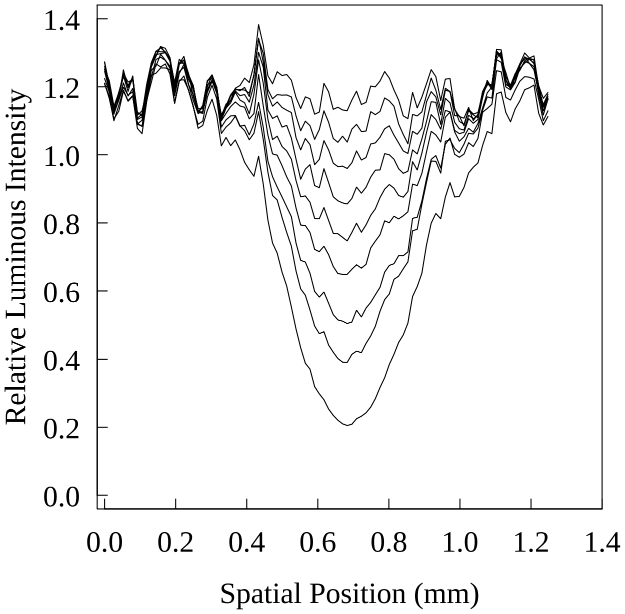

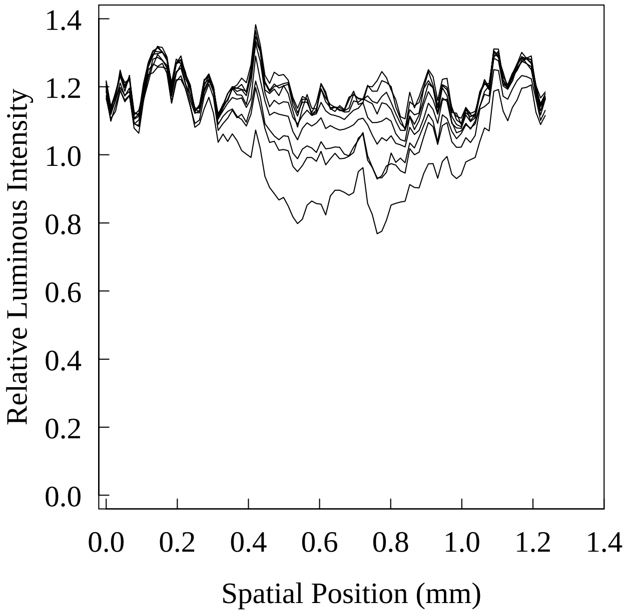

Figures 15: Horizontal Profiles of Dark Pixels (left) & 16: Sum of Horizontal Profiles of Dark and Light Pixels (right). The curves on the left are the horizontal profiles through the centroid of a single pixel presented at 9 nominal luminances, ranging from 0 to 1 in steps of 1⁄8, on a white background. (The mean of the white portion of the scan line is greater than 1.0 because our normalization factor corresponds to the mean luminance of a white screen, including the dark furrows between scan lines. As these curves are in the middle of a scan line, the relative luminous intensity is greater than the average luminance.) The curves in the right figure are the horizontal profiles of the sum of light pixels on a black background and dark pixels on a white background, for 9 pairs of complementary nominal luminances. The bottom (i.e., most deviant) curve represents the case of a white pixel on a black background added to a black pixel on a white background.

horizontal axis). As the width of the stripes increases, both the maximum and the minimum luminous intensities become more extreme. This, then, represents a form of non-linear spatial interaction affecting luminance, that occurs even when the total amount of light called for at the DAC is constant. It is undoubtedly at least partially due to low-pass filtering by the electronics interposed between the output of the DAC and the signal applied to the electron gun. The asymmetry of the effect – i. e., the greater decrease in luminance of the white stripes than the increase in luminance of the black stripes – can be fully accounted for by the non-linearity of the phosphor. Note also that, at half height, all the white stripes are narrower than the black stripes; this corresponds to the finding that black pixels are larger than white pixels. These local spatial non-linearities will be discussed further below, but first it is necessary to treat the effect of more global non-linear spatial interactions on luminance.

Figure 19 shows that the proportion of the screen illuminated also affects luminance at a fixed position. The solid lines show the observed luminance at the center of a central white rectangle as the nominal luminance setting varies (the rest of the screen is set to black). Different lines are for rectangles of different size, but all rectangles maintain the full-screen's aspect ratio and are larger than the area measured with the photometer. We are interested in the integrated luminous flux associated with isolated pixels as well. Our photometer, however, could not measure such a small area. The dashed line shows the integrated luminous flux measured with the CCD array, scaled by the integrated luminous flux of a white screen as recorded by the CCD array. Although the term luminance is not appropriate for point sources, we consider the dashed line to be the inferred luminance of a single pixel – the limiting case (for such displays) in the process of reducing the area of illumination.

The straightness of the lines shows that the system's luminance non-linearity has been corrected; the variations of lines' slopes show how the proportion of the screen illuminated affects luminance at its center. Thus, at any nominal luminance setting for the rectangle (i. e., at any position on the horizontal axis), the *size* of the rectangle affects the luminance at its center. So even when the nominal luminance is set at a constant value, the total light emitted by the display affects the actual luminance at a given location.

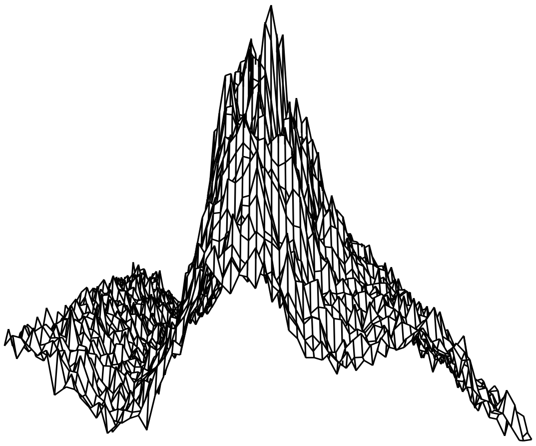

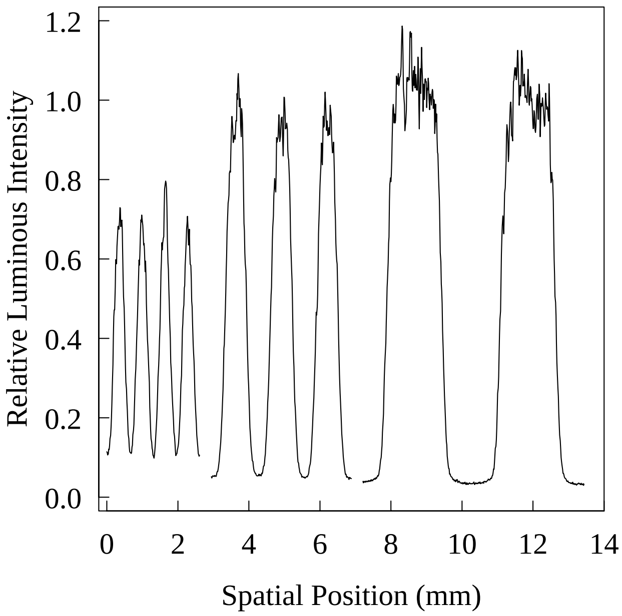

Figures 17: Two-Dimensional Difference of White Background and Sum of Black and White Pixels (left) & 18: Horizontal Profiles of Vertical Stripes (right). On the left is the two-dimensional luminous-intensity profile of the arithmetic difference between a completely white screen (with its ridges and furrows) and the sum of a black pixel and a white pixel, peering down the scan lines of the display. If a black pixel were the exact complement of a white pixel, then this difference would be a uniform flat surface. Compared to the profile of a white pixel on a black background (Figure 2), the difference plot shown here (which has been scaled for clarity) ranges from 6% below the baseline of a white pixel to 65% of the white pixel's peak (the mean of the difference plot is 14% of the white pixel's peak). This implies that a black pixel is larger than a white pixel. There are 64×64 CCD samples in this height field. On the right, the horizontal profiles of vertically-oriented black and white stripes that are 1, 2, and 5 pixels wide are shown. Note that the minima and maxima are offset from 0 and 1 luminous intensity, and that the deviations are worse for the thinner stripes. While the deviations may be attributable to low-pass filtering, the asymmetric growth of these deviations is caused by the luminance point non-linearity of the display system. In this figure, there are 1,024 CCD elements represented along the x-axis.

Figure 20 shows how the luminance at the center of a central white rectangle of fixed size (50% of the area of the entire screen, in this case) is affected by the settings of a surrounding area of fixed size (also 50%, in this case). This further demonstrates that the luminance at one location depends on the total amount of light emitted over the entire display and not simply on the local area emitting light. Note that in both of these figures, increasing the total light emitted by the display also increases the luminance at the center of the screen. This constitutes an over-compensation for the added current required to increase the amount of light produced (cf. [Mulligan and Stone 89]). Other displays exhibit an opposite effect: increasing the total light emitted by the display *decreases* the luminance at the center of the screen [Pelli and Zhang 91]. These figures show the results of global interactions; we now return to the effect of pixel settings on the local distribution of light.

Luminance Superadditivity Along a Scan Line

Figure 21 shows the horizontal profiles of 1, 2, 3, 4, or 5 horizontally-adjacent pixels presented simultaneously at maximal luminance (solid lines). Note that, as more pixels are added, the luminous intensity grows exponentially (as represented by the heavy dashed line), with a space constant of 0.62 pixel. The dotted line shows the arithmetic sum of the individual profiles of the five pixels. The increased luminous intensity with increasing number of adjacent pixels, and the difference between the summed (dotted line) and observed (solid line) luminous intensities, can be called luminance *superadditivity*. Evidently, the second pixel accounts for most of the superadditivity, but some increase (in luminous intensity, and in the area between the solid and dotted lines) is also seen with the addition of the third pixel.

The superadditivity in the pixel profiles presented in Figure 21 is relevant if pixel structure is visible. However, if pixel structure is not discernible, then it is more appropriate to consider how luminous flux per pixel grows with increasing number of adjacent pixels (Figure 22). The asymptotic luminance is almost 30% greater than that of a single pixel. An exponential curve with the space constant 0.62 pixel fits these data well.

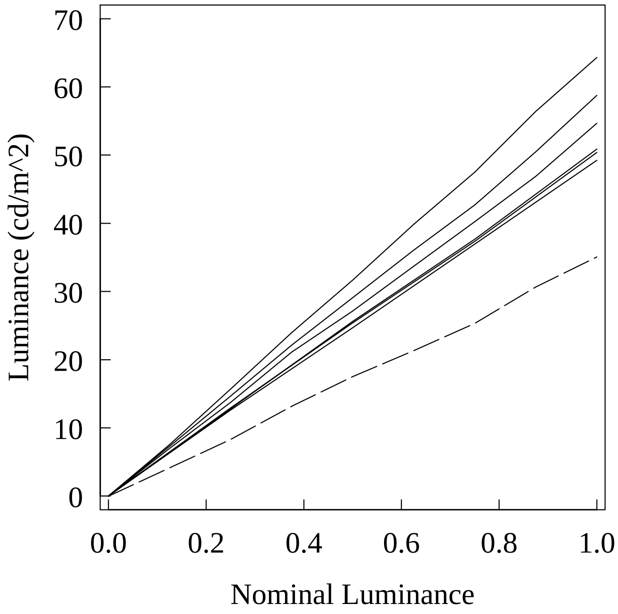

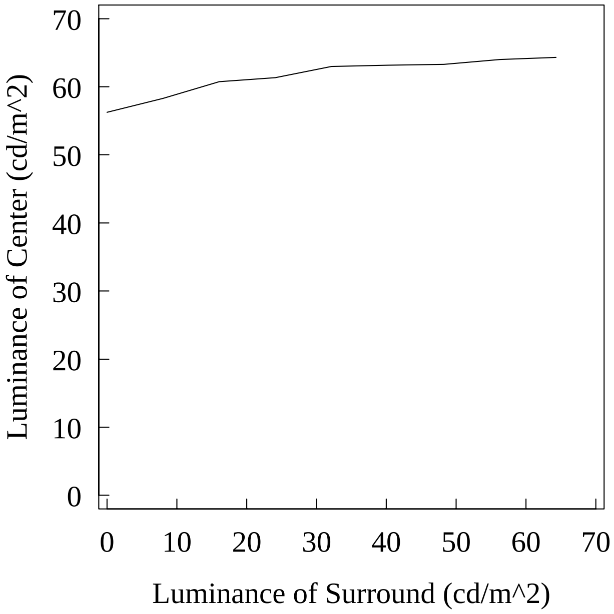

Figures 19: Luminance Dependence on Area Illuminated (left) & 20: Luminance Dependence on Surround Illumination (right). To compensate for the display's luminance point non-linearity, an inversion table was derived from luminance measurements at the center of the screen for full-screen illumination at various voltages. The top line on the left shows the luminance response of the display system for full-screen illumination after loading the LUT with the compensation table. The remaining solid lines show the luminance response as a function of nominal, normalized luminance requested for displays illuminated only in the central 50, 25, 12.5, 6.25, and 1.5625 percent of the screen. In all of these cases, the area illuminated maintained the aspect ratio of a full screen, and the measurements were taken with the photometer. The dashed line corresponds to the inferred 'luminance' of a single pixel, derived from measurements taken with the CCD array (see text). For the measurements shown on the right, the central 50% of the screen was illuminated with white pixels. The curve plots luminance measurements taken at the center of this rectangular region as the surrounding area (also 50% of the screen) was varied from black to white.

The exponential growth of luminance is just what one would expect if the step function intended for the output of the electron gun had been low-pass filtered at an earlier stage of electronics, as suggested by [Mulligan and Stone 89]. Moreover, the space constant reported here is close to the value of 0.65 inferred indirectly by [Mulligan and Stone 89] for a different CRT display. We can think of no way to account for these effects solely on the basis of the properties of point non-linearities attributable to the phosphor. Thus it is attractive to follow [Mulligan and Stone 89] in the speculation that the superadditivity is due to a low-pass filter in the electronics.

Note that, according to this view, it is not the luminance of multiple pixels that is 'supernormal', but the luminance of solitary pixels that is 'subnormal'. A single pixel is subnormal because the maximal slew rate of the amplifier does not allow the voltage to attain full value during the time allowed for a single pixel during the raster scan. *Superadditivity* arises because two subnormal pixels do not add up to match the profile of a string of adjacent pixels, during the presentation of which the amplifier has time to reach full value. Nevertheless, the term 'superadditivity' is appropriate inasmuch as the light emitted by multiple pixels is greater than the corresponding multiple of the light emitted by a single pixel.

Next we examine the superadditivity in more detail by studying interactions between pairs of pixels.

Vertically-Adjacent Pixels

Figure 23 shows the luminous intensity profile of two vertically-adjacent white pixels (i.e., perpendicular to the scanning direction) in an otherwise black field. The dotted lines show the profiles of each pixel when turned on alone, and the solid line is the profile when both are turned on at the same time. Superimposed on the solid line is a dashed line representing the point-by-point sum of the dotted curves; the two are so nearly coincident that they are hard to separate. This means that the profile of two vertical pixels is simply the arithmetic sum of the luminous intensity of each pixel at each point. Hence, there are no detectable spatial interactions in the dimension perpendicular to the scan lines of the display, as has been reported previously [Lyons and Farrell 89].

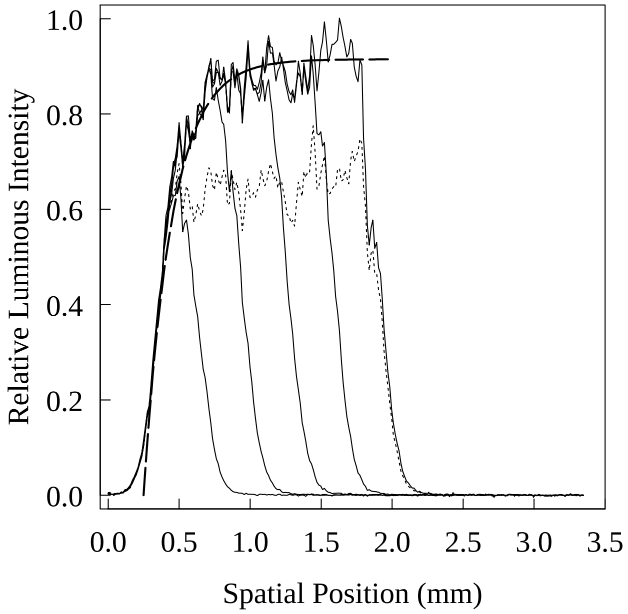

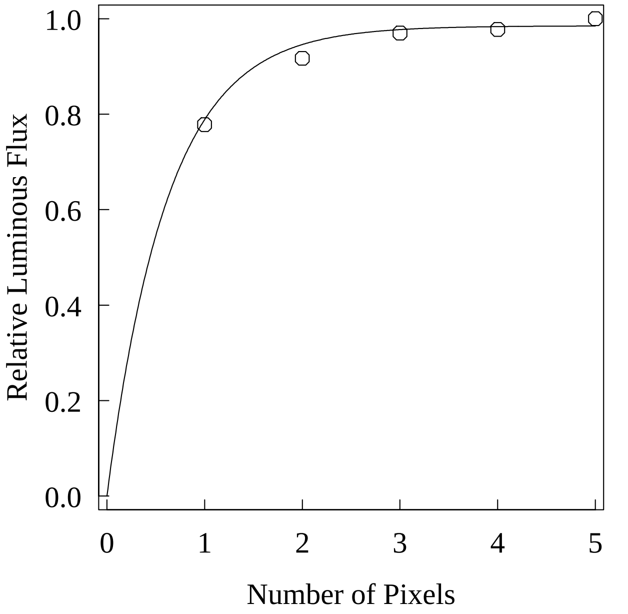

Figures 21: Superadditivity of Luminous Intensity (left) & 22: Superadditivity of Luminous Flux (right). On the left, horizontal profiles of 1, 2, 3, 4, and 5 horizontally-adjacent pixels set to maximum luminance are shown (solid lines). The dotted line represents the sum of the individual profiles of the five pixels. The heavy, dashed line is an exponential curve with a space constant of 0.62 pixel (0.2 mm) and an asymptote of 0.915. The curve does not approach an asymptote of unity because a single row never has the same luminance as an entire screen (cf. Figure 19), but maximal luminous intensity is defined here according to measurements made with full-screen illumination. In this figure, there are 256 CCD elements represented along the x-axis. The figure on the right shows the per-pixel relative luminous flux corresponding to the pixel profiles shown on the left. The solid line is an exponential curve with the same space constant as on the left.

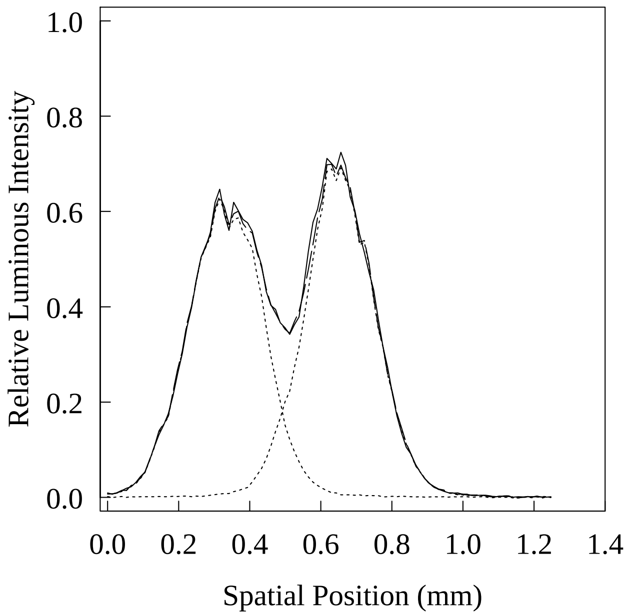

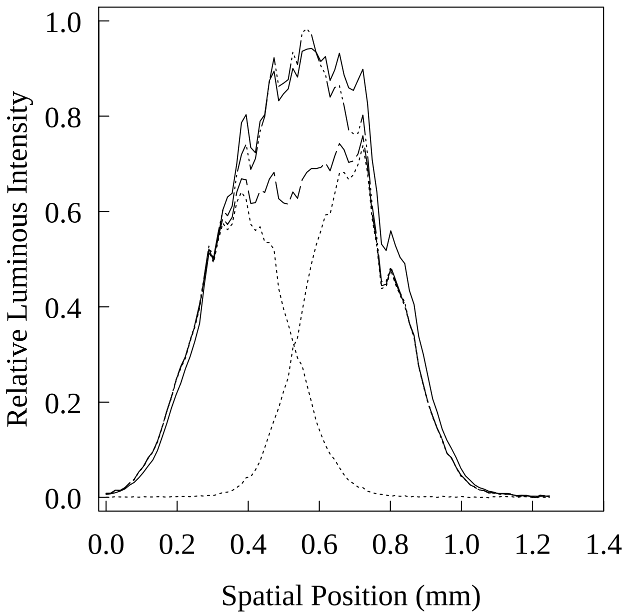

Figures 23: Profiles of Vertically-Adjacent White Pixels (left) & 24: Profiles of Horizontally-Adjacent White Pixels (right). The solid curve in the figure on the left is the vertical profile through the horizontal centroid of two vertically-adjacent white pixels. The dotted curves are the profiles of the two pixels when illuminated individually. The dashed curve is the arithmetic sum of the two dotted curves, and straddles the solid curve so closely as to make them hard to separate. The figure on the right presents analogous curves for two horizontally-adjacent white pixels. In this instance, the dashed curve falls well below the solid curve. The dotted and dashed curve is the sum of the dotted curves plus 2.86 times their product.

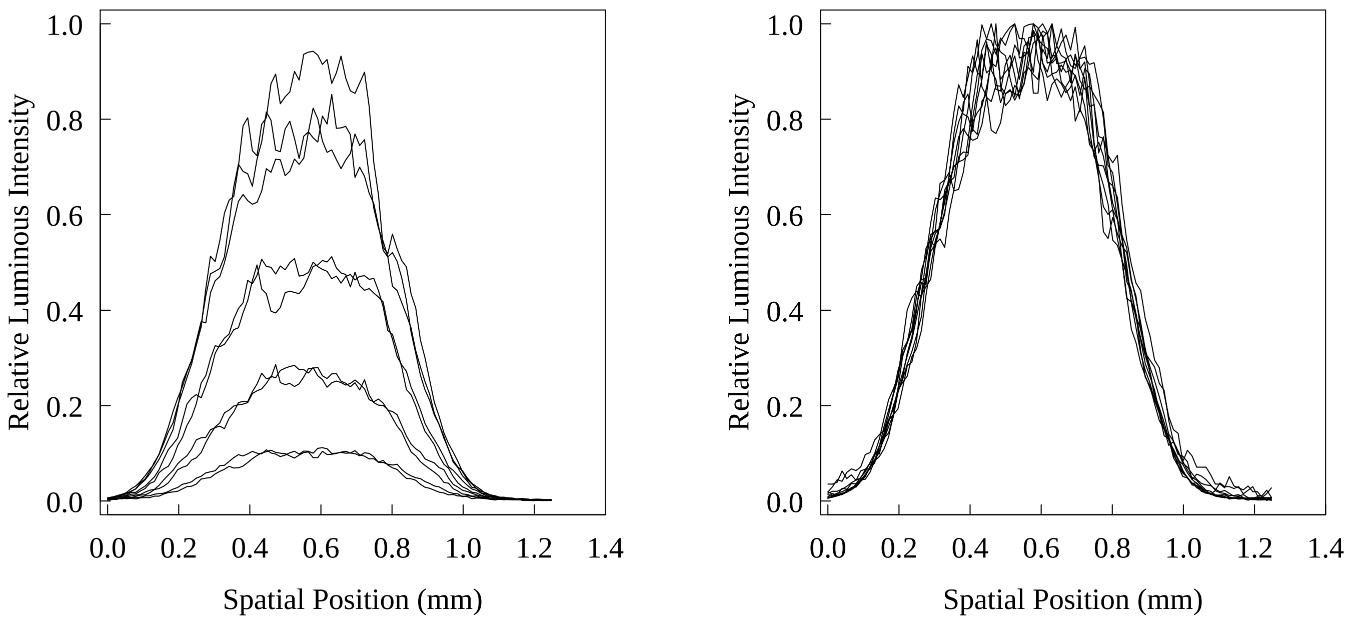

Figures 25: Horizontal Profiles of Two White Pixels (left) & 26: Normalized Horizontal Profiles of Two White Pixels (right). The top curve on the left is the horizontal profile of two white pixels taken through the vertical centroid of the pixels. The pairs of curves below that one represent equi-distant horizontal profiles above and below the vertical centroid at distances of 0.0525, 0.105, 0.1575, and 0.21 mm. These profiles have been scaled on the right to demonstrate that the horizontal profile of a pair of white pixels is independent of its vertical position within the pixels. Note that, in this case, the curves have been normalized to the height of the horizontal profile through the pixels' vertical centroid (i.e., the highest curve in the figure on the left).

Horizontally-Adjacent Pixels

Figure 24 shows the luminous intensity profile of two horizontally-adjacent white pixels (i.e., along the scanning direction) in an otherwise black field. The dotted lines show the profiles of each pixel when turned on alone, and the solid line is the profile when both are turned on at the same time. The dashed line is the point-by-point sum of the dotted curves. However, in this case, the dashed line does not coincide with the solid line but falls well below it (as in Figure 21), exhibiting the superadditive spatial interaction introduced there. Moreover, the interaction is multiplicative, in the sense that the superadditivity – i. e., the difference between the solid and dashed lines – is approximately proportional to the product of the individual pixel luminances. The curve drawn in dots and dashes represents the sum of the individual pixel profiles plus 2.86 times their product. The left half of this curve fits the observed profile (solid line) reasonably well, but there is a small but consistent deviation on the right side.

Figure 25 shows the profiles of a pair of pixels at different vertical positions in the scan line. Only the height of the profiles varies across the scan line; the superadditive component remains proportional to the quantities being added, for these profiles are identical after appropriate scaling (Figure 26). This establishes that the same constant of proportionality applies at all vertical positions along the (horizontal) scan line.

The constant of proportionality, however, does vary slightly with the luminance of the pixels. Figure 27 shows profiles of the pixel pair at different luminances. Figure 28 shows that after these profiles have been scaled up appropriately, the profiles for the three lowest luminances fall slightly above the rest of the curves. That is, the superadditivity for the lowest luminances is slightly greater than that of the higher luminance pixels. The peaks of the profiles of the dimmest pixels are proportionately about 10% greater. Since individual pixels do scale proportionately (Figure 5), this excess must be attributed to the interactions. Since about 30% of the peak height of the profiles is due to the superadditivity, an additional 10% added to the peak corresponds to about a one third increase in the superadditive component. This *super*-superadditivity of dim pixels can also be seen in the next few figures.

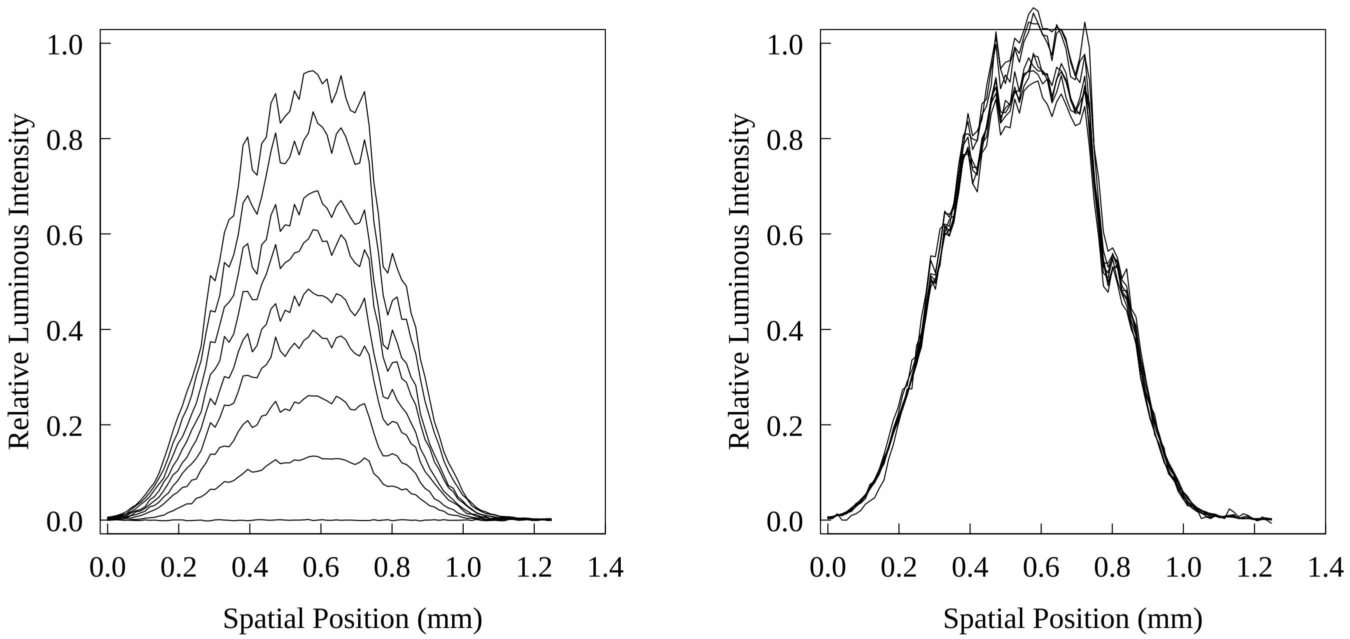

Figures 27: Horizontal Profiles of Two Light Pixels (left) & 28: Normalized Horizontal Profiles of Two Light Pixels (right). The curves on the left are the horizontal profiles through the centroid of two horizontally-adjacent pixels presented simultaneously at 9 nominal luminances, ranging from 0 to 1 in steps of 1⁄8. On the right, the curves have been scaled by their nominal luminances. The three highest curves correspond to the three lowest luminances, indicating a *super*-superadditive characteristic of the spatial interaction for dim pixels.

Figure 29 shows the luminance profiles of a simultaneously-presented pair of horizontally-adjacent pixels with the left pixel at 5 different, equally-spaced nominal luminances, and the right pixel at maximum luminance. Figure 30 shows the same data after the luminance profiles of the left pixel alone (shown at maximum luminance by the dotted line) has been subtracted from the overall profiles. The profile of the right pixel alone is shown by the dashed line. If there were no interactions between the pixels, the solid curves would coincide with one another and with the dashed curve. The differences between the solid curves and the dashed curve, then, demonstrate the superadditivity. This figure shows that half the superadditivity is observed when the left pixel is only at 25% of maximum, owing to the super-superadditivity. Hence, the superadditivity is a different proportion of the product of the two pixel profiles depending on the luminance settings of the pixels.

The asymmetry of the fit between the upper two curves in Figure 24 suggests that the superadditivity is asymmetrical. Figure 31 demonstrates this asymmetry. The dashed lines show the individual profiles of two pixels where the left one is at maximum luminance and the right one is at 75% of maximum luminance; the dotted lines show the reverse situation, where the left pixel is at 75% of maximum and the right pixel is at maximum luminance. The solid curves show the superadditivity observed when these pixels were presented simultaneously: they are the profiles of the pixels presented together, after the profiles of the individual pixels presented alone have been subtracted. The asymmetry of the superadditive component is evident in both curves, and the superadditive component is about one third of the total luminous intensity at the intersection of the pixel profiles.

Figure 32 shows the profile of two simultaneously-presented white pixels along a scan line, when separated by a single black pixel. Although no detectable interaction or superadditivity is evident in the spatial profile, we know from above (cf. Figures 21 and 22) that the superadditivity for consecutively-illuminated pixels along a scan line extends to the third pixel.

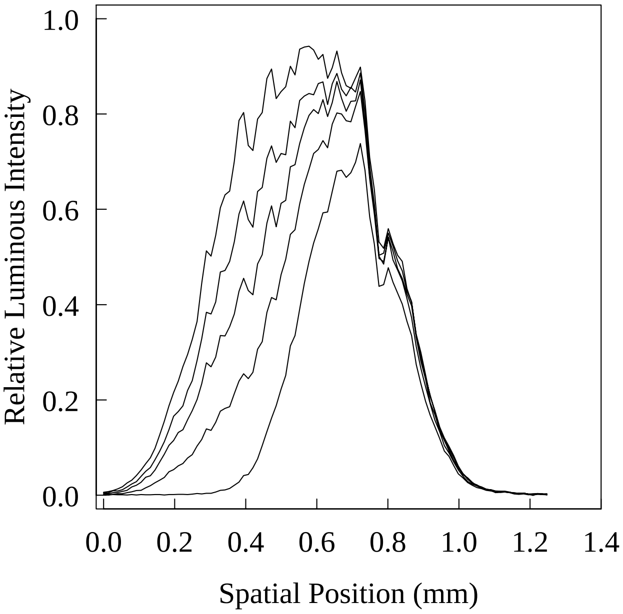

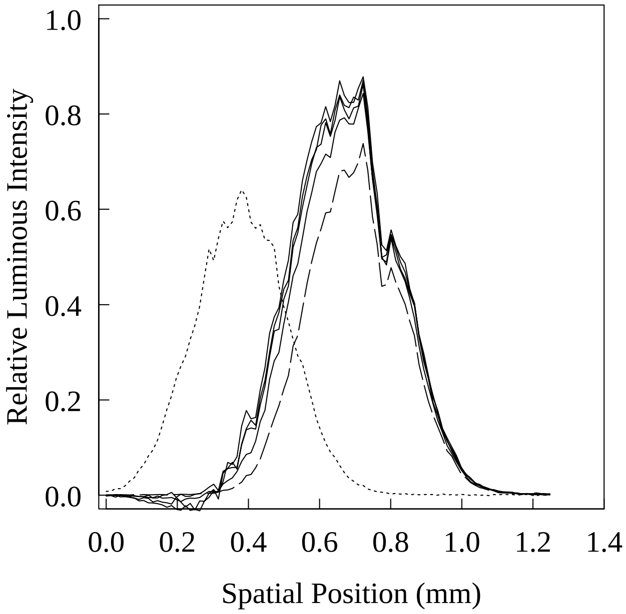

Figures 29: Profile of White Pixel and Neighbor as the Neighbor Varies (left) & 30: Profile of White Pixel as the Neighbor Varies (right). The lowest curve on the left is the horizontal profile of a white pixel. The four other curves are the horizontal profiles of the pixel and its left neighbor, as the neighbor is changed from black to white in equal steps. From these four curves, the profiles of the left neighbor at the corresponding setting have been subtracted, resulting in the four solid curves on the right. The profile of the neighboring pixel at maximum luminance is shown as the dotted curve, and the profile of the pixel at maximum luminance is shown as the dashed curve. The fact that the solid curves do not coincide with the dashed curve demonstrates the superadditivity (see text for further details).

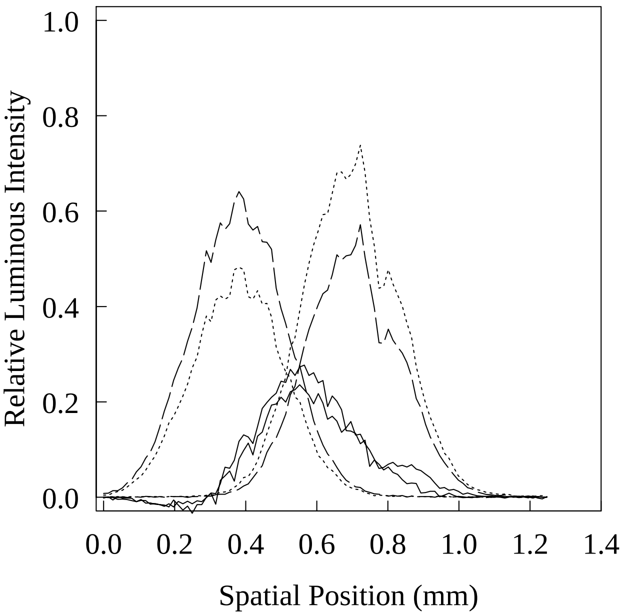

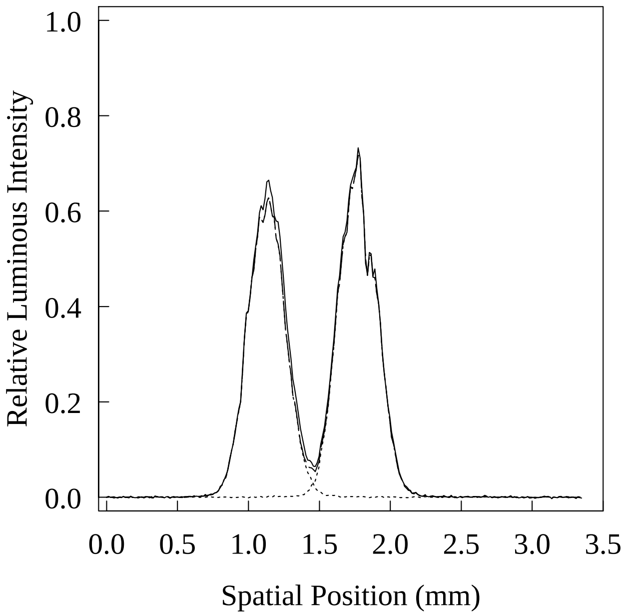

Figures 31: Asymmetry of Superadditivity (left) & 32: Non-Interaction When White Pixels are Separated by a Black Pixel (right). On the left, the dashed curves are the individual profiles of a left pixel at maximum luminance and a right pixel at 75% of maximum, while the dotted curves are the analogous profiles for the reverse situation. Each solid curve is the difference between the measured profile of the pair of pixels presented together and the sum of the profiles of the individual pixels presented alone. The asymmetry of these curves demonstrates that superadditivity is asymmetric. The solid curve in the figure on the right is the horizontal profile through the vertical centroid of two white pixels along a scan line that are separated by a single black pixel. The dotted curves are the profiles of the two white pixels when illuminated individually. The dashed curve is the arithmetic sum of the two dotted curves, and straddles the solid curve so closely as to make them hard to separate.

Conclusions

It is typical in computer graphics to represent an image as a matrix of independent points corresponding to the screen pixels, but this disregards the interactions among pixels observed here. These interactions can pose problems for demanding applications, and they raise questions about work in which display calibrations or sensitivity to these problems are not shown.

We can summarize the results we obtained for the particular grayscale monitor we investigated as follows:

- The luminance at the center of a fully-illuminated screen exhibits a non-linear dependency on the voltage applied to the electron gun. This relationship is well-characterized by a power function with an exponent of 1.45.
- The luminance at the center of the screen depends both on the proportion of the screen illuminated and on the luminance of a distant surround.
- The shape of a white pixel on a black background is independent of luminance, and horizontal profiles are independent of vertical position in the pixel.
- The shape of a black pixel on a white background depends on luminance, and the shapes of different cross-sections change in different ways.
- White pixels are slightly elongated in the direction of the scan line, but black pixels have complex shapes with many inflections.
- Even in the central-most region, where black pixels resemble inverted white pixels, the white pixels are smaller and more leptokurtic.
- The contrast of vertical black and white stripes decreases as their widths decrease from five pixels to a single pixel, and the mean luminance of the entire display decreases as well.
- The luminance and peak luminous intensity grow exponentially (τ=0.62 pixel) with increasing number of simultaneously-presented adjacent pixels along the raster. Hence, detectable spatial, non-linear interactions extend over three neighboring pixels.
- The interactions along the scan line between a pair of adjacent pixels are superadditive and asymmetric. The superadditive component increases the total luminance flux by about one third.
- The superadditive contribution to each point in the profile of a pair of adjacent pixels is proportional to the product of the individual pixel luminances, but the constant of proportionality depends on the luminance of the pixels.
- There are no significant local, spatial interactions between pixels separated by a black pixel.
- There are no significant local, spatial interactions in the dimension perpendicular to the scan lines.

These observations demonstrate different types of non-linearities associated with CRT displays: a non-spatial, or point, non-linear relationship between phosphor luminance and voltage applied to the electron gun (Figure 1); a global, spatial non-linearity exhibited as a dependence of luminance at the center of the screen on the proportion of the screen illuminated (Figure 19) and the luminance of distant regions (Figure 20); and a local, spatial, non-linearity observed as a superadditive interaction between neighboring pixels within a single scan line (e. g., Figures 24 and 31). The spatial non-linearities would have been even more pronounced had the point non-linearity of the phosphor been more typical of other displays we have encountered.

These observations have several clear implications:

- The use of lookup tables cannot compensate for the spatial interactions.
- Dithering (also called halftoning) attempts to simulate intermediate luminances (or colors) by distributing a set of pixels that cannot be individually distinguished (see, for example, [Foley *et al.* 90, Ulichney 88]). The spatial non-linearities reported here can cause relative errors as great as 60% in the results of dithering techniques that assume linear spatial summation [Naiman 91].
- Each of these non-linearities violates the assumption of superposition on which the use of modulation transfer functions depends.
- A single white pixel on a field of black is *not* simply the reverse polarity of a single black pixel on a white background. Hence, reversing the polarity of high spatial-frequency stimuli is likely to affect image quality.

These results are restricted to the physical properties of a particular CRT display. How these appear to the eye is an entirely separate question that we have not addressed here. Moreover, although we believe that the spatial non-linearities presented here are common to most CRT display systems, it remains to repeat these measurements on more displays of the same type we used, on other grayscale displays, on color displays, and on other display technologies. These non-linearities demand special care where precise control over image quality is required, as, for example, in experiments on spatial vision.

Acknowledgements

This research was supported by US Public Health Service Grants EY-04885, EY-07125, and EY-01319, and the Natural Sciences and Engineering Research Council of Canada. We thank the Eastman Kodak Company for supplying the CCD array and David R. Williams for the loan of his CCD camera system.

References

Conrac 85 — Conrac Corporation, *Raster Graphics Handbook, Second Edition,* Van Nostrand Reinhold, New York, 1985.

Cowan 83 — William B. Cowan, ‘‘An Inexpensive Scheme for Calibration of a Colour Monitor in Terms of CIE Standard Coordinates,’’ *Computer Graphics*, Volume 17, Number 3, July 1983, pp. 315-321. SIGGRAPH 1983 Proceedings.

Foley *et al.* 90 — James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes, *Computer Graphics, Principles and Practices,* Addison-Wesley Publishing Company, Menlo Park, California, 1990.

Lyons and Farrell 89 — Nic P. Lyons and Joyce E. Farrell, ‘‘Linear Systems Analysis of CRT Displays,’’ *Society for Information Display 89 Digest*, Volume 20, 1989, pp. 220-223.

Mulligan and Stone 89 — Jeffrey B. Mulligan and Leland S. Stone, ‘‘Halftoning Method for the Generation of Motion Stimuli,’’ *Journal of the Optical Society of America A*, Volume 6, Number 8, August 1989, pp. 1217-1227.

Murch and Weiman 90 — Gerald M. Murch and Novia Weiman, ‘‘Assessing Visual Grey Scale Sensitivity on a CRT,’’ *SPIE Proceedings: Human Vision and Electronic Imaging: Models, Methods, and Applications*, Volume 1249, February 1990, pp. 214-223.

Naiman 91 — Avi C. Naiman, ‘‘CRT Spatial Non-Linearities and Luminance Linearization,’’ in *Raster Imaging and Digital Typography*, Edited by Robert Morris, Cambridge University Press, Cambridge, England, October 1991. Proceedings of the RIDT91 International Conference, Boston, Massachusetts, In Press.

Pelli and Zhang 91 — Denis G. Pelli and Lan Zhang, ‘‘Accurate Control of Contrast on Microcomputer Displays,’’ *Vision Research*, Volume 31, Number 7/8, 1991, pp. 1337-1350.

Ulichney 88 — Robert Ulichney, *Digital Halftoning,* MIT Press, Cambridge, Massachusetts, 1988.