…a journey through the magic realm of color spaces, gray shades, color temperatures, just noticeable differences and more

Goran Stojmenovik
Product Manager Barco Control Rooms division

The whole story behind display specifications and human vision (2)

Abstract

Our world and the displays representing it are not black and white, as someone might think when seeing a display spec consisting only of luminance and contrast. There is much more, there is a whole world of gray shades, color spaces, color temperatures, contrast sensitivities, just noticeable differences. Knowing what “brightness” and “luminance” are and being able to distinguish them, knowing what influences on-screen contrast of a display, and knowing how the eye adapts to changes in luminance levels – these are factors you should have learned from the first part of this white paper. Here, we will go into deeper waters of human vision, to explain when we perceive a display as a good display, and what are the requirements posed by the human visual system to create a good display.

Display requirements and human vision

Gamma

THE GAMMA DETERMINES THE BRIGHTNESS OF THE MID-GRAY SHADES

When we talk about brightness and contrast, we cannot skip the famous "gamma." For one specified contrast, we know it's the brightest white and darkest black that matter. But, what about all the in-between? This is the region of the gray shades. They are produced by tilting the liquid crystal molecules with a voltage (in LCD displays), or by modulating the amount of time a certain DMD mirror is on during one frame (in DLP projection displays). The amount of the "stimulus" that produces the gray shade is called the gray level, such as a changing voltage over an LCD display. The ordered set of the produced gray shades is the gray scale (Fig. 1).

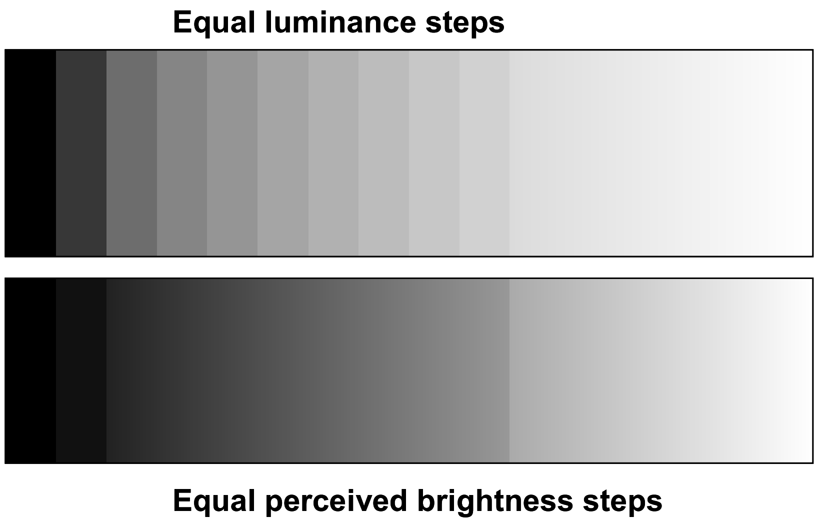

Fig. 1

Now, what is a gray shade? The luminances of the white and black levels don't tell us anything about how the gray shades are organized. We may think, well, they should be uniform. But, uniform in what? Gray level, luminance or perceived brightness? We saw before that luminance and brightness are not the same, and nobody guarantees that equal stimulus (such as voltage) steps will have any linear effect on the brightness or luminance. That's why display technology chooses equal brightness steps for creation of a gray scale. Take a look at Fig. 1. The upper gray scale is built with equal luminance steps (as would be measured by an instrument). The lower scale is with equal brightness steps (as perceived by a viewer). If we are asked to find the gray that is half way between black and white (the so-called 50% gray) on the bottom scale it would be approximately in the middle, whereas on the top scale it would be toward the left.

The infamous gamma is just a number that describes how fast the gray shades change from black to white.

This is of course linked to the perceived contrast of an image (say a photograph, or anything with lots of different colors and shades). Because white remains white and black remains black depending on the gamma, the mid-tones will be brighter or darker, and of course the contrast ratio between the full white and these mid-tones will be different. So, many displays today have a gamma control that can be adjusted depending on the display application.

Contrast sensitivity

WHAT SHOULD THE CONTRAST AND SPATIAL RESOLUTION BE IF WE WANT TO DISCERN DETAIL?

Now we go a step deeper. What is the minimum contrast necessary for the eye to distinguish between two small objects (two pixels, or two lines with different brightness)? From experience, you know there is no straight answer to this question: it will depend on the object size, luminance, distance from the viewer and probably some other factors.

So let's start reducing the question. The object size *h* and the distance *d* from the viewer can be combined in one quantity: visual angle $\alpha=\tan(h/d)$ (Fig. 2). Most experiments in resolving small features have been done with periodic gratings (subsequent dark and bright lines with sinusoidal modulation of luminance). We talk about a cycle, meaning one dark and one bright line with sinusoidal transition between them. The "spatial frequency" of the grating can be expressed in terms of cycles per degree of visual angle, which in the above case is $2h/\alpha$.

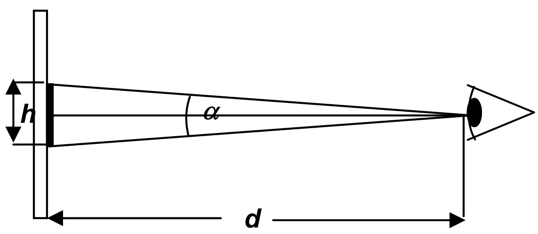

Fig. 2

Of course, this value depends on the number of cycles per millimeter on the display, as well as on the distance from the viewer, and both matter to the eye at the same time. So our question is reduced to: what is the spatial frequency (in cycles per degree) of a grating that we can distinguish as separate lines? It turns out that there is no one such value, but that this threshold spatial frequency depends on the contrast, and vice versa: the lowest distinguishable contrast depends on the spatial frequency.

Fig. 3 explains this last claim. It represents a periodic pattern (dark and bright lines) whose density (spatial frequency) increases as you go to the right, and the contrast decreases as you go upwards. The idea is to visualize the line above which you can't see any difference between the stripes (everything is gray). This line is the contrast threshold that depends on the spatial frequency. You will notice that this line is highest somewhere in the middle, and then drops again as you go to higher spatial frequencies. We say that the contrast sensitivity of the eye (the reciprocal of the contrast threshold) is highest in this mid-region of spatial frequencies (typically around 2-4 cycles per degree). Of course, the spatial frequency depends on the viewing distance. If you move closer and farther from this page, you will see that this contrast sensitivity curve changes shape and the maximum changes position.

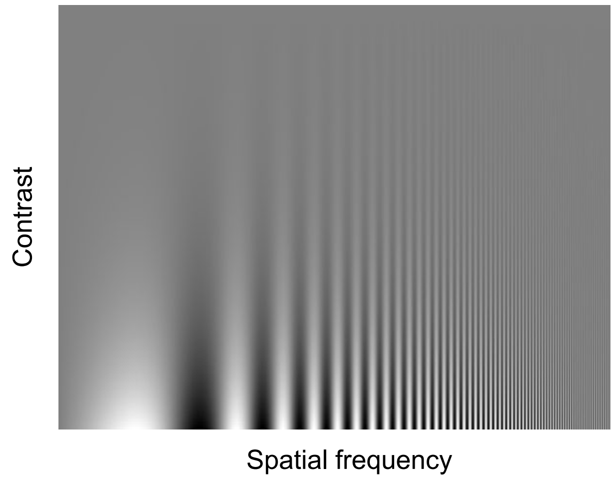

Fig. 3

Now it is clear where we are heading. People who build displays and make applications for them have to know what we can see and what we can't when looking at the display. Knowing the eye's limit for a given application surely helps avoid going into expensive overkill with fancy display features that don't contribute anything.

Colors

Colors are something wonderful, we all agree. Imagine the world looking like a film noir from the '40s, where all we have is different shades of gray. The color is connected to the property of the light called wavelength. We all have at least the vague notion that a color of an object is a bunch of light rays with one or many different wavelengths (which, in turn, originate from the different chemical, physical and biological structures of the object). However, in the world of displays it is impossible to reproduce these colors dynamically in terms of wavelength, because this means that we have to reproduce the object's structure.

Luckily, the human visual system helps. Our eye has photoreceptors that respond with maximums in the red, green and blue region of the wavelength spectrum. And scientists have found out that, within a certain region of colors, you can mix these three 'primary colors' or "primaries" (red, green and blue -– RGB), in different ratios (with different luminances), and still give the eye the impression that it sees a color consisting of a spectrum of wavelengths. In other words, we can cheat the eye.

This is exactly what is being done in display technology. Just take a close look at your CRT TV at home (if you still have one). An LCD pixel also consists of three sub-

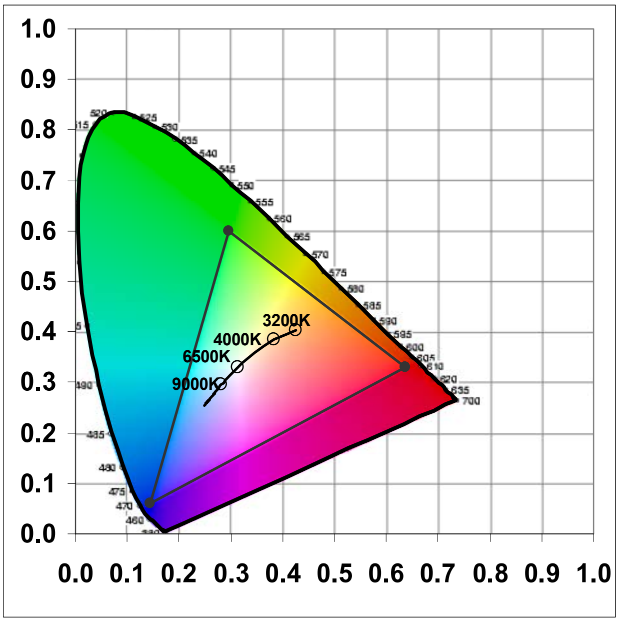

Fig. 4

pixels -- red, green and blue. The luminance of each can be controlled to obtain "gray shades," but now with a certain color. For example, the darkest shade of each color will be black, but the brightest will be full red, green, or blue, respectively. Because these sub-pixels are so small, our eye mixes their colors and creates an impression of one pixel with one color. How many colors can we create in this way? Well, if you have 256 different gray shades within one sub-pixel, there will be 256x256x256=16.8 million color combinations! This is where this fantastic number comes from.

Now look at Fig. 4. This is a representation of colors in terms of two coordinates, *x* and *y*, which we call a color space. All the colors that the human eye can see are shown in the curved colored region. Examples of a red, green and blue primary (the sub-pixel colors) of a certain display are represented on this chart, and they form a triangle. By mixing 256 shades (from black to the full color) of each primary, we can reproduce 16.8 million colors that are scattered within this color triangle (gamut). We see that the wider the gamut (triangle), the larger the color range a display can cover.

Mixing the three primaries with full gain (maximum luminance), one obtains white. However, this white is badly defined, because we can imagine color primaries varying from display to display, which results in different white points for various displays. That is why we define the white point in terms of *color temperature*.

Color temperature

Say we have a white piece of paper and look at it outside on a bright sunny day. Now imagine the same paper when we look at it indoors, under a lamp, or even under candlelight. The same piece of paper will look differently white, no doubt about that. It is possible to connect the "color" of this white to the temperature of the light source: the temperature of the sun surface is around 6500 K (degrees Kelvin), so we say that daylight (when reflected from a white, matte paper) has a "color temperature" of 6500 K. Tungsten lamps would respectively be heated to about 3200 K, so the color temperature of the tungsten white light will also be 3200 K. Now, this is a quantitative way to quantify the different whites we have, instead of saying "warm" or "cold," "reddish" or "bluish."

Looking back at display technology, we have seen that we can cheat the eye by mixing only three primary colors instead of a whole spectrum in order to create a wanted color, or white. The white created in this way can also be said to have a certain color temperature. We can also call it "the white point" or "white balance," meaning that all colors are balanced with reference to the chosen color temperature of white.

But this color temperature is display-specific, since the color filters used depend on the display and technology. Mixing the R, G and B with nicely adjusted gains (luminances) can produce white points with previously selected color temperatures, and this can be done for different displays or display technologies independently. This ensures that the white balance on one display is the same as the white balance on another display. All the possible white points lie on a curve called the "Planckian locus," which is shown in Fig. 3. There we can see that 3200 K white balance is in the reddish region and 9000 K in the bluish one.

Why, then, would different color temperatures be important for a display? Well, say you have a studio environment that is illuminated with tungsten lights. The objects in the studio will consequently have a 3200 K white balance. The cameras will record this, and will then convert it to a daylight 6500 K for us to view. But a display usually radiates its own light. Imagine this being at 6500 K daylight white point. When the camera records this (with the 3200 K environment), and then converts it to a higher color temperature, the display will inevitably look bluish. To avoid this, the white balance of the display has to be tuned to the environment, and this is 3200 K in broadcast studios.

Just noticeable differences (JNDs)

The last thing we want to clarify in this white paper is the notion of color differences. This is also the most complicated part. But if you have come this far, you might as well read it.

Color spaces. The *x*-*y* graph in Fig. 4 represents a color space, but we notice that there is no black. This is because luminance is not represented. However, the interplay between luminance and colors is important, because we don't perceive a same color equally when it has different luminance levels. So people have created color spaces that contain both luminance (L) or perceived brightness (L^*) and a pair of color coordinates (called fancy names as $[u',v']$, or $[a^*,b^*]$, or $[u^*,v^*]$). Now, in such a color space, when we talk about color difference, we mean a "distance" between two points that combine luminance and color. To avoid writing color/luminance below we will say color, but will mean *both* color and luminance.

Just noticeable differences (JNDs). As its name states, one just noticeable difference is the distance between two colors (in a given color space) that the human eye can barely notice. Two colors closer than 1 JND would then be indistinguishable. Of course, we have to soften this requirement, since not all people have the same sight, and even one person can sometimes see this difference and sometimes not. So we say that 1 JND is the difference between two colors that can just be noticed in 50% of the cases (or by 50% by the subjects).

So far so good. But how do we translate this into more quantitative terms? Well, we use a color space, and then it is relatively easy to express the color difference as the distance between two color points. But which color space? The *x*-*y* in Fig. 4 doesn't treat luminance. Furthermore, experiments have shown that the actual length of an observed 1 JND is not uniform in this color space, meaning that the distance between two just distinguishable greens is much larger than between two blues.

So we need a much more uniform color space, where a just noticeable difference between two colors will be the same for all (or most) colors. An example is the so-called $L^*u^*v^*$ system. L^*, u^* and v^* are connected to the luminance L and color coordinates u' and v' (derived from *x* and *y*, see formulas below). The L^* is actually a measure of the perceived brightness, as represented in Fig. 2. The color difference (ΔE) between two colors with coordinates L^*, u' and v' in this new system is calculated with the following equation (remember Pythagoras?):

$$\Delta E = \sqrt{(L_1^* - L_2^*)^2 + (u_1^* - u_2^*)^2 + (v_1^* - v_2^*)^2}$$

A value of $\Delta E = 1$ represents 1 JND, and this is fairly uniform for all luminances and colors. The rest of the formulas necessary to calculate the color difference are given below. L_n, u_n' and v_n' are the reference luminance and color coordinates of the full white of the display in question. So, what you need to calculate ΔE are the luminance *L*, the *x* and *y* color coordinates of the two adjacent colors and those of the reference white.

$$L^* = 116\left(\frac{L}{L_n}\right)^{1/3} - 16$$

$$u' = \frac{4x}{-2x + 12y + 3}, \quad v' = \frac{9y}{-2x + 12y + 3}$$

$$u^* = 13L^*(u' - u_n'), \quad v^* = 13L^*(v' - v_n')$$

A final note: We now see that the human eye cannot discern an infinite number of colors. The value of 16.8 million colors used in most of today's displays is more than enough to cover the color triangle (gamut) with points that are much closer than 1 JND, so that we can barely distinguish them, and this is why this value has somehow "stabilized" on the market now. Going into more colors means more engineering, more system requirements and higher cost, while we cannot distinguish those colors anyway.

Summary

Summary and conclusions

In Barco Control Rooms we use the knowledge of the color theory and the human visual system to engineer displays with good visual characteristics that are favorable for the operators. This encompasses known standards of a nicely regulated display gamma, contrast sensitivity, color triangles and color temperatures for different applications, but we also go a step further in new territories. For example, while 1 JND is just noticeable for half of the observers, a few JNDs are tolerable for most people. That means that people may see a difference, but they still value a display (or two displays next to each other for that matter) as good. The question here is how many JNDs are tolerable? Knowing this saves unnecessary expenditures in engineering the ideal of unnoticeable differences, and produces a good and yet affordable product.

References and further reading

International standard ISO 13406-2: Ergonomic requirements for work with visual displays based on flat panels – Part 2: Ergonomic requirement for flat panel displays

Contrast Sensitivity of the Human Eye and its Effects on Image Quality, Peter G. J. Barten, SPIE Optical Engineering Press, 1999

Color Science: Concepts and Methods, Quantitative Data and Formulae, G. Wyszecki, W. S. Styles, Wiley Interscience 2000

http://www.sid.org/publications/bookstore.html