Christopher Tsai
January 26, 2008
EE 362 – Applied Vision and Image Systems

Tutorial II – Color Matching

Problem #1 – Using the Color Matching Functions

(a.) In order to match a monochromatic test light with a given wavelength, we must match the XYZ representation of that test light. After extracting the proper XYZ values that we seek to match, we convert our CRT monitor phosphor spectral power distributions (SPDs) into the same XYZ space and determine the necessary weights (or coefficients) ***n*** of each phosphor type:

$$\text{XYZ}_{\text{target}} = [\text{XYZ}]\ [\text{phosphors SPD}]\ [\text{weights}]$$

$$\begin{bmatrix} \text{X}(550\text{ nm}) \\ \text{Y}(550\text{ nm}) \\ \text{Z}(550\text{ nm}) \end{bmatrix} = \begin{bmatrix} \dots & \text{X}(\lambda) & \dots \\ \dots & \text{Y}(\lambda) & \dots \\ \dots & \text{Z}(\lambda) & \dots \end{bmatrix} \begin{bmatrix} \vdots & \vdots & \vdots \\ p_1(\lambda) & p_2(\lambda) & p_3(\lambda) \\ \vdots & \vdots & \vdots \end{bmatrix} \begin{bmatrix} n_1 \\ n_2 \\ n_3 \end{bmatrix}$$

$$(3 \times 1) \quad = \quad (3 \times \lambda) \quad (\lambda \times 3) \quad (3 \times 1)$$

To determine the appropriate phosphor SPD weights, we must invert the matrix pre-multiplying the coefficient matrix; luckily, $[\text{XYZ}]\ [\text{phosphors SPD}]$ is square (3 × 3) and full rank, so we can invert it without issue to obtain the following weights for 550 nm and 430 nm matching monitor phosphor distributions, which we compute both in raw weight and in normalized intensity weight:

WAVELENGTH	**COMPONENT**	**COEFFICIENT**	**INTENSITY**
$\lambda = 550$ nm	Phosphor p_1	-0.003039	-0.324728
	Phosphor p_2	-0.015098	1.613189
	Phosphor p_3	-0.001450	-0.155028
$\lambda = 430$ nm	Phosphor p_1	0.004142	0.442530
	Phosphor p_2	-0.002641	-0.282234
	Phosphor p_3	0.011835	1.264945

(b.) By weighing the CRT monitor phosphors' spectral power distributions by the pre-calculated amounts, we achieve the matching spectral distributions:

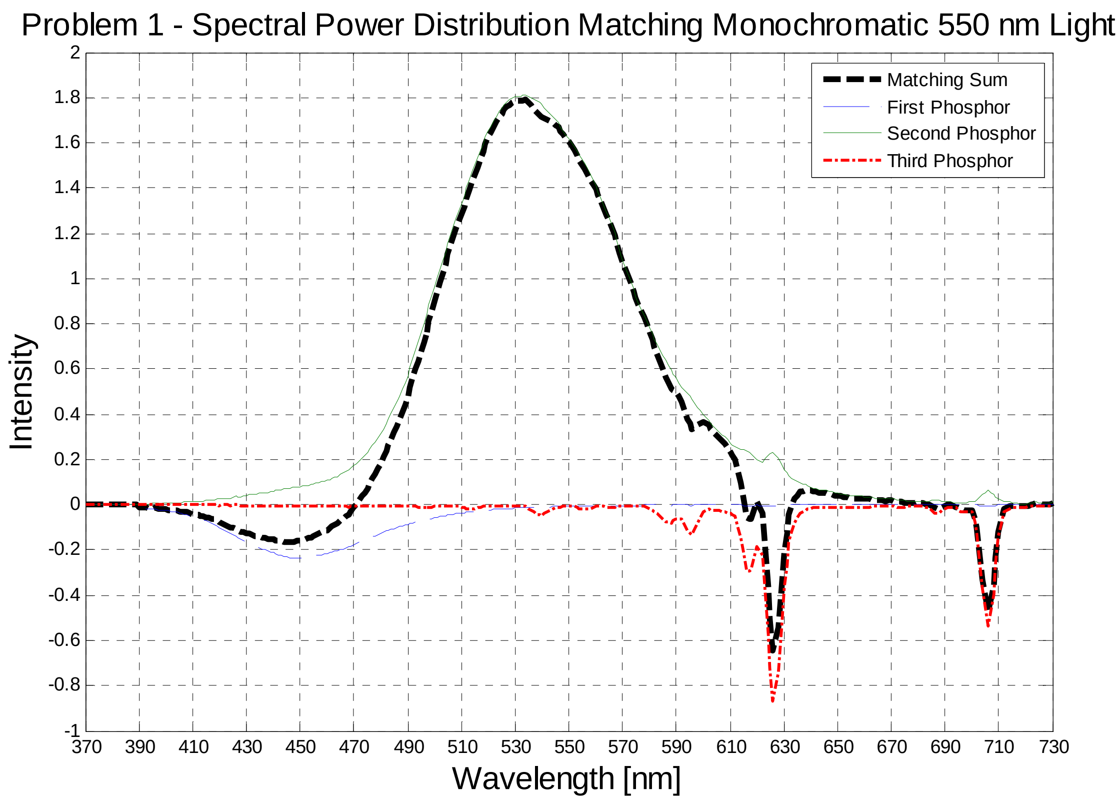
Problem 1 - Spectral Power Distribution Matching Monochromatic 550 nm Light
Matching Sum
First Phosphor
Second Phosphor
Third Phosphor
Intensity
Wavelength [nm]
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
-1
370
390
410
430
450
470
490
510
530
550
570
590
610
630
650
670
690
710
730

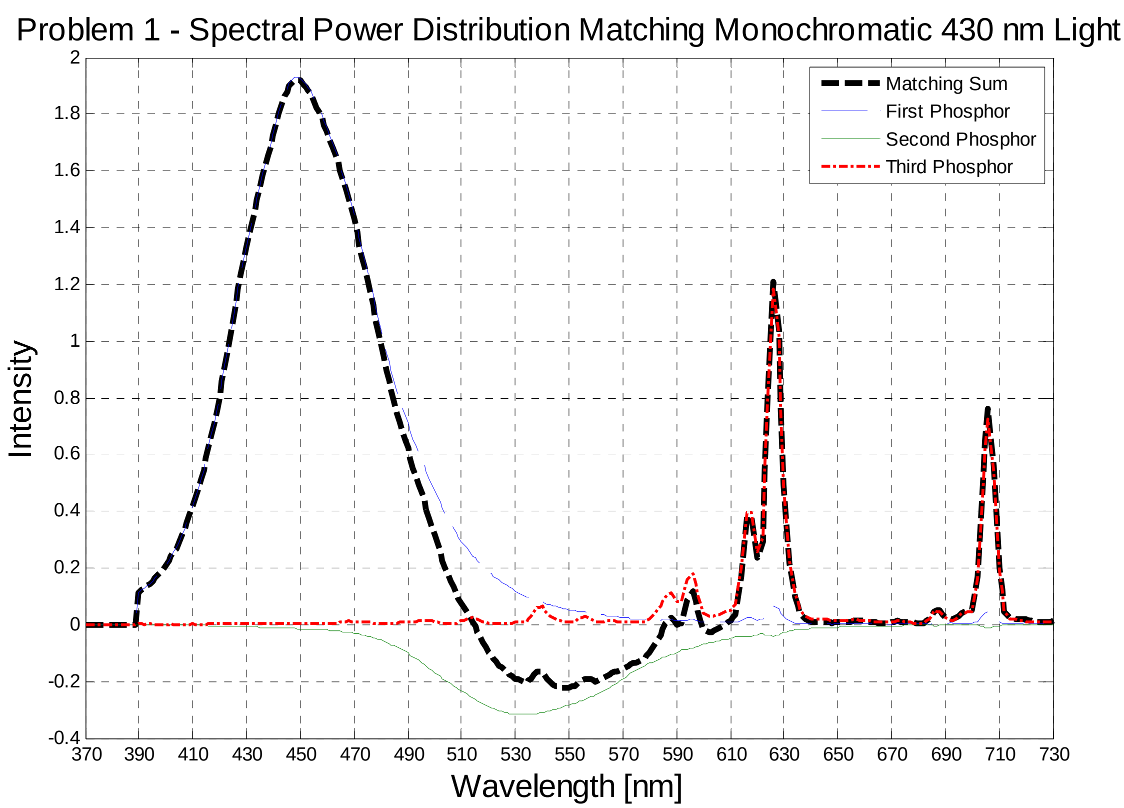
Problem 1 - Spectral Power Distribution Matching Monochromatic 430 nm Light
Matching Sum
First Phosphor
Second Phosphor
Third Phosphor
Intensity
Wavelength [nm]
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
370
390
410
430
450
470
490
510
530
550
570
590
610
630
650
670
690
710
730

However, we notice that these intensities are not physically realizable because their spectral power distributions dip below zero. As mathematically sound as a negative coefficient might seem, no laboratory display device can produce a negative power spectrum. As a result, additive mixtures of the phosphors cannot produce certain test light configurations, such as the ones proposed in the problem (monochromatic 550 nm and monochromatic 430 nm); if we want to generate situations with these configurations, then we must add the negatively weighted spectral power distributions to the test light rather than subtracting them from the CRT monitor phosphors. In other words, because intensity cannot be negative and therefore cannot be subtracted, the only feasible way to match these monochromatic (550 nm and 430 nm) test lights is to *add* the sub-zero portions of the CRT monitor primaries to the test light; in this manner, the test side, though modified, can be altered to match the positively weighted phosphor primaries.

(c.) If we wanted to perform an experiment to test the predicted match, then we would have to arrange a bipartite field, with the test light to one side and the primaries on the other. Because we have concluded that additive combination of the primary lights cannot achieve a perfect match due to negative weights, we must also allow the subject to add portions of the primary to the test light side of the bipartite field, effectively implementing primary power subtraction. Thus, the subject would be given the freedom to adjust the relative amounts of CRT monitor phosphor primaries on *both* sides of the screen until the two lights of the bipartite field look identical, forming a metamer combination. At this point, to compute and compare the weights necessary for the subject to perceive a color match, we numerically subtract the weight of CRT phosphor intensities added to the test light side from the CRT phosphor intensities on the phosphor primary side to form the isomeric phosphor spectral power distributions plotted above.

Problem #2 – Color Monitor Calibration and Chromaticity

(a.) If the spectral power distributions of the three phosphors of a color monitor are *R(λ), G(λ),* and *B(λ)*, then the spectral power distribution of a monitor pixel with color vector [r, g, b] is given by the matrix equation:

$$X(\lambda) = \begin{bmatrix} \vdots & \vdots & \vdots \\ R(\lambda) & G(\lambda) & B(\lambda) \\ \vdots & \vdots & \vdots \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix} = rR(\lambda) + gG(\lambda) + bB(\lambda)$$

$$(3 \times 1) = \quad (\lambda \times 3) \quad (3 \times 1)$$

In other words, we arrange the spectral power distributions of the color monitor as the columns of a system matrix and then linearly combine weighted portions to obtain the spectral power distribution of an individual pixel. The pixel weights determine the relative proportions of the monitor SPDs required to represent the pixel in the desired color.

(b.) The spectral power distributions emitted from a pixel do not match the spectral power distribution of the environment, simply because the color monitor has its own limited set of thtree basis functions from which it generates the pixel SPD. For example, consider the color monitor:

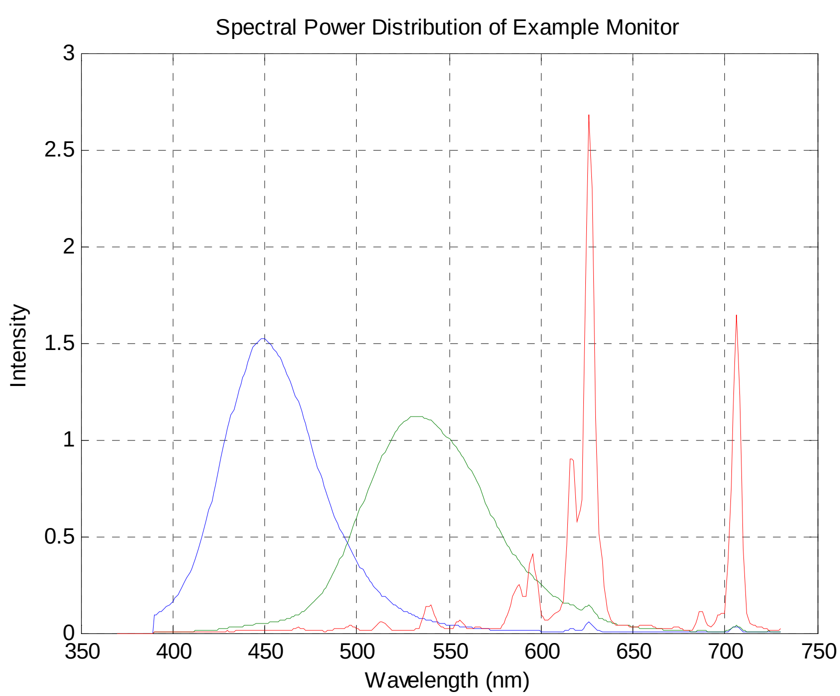

Let us juxtapose this set of phosphor SPDs with the set that Judd proposed for modeling daylight:

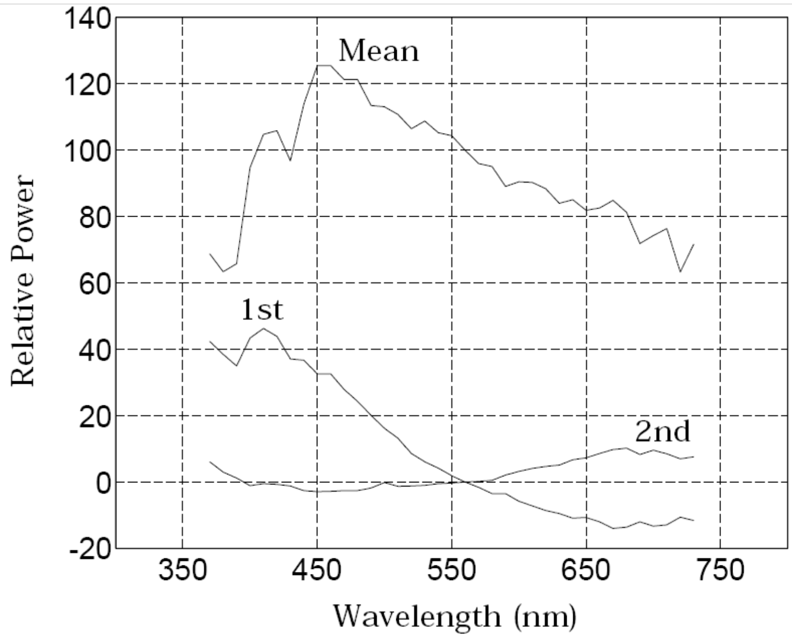

Figure 1.7: *A linear model for daylight spectral power distributions.* The curve labeled mean is the mean spectral power distribution of a set of daylights whose spectral power distributions were normalized to a value of 100 at 560nm. The curves labeled 1st and 2nd show the two basis curves used to define a linear model of daylights. By adding together the mean and weighted sums of the two basis functions, one can generate examples of typical relative spectral power distributions of daylight.

As Wandell explains in *Foundations of Vision* (Chapter 9, page 297-298), daylight possesses a distinct spectral power density; Judd's experiments ascertain that even different daylight conditions share roughly similar spectral power distributions. Thus, we can devise a linear model of daylight with a set of only three basis functions – the mean and two additional bases – to build the spectral power distribution of daylight under different conditions. However, the color monitor is not built to replicate these three basis functions; for example, notice that both the "1^{st}" and "2^{nd}" SPD basis functions displayed above dip below zero. Hence, even if we can perfectly represent any possible SPD in the environment using Judd's linear model, a color monitor cannot reproduce the constituent functions because it can generate only positive intensities. Furthermore, while *daylight* might be replicable with a small, limited basis, the environment is not an empty void with nothing but light from the sun; other objects populate our world, and light reflecting from and through their surfaces complicates the spectral power density of light circulating the natural environment. This wide variety of natural objects in the environment – leaves, branches, grasses, rocks, and minerals – complicate the reflectance spectrum; as a result, even pure daylight acquires more inimitable subtleties. According to Wandell, this additional complexity demands seven or even eight basis functions to duplicate on the laboratory monitor. Thus, whereas a pixel *could* parallel daylight's SPD using three basis functions, color monitors neither use the same three basis functions that theorists have developed for modeling daylight nor boast a sufficient number of basis functions to capture the subtleties that arise in the object-populated environment's reflectance spectra.

The chromaticity diagram reveals the number of possible spectral energy distributions we can hope to reconstruct from a typical color monitor:

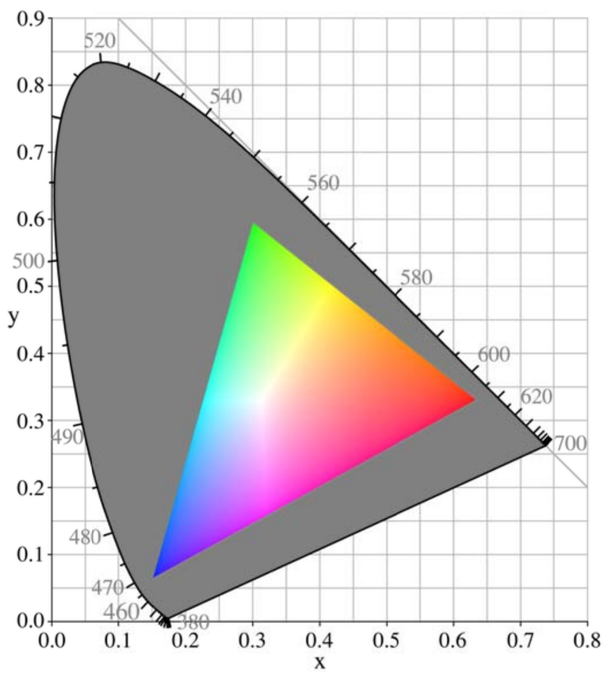

The rainbow triangle represents the range of recoverable colors from a phosphor set given by the three vertices of the triangle. Whereas daylight white falls into this gamut, the presence of natural objects such as plants and animals ensures a much wider variety of colors than those found in the triangular gamut, so, working from only the monitor's three phosphors, the monitor pixels would inevitably exclude some of the SPDs that could arise from object reflectance in the environment.

(c.) The number of hue and saturation levels that the human eye can distinguish remains largely recondite. In other words, given a continuum of colors, how finely can the human eye distinguish quantized levels along that continuum? The answer remains debatable. In general, vision scientists agree that humans can distinguish a number of different colors (hues) approaching the order of millions, but the eye cannot differentiate more than a few hundred shades or tones. In that sense, given that a color monitor produces colors from a set of three phosphors that roughly approximate the three primary colors (red, green, and blue), the images that we view under 16-bit modulation (65,536 different intensities per phosphor channel!) of intensity do not represent a noticeably improved representation than the more common 8-bit modulation (256 levels per color). Despite the fact that we discern a wide range of colors, the number of levels of each color (and their mixtures) that we can perceive likely does not exceed 256; even if it does, the amount of information lost in an 8-bit quantization would not severely degrade the accuracy of the image representation. Thus, if each primary color (phosphor) has an 8-bit quantization (24-bit total), then further improving the intensity resolution with 16-bits (48-bit total) should not vastly improve color appearance. Although some aficionados may disagree, 8 bits of shade/tint are more than adequate for quantizing the range of colors discernable to the human eye. In total, this 8-bit representation provides a set of $2^8 \times 2^8 \times 2^8$ = 16,777,216 (nearly 17 million) different colors to view on the display. This gargantuan number well approximates – or even exceeds – the number of colors that the human eye can discern. Incrementing the bit representation to 16 bits, the number of displayed colors grows beyond 2.8×10^{14}, which seems like inordinate overkill since it exceeds the limit of discernable colors in the human eye; such a fine scale is definitely unnecessary!

(d.) If we are able to build a monitor with a fourth phosphor, then the number of reproducible colors would grow to accommodate the additional basis function and the augmented region of possible linear combinations that we could form using this additional flexibility. In terms of the chromaticity diagram, we would be adding a fourth vertex, thereby possibly increasing the gamut of spectral power distributions that our color monitor could represent. We strategically position the fourth vertex so that the resultant quadrilateral encloses the maximal area possible within the chromaticity diagram:

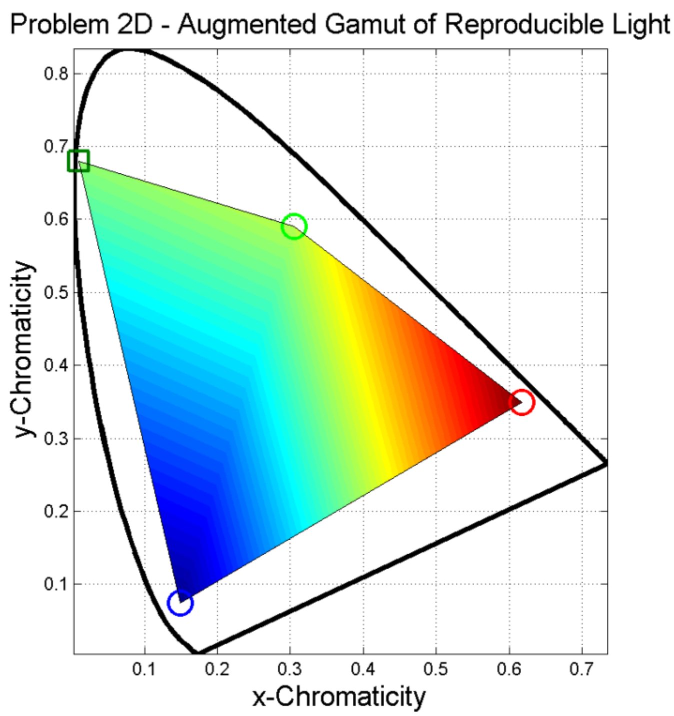

This monitor is better than the three-primary monitor in the sense that it can realize a wider variety of colors through additive mixture. By combining weighted combinations of the spectral power distributions at the *four* vertices, we can produce a greater range of intermediate colors, which the augmented gamut area represents. By placing the fourth phosphor as far from the other three as possible from a chromaticity diagram standpoint, we maximize the gamut area and hence maximize the width of our reproducible color range. The point pictured in the upper-left corner of the chromaticity locus lies at the chromaticity coordinates (x, y, z) = (0.01, 0.68, 0.31).

Problem #3 – Real LCD Display

In a real LCD display, the liquid crystal gates are unable to completely block the backlight. Due to such leakage, the contrast ratio between a fully black and fully white display is 100:1. Because the backlight is its own source of light and intensity, it contributes to the range of achievable colors in the chromaticity diagram. If the backlight is a light source that is always turned on, then it may warp the gamut of producible colors due to its incessant, irremovable presence. For example, suppose that the backlight produces white light with chromaticity values comparable to those at the white point; in that event, we cannot extinguish the addition of white to our final image, however small it may be, so the color gamut contracts to reflect the additive white input:

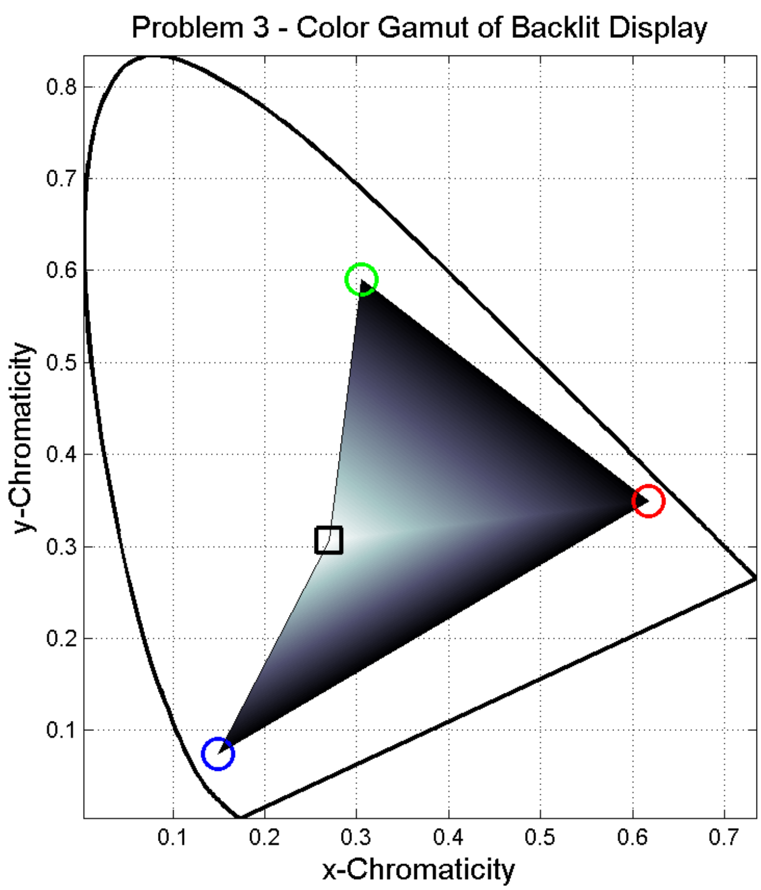

Effectively, the backlight limits the range of possible colors that our monitor can display.