Characterizing and Controlling the

Spectral Output of an HDR Display

Ana Radonjić, Christopher G. Broussard, and David H. Brainard

Department of Psychology,

University of Pennsylvania, Philadelphia, PA 19104

August, 2011

This report provides technical detail related to the calibration of the an HDR display used in the Brainard lab.

If you would like to cite this work, please use the following:

> Radonjić, A., Broussard, C. G. & Brainard, D. H. & (2011). "Characterizing and Controlling the Spectral Output of an HDR Display," UPenn Brainard Lab Tech Report 2011-1, Department of Psychology, University of Pennsylvania, Philadelphia, PA. http://color.psych.upenn.edu/brainard/papers/hdrcharacterize.pdf.

Display

The basic design of the HDR display is provided in several publications. See for example Olkkonen, M. and D. H. Brainard (2010). "Perceived glossiness and lightness under real-world illumination." Journal of Vision 10(9:5). These procedures were not implemented at the time that paper was published however. They were developed as part of implementing studies of lightness perception in high-dynamic range images that used the same hardware. MATLAB code that implements these procedures is available upon request.

Basic Method

Denote the spectrum coming from a location of interest in the back display (projector) with B(λ, b), where b is a column vector that specifies the settings of the back display for that location.

$$b = \begin{bmatrix} b_r \\ b_b \\ b_g \end{bmatrix}.$$

Denote the filtering applied by the front display (LCD Panel) at the location with F(λ, f), where f is a column vector that specifies the settings of the front display at that location.

$$f = \begin{bmatrix} f_r \\ f_b \\ f_g \end{bmatrix}.$$

Then, we can express the spectrum coming from the HDR display to the observer's eye from the location of interest as:

$$C(\lambda,b,f) = B(\lambda,b)F(\lambda,f) + S(\lambda) \quad (1)$$

where $S(\lambda)$ is scattered light. Note that scattered light is different from the ambient light due to the back [defined as $B(\lambda,[0,0,0]')F(\lambda,[1,1,1]')$] or front [defined as $B(\lambda,[1,1,1]')F(\lambda,[0,0,0]')$] displays themselves. Ambient light is measured and accounted for as part of the characterization of the front and back displays, using standard methods.

In the actual display, there are a number of possible sources for the scattered light. For example, light could scatter from the surrounding image within the diffuser and emerge from the location of interest. Or, it could scatter from any location on the display off the interior of the black enclosing box back onto the LCD and then to the observer's eye. We have made some measurements to characterize this, which are not described here. We assume that $S(\lambda)$ depends on the surround of the region of interest (which it definitely does) and is independent of the display settings at the region of interest (which is an approximation). Because of the dependence of $S(\lambda)$ on the surround, we calibrate separately for each choice of surround when it is feasible in a given experiment.

In the calibration procedure, we characterized the back display by measuring the spectrum coming from the HDR display, as a function of b, when all pixels of R, G and B channel of front display were set to maximum level (corresponding to the maximal transmission through the panel).

$$B_{cal}(\lambda, b) = B(\lambda, b)\, F(\lambda, [1,1,1]') + S(\lambda), \quad (2)$$

Also, we measured the light coming through the full system for various R, G, and B settings of the front display when the all projector pixels are set to maximum input value (full light output).

$$F_{cal}(\lambda, f) = B(\lambda, [1,1,1]') F(\lambda, f) + S(\lambda). \quad (3)$$

Rewriting Equation (1) in terms of the calibration data provided by Equations (2) and (3) allow us to compute the light $C(\lambda,b,f)$ emitted from the overall HDR display for any combination of back and front settings.

$$\begin{aligned} C(\lambda,b,f) &= B(\lambda,b)F(\lambda,f) + S(\lambda) \\ &= \left(\frac{[B_{cal}(\lambda,b) - S(\lambda)]}{F(\lambda,[1,1,1]')}\right)\left(\frac{[F_{cal}(\lambda,f) - S(\lambda)]}{B(\lambda,[1,1,1]')}\right) + S(\lambda) \\ &= \frac{[B_{cal}(\lambda,b) - S(\lambda)][F_{cal}(\lambda,f) - S(\lambda)]}{F(\lambda,[1,1,1]')B(\lambda,[1,1,1]')} + S(\lambda) \\ &= \frac{[B_{cal}(\lambda,b) - S(\lambda)][F_{cal}(\lambda,f) - S(\lambda)]}{B_{cal}(\lambda,[1,1,1]') - S(\lambda)} + S(\lambda) \end{aligned} \quad (4)$$

Note that in the last expression in Equation (4) we could also use the quantity $F_{cal}(\lambda,[1,1,1]') - S(\lambda)$ in the denominator.

Our goal was to develop the algorithm that will compute the settings *b* and *f* that will produce a stimulus that has specified tristimulus coordinates *t*, where the entries of this vector provide the individual coordinates. (The same methods work for any choice of tristimulus coordinates, as long as the corresponding color matching functions are known.) The k^{th} entry of *t* to the stimulus produced by settings *b* and *f* is given by the expression

$$t_k = \int_\lambda T_k(\lambda)C(\lambda,b,f)d\lambda \quad (5)$$

where $T_k(\lambda)$ is the spectral sensitivity of the k^{th} color matching function. We can use matrix/vector form to represent computation of tristimulus values, using standard conventions.

$$t = \mathbf{T}\frac{[\mathbf{B}_{cal}(b)-s][\mathbf{F}_{cal}(f)-s]}{\mathbf{B}_{cal}([1,1,1]')-s}+\mathbf{T}s \tag{6}$$

It is convenient to handle the scattered term simply by subtracting it from the desired tristimulus coordinates, defining $t' = t - \mathbf{T}s$ and writing

$$t' = \mathbf{T}\frac{[\mathbf{B}_{cal}(b)-s][\mathbf{F}_{cal}(f)-s]}{\mathbf{B}_{cal}([1,1,1]')-s}\,. \tag{7}$$

It is also convenient to similarly define $\mathbf{B}^{'}_{cal}(b) = \mathbf{B}_{cal}(b) - s$ and $\mathbf{F}^{'}_{cal}(f) = \mathbf{F}_{cal}(f) - s$. This then gives us

$$t' = \mathbf{T}\frac{\mathbf{B}^{'}_{cal}(b)\mathbf{F}^{'}_{cal}(f)}{\mathbf{B}^{'}_{cal}([111]')}\,. \tag{8}$$

To find the desired settings *b* and *f*, we need to invert Equation (8), given t' and the calibration measurements. There are a number of options for how to do this, since there may be many pairs of settings *b* and *f* that yield the same (or very similar) t'. Here we proceeded in two steps. First, we used a simple heuristic to find settings *b* such that all entries of *b* were equal to each other and allowed for a set of in-gamut settings *f* that produced t' through Equation (10). This heuristic was as follows.

- First, we asked what settings f_0 would be required for the front LCD display, if the back display were set to its maximum input values ([1,1,1]'). This calculation was performed using the front calibration data and standard methods. We then used the gamma measurements for the front display to obtain linearized versions of the settings f_0^{lin}. By convention, these linearized settings are normalized so that 1 represents the maximum amount of light that can be obtained from each of the front panel's color channels.

- We found the maximum entry of f_0^{lin} and took its square root. This provided a factor between 0 and 1 that indicated roughly how much we should attenuate the back display (relative to when it is set to maximum input values ([1,1,1]') to portion the overall attenuation between the front and back displays.
- We scaled the desired stimulus (*t'*) by this factor and used the back calibration data to find the linearized settings b_0^{lin} that would produce this scaling. We gamma corrected these to find the RGB settings that produced this linear attenuation in each channel. Since these did not necessarily have the property that all three settings were the same as each other, we averaged the obtained settings values and constructed the settings vector b such that all of its entries had this value.[1]
- Given the choice of b we could then use Equation (8) together with standard methods to obtain the front settings f that produce t', given the back settings b. In particular, to obtain spectral calibration measurements for the front display from $\mathbf{F}_{cal}^{'}(f)$ for the case where the back display is set to settings b, note that we need only to multiply each calibration measurement by the biasing spectrum

[1] In our code, this method is called "backequalsettings". We tested other methods as well; they are also available in the code.

$$\mathbf{B}^{'}_{bias} = \frac{\mathbf{B}^{'}_{cal}(b)}{\mathbf{B}^{'}_{cal}([111]')}. \qquad (9)$$

Estimating scattered light $S(\lambda)$

When both b and f are set to zero, there is still some light that reaches the observer from the corresponding location in the HDR display. Some of this light arises because the back projector still emits some light when its input settings are zero, and some of this light is transmitted directly to the observer even when the front display settings are zero. Some, however, is scattered light $S(\lambda)$. For example, when an spatially extended stimulus is presented on the HDR display (e.g., a white background surrounding a test location), some of this light scatters off of surfaces in the display chamber and the front panel, and eventually reaches the observers' eye. To apply the calibration procedure above, we need to estimate $S(\lambda)$. It is difficult to do this through direct measurements because $S(\lambda)$ is too small to be measured directly by our PhotoResearch PR-650 spectral radiometer, at least for some of our experimental conditions. None-the-less, the effect of $S(\lambda)$ light can be observed as deviations of measured spectra from predictions of such spectra made on the assumption that $S(\lambda) = 0$, for input values greater than 0 (Figure 1A).

To estimate $S(\lambda)$, we chose a set of grayscale stimuli and derived front and back settings on the assumption that the amount of scattered light $S_0(\lambda) = 0$. We then measured these stimuli. Because the sensitivity of the PR-650 radiometer is limited, we could not obtain measurements for stimuli with luminances below ~0.07 cd/m^2. We then compared the measurements to the predictions from the calibration data.

Figure 1 (panel A) shows the comparison for measurements for two different experimental conditions. In one, the test patch positioned in the center of the display, which we measure in the course of the calibration, is embedded in a 5x5 checkerboard (the checkerboard surround condition). In another, the test patch is embedded in a uniform white surround (the white surround condition). Note that predictions are made separately for the full and white conditions, as separate calibration measurements for each were obtained with the background set for those conditions. Also note that the predictions were made after application of the correction for display non-additivity described below.

There are deviations from the predictions at both the low and high ends of the luminance range. We posit that the deviations at the low end are due to the scattered light $S(\lambda)$. We chose by eye a range of n lowest luminance stimuli from which to derive $S(\lambda)$: we used the stimuli at the low-luminance end of the range for which the deviations from predictions were judged to be consistent in magnitude and direction. $S(\lambda)$ was then obtained by using numerical parameter search to find the spectrum that when added to the predictions brought the result as close as possible to the n measurements. The resultant value of $S(\lambda)$ was then used in the procedures described above (see Equations 1-4) to obtain the settings for the experimental stimuli.

We do not know the source of the deviations from the predictions at the high end of the luminance range, but these deviations are not of critical importance because we can directly measure these stimuli (which we do when we report our experimental data).

Display Non-Additivity

The individual front and back displays in our apparatus exhibit a certain degree of additivity failure (that is, the light emitted for combinations of settings is not always well predicted from the measurements for each channel alone.) To minimize the effects of this, we developed procedures for calibrating the displays using settings as close as possible to those that would be used in an experiment. This improves our ability to predict the stimuli from the calibration measurements. This procedure would not be needed if the components of the HDR display were better behaved.

We implemented the procedure in two steps. The first is simple. For each channel, we characterized the gamma curve by measuring the spectra at different intensity levels for that channel, while keeping the remaining two channels set to a constant intensity, approximately in the middle of the range. In other words, to obtain a gamma curve for a red channel, we first measured the red channel at a set number of levels r between minimal (0) and maximal (1) intensity, while the blue and green channel were set to a fixed intensities of g_o and b_o. Call these measurements $R_{gam}(\lambda,r;g_o,b_o)$. From each of these we subtract measurements $R_{gam}(\lambda,0;g_o,b_o)$ to obtain the incremental effect of the red channel around the operating point defined by g_o and b_o. A similar procedure provides the measurements for the green and blue guns. We fit these measurements with smooth sigmoidal curves derived from the cumulative distribution of the beta probability distribution, which captured the particular shape of our device's gamma curves reasonably well.

The above procedure provides a first order correction for gun interactions, but does not capture effects caused by large deviations from the calibration operating point. To do this for our grayscale stimuli, we thus implemented a second step. We computed the settings that were required for the front and back displays for the set of grayscale stimuli we wished to produce. This depends both on choice of stimulus chromaticity and stimulus luminances. This gave us a set of settings for both displays. For each display, we measured the stimuli corresponding to the computed settings, while the other display was set to its maximum. We will refer to these measures as yoked, because they are identical to our desired stimuli measurements. We then fit each measurement as the sum of the three channel primaries and extracted the weight required for each channel in this sum. These weights were then used as the gamma functions for each channel. They represent how much of each channel we get when the other two channels are at the values used in the experiments. We also used linear model methods to correct for variation across input levels in the relative spectrum produced by each channel. Again, this would not be necessary if the component devices were better behaved. For additional detail on these procedures, see the code itself.

We refer to this whole second step in the course of which the gamma functions are refit using the yoked measurements as "yoked calibration". The settings for the test stimuli are computed from the gamma functions, which are result of yoked calibration. To maximize predictive performance, we performed the yoked calibration in the context of surrounding stimuli similar to the experimental stimuli.

The above step could be iterated, but we did not find additional improvement when we tried that.

Overall Predictive Performance

Figure 1 shows the predictive performance of our algorithm for a set of 24 stimuli in the checkerboard surround and white surround conditions: deviations of measures (circles) from predictions (solid lines) are plotted for stimulus log-luminance (panels A and B), x (panel C) and y chromaticity (panel D) for the checkerboard surround (in red) and the white surround condition (in green). The comparison of log-luminance deviations between panel A (no scattered light assumed) and B (with estimated scattered light) shows far better agreement between measures and predictions when scattered light is estimated and taken into account when computing the settings. The chromaticity plots also show predictions when scattered light is taken into account. The quality of the fit of the scattered light predictions is not always this good, but we do believe that the estimation procedure leads to better predictions than omitting it.

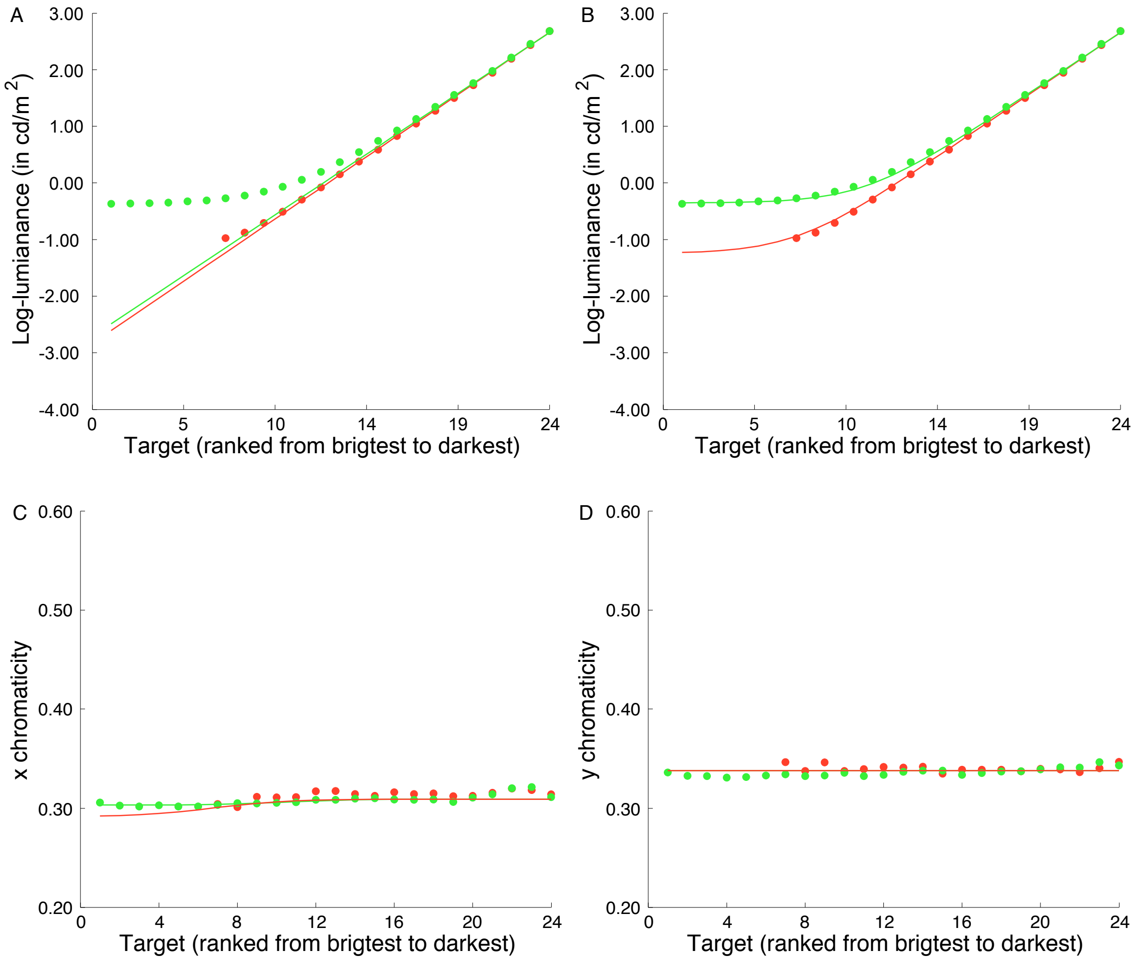

Figure 1. Deviations of measured from predicted test log-luminance and chromaticity for checkerboard and white surround conditions. Luminance and chromaticity is measured for a set of 24 test stimuli in each condition when settings are derived from a yoked calibration. Test stimuli are ranked in the ascending luminance order and plotted on x-axis (1 is the darkest, 24 is the brightest test). The y-axis plots log-luminance in cd/m^2 (A, B), x (C) and y (D) chromaticity. Scatter light is set to 0 in Figure A and to best estimate of scattered light in Figures B-D. The circles represent the actual stimulus measures while the solid lines represent the predictions. White surround condition is shown in green, checkerboard surround in red. Six darkest stimuli in the checkerboard surround condition are not plotted because they were below PR-650 colorimeter luminance range.