Face-based Luminance Matching for Perceptual Colormap Generation

Gordon Kindlmann*
School of Computing
University of Utah

Erik Reinhard
School of Electrical Engineering and Computer Science
University of Central Florida

Sarah Creem
Dept. of Psychology
University of Utah

ABSTRACT

Most systems used for creating and displaying colormap-based visualizations are not photometrically calibrated. That is, the relationship between RGB input levels and perceived luminance is usually not known, due to variations in the monitor, hardware configuration, and the viewing environment. However, the luminance component of perceptually based colormaps should be controlled, due to the central role that luminance plays in our visual processing. We address this problem with a simple and effective method for performing luminance matching on an uncalibrated monitor. The method is akin to the minimally distinct border technique (a previous method of luminance matching used for measuring luminous efficiency), but our method relies on the brain's highly developed ability to distinguish human faces. We present a user study showing that our method produces equivalent results to the minimally distinct border technique, but with significantly improved precision. We demonstrate how results from our luminance matching method can be directly applied to create new univariate colormaps.

CR Categories: I.3.3 [Computing Methodologies]: Computer Graphics—Picture/Image Generation I.3.4 [Computing Methodologies]: Computer Graphics—Graphics Utilities I.4.10 [Computing Methodologies]: Image Processing and Computer Vision—Image Representation

Keywords: Colormaps, Color Scales, Isoluminance, Brightness Matching, Perceptually-based Visualization

1 INTRODUCTION

Visualization tasks routinely involve the mapping of some aspect of the data to a color scale or colormap. In most cases, the luminance of the colormap should be properly controlled. Luminance is an important quantity in visualization because it plays a central role our perception of image structure and surface shape [16]. When applying univariate colormaps to visualize data on a flat two-dimensional domain, luminance can be used to enhance display of large-scale structural composition and variation. Or, luminance can be held constant so as to minimize interpretive errors caused by perceptual effects such as simultaneous contrast [25]. In the context of bivariate colormaps using separate axes for color and "brightness", carefully controlling luminance can help maintain orthogonality between the visual representations of different data components [21]. A similar constraint applies when using univariate colormaps in combination with surface shading, since (with the exception of specular highlights, and assuming white lights) variations in reflected light due to changes in surface orientation are not variations in color, but only overall intensity [10]. Misleading cues of

surface shape could come from luminance variations in the univariate colormap itself.

Exerting control of luminance in colormap-based visualizations is an interesting problem, due to at least three confounding issues. Most importantly, the display device tends to be uncalibrated (proper calibration would require an external measurement device [7]). The chromaticities, intensities, and response functions of the primary colors are often not known, and can vary significantly between display devices [13]. Also, the lighting conditions and configuration of the room are often unknown (or uncontrolled), contributing to factors such as light reflecting off the display device surface, and differences in brightness and color perception caused by variations between foveal and peripheral luminous sensitivity [26]. Finally, yellow pigments in the ocular media such as the lens and the macular area of the retina, can cause non-trivial differences between individuals in spectral sensitivity [26]. Our experience has been that there is so much variation between monitors that it is not sensible to simply assume a "standard" monitor and then work in a CIE colorimetric space such as XYZ, or the approximately perceptually uniform spaces CIELAB and CIELUV.

We address the general problem of colormap luminance control by proposing a novel technique for luminance matching. Given a fixed reference color, and a test color with lightness varied by a user interface, our technique facilitates matching the luminance of the two colors. The technique is based on the brain's special capacity to recognize images of human faces. By picking a fixed gray level, we can adjust the intensity of control points in a univariate colormap to equalize their luminance, thereby creating an isoluminant colormap. It is also simple to create univariate colormaps with other patterns of luminance variation, such as monotonically increasing. Our approach works on any display device with additive color that obeys a power-law response function, such as a CRT.

The remainder of the paper is as follows. After a review of previous work, the rationale for our luminance matching technique is explained in Section 3, in relation to a previous method called "minimally distinct boundary". Section 4 presents the interface and usage of the new method, as well as a demonstration. To validate our method against minimally distinct boundary (that is, to show that we measure the same quantity, but do so more precisely), we conducted a user study, described in Section 5. Given the ability to match luminance for individual colors, we describe in Section 6 how to interpolate between colormap control points while controlling luminance, so as to produce colormaps with constant or monotonically increasing luminance. Conclusions, discussion, and directions for future work are given in Section 7.

2 PREVIOUS WORK

There is extensive literature on the generation and application of perceptual colormaps, especially concerning the relationship between axes of color perception and the representation of the underlying data [3, 18–20, 25]. In general, a controlled or standardized system for visualization display is assumed, i.e., calibrated monitors, standard observers, and/or controlled viewing environments. Our work aims to relax these constraints without invalidating cur-

---

*Contact information: SCI Institute, University of Utah, 50 S. Central Campus Dr. #3490, Salt Lake City, UT 84112. gk@cs.utah.edu.

rent practice. The ability to make accurate comparisons of luminance is useful for correctly implementing any perceptual colormap in an uncalibrated environment.

Using humans' highly developed skills in recognizing faces for specific tasks in visualization is a relatively new concept. The work closest to ours is called the "Which Blair Project" [22]. Here, the user evaluates the monotonicity of a given colormap's luminance by applying the colormap to a continuous grayscale image of a face. By doing this a number of times, using a sequence of small segments from the colormap domain, a sequence of colormapped images is created. If in all the images a positive face is seen, the colormap is said to be monotonically increasing in luminance. If some of the images are viewed as positive and some as negative, the luminance of the colormap is not monotonically increasing, but contains reversals. This method of colormap evaluation is useful in part because it does not require any display calibration.

Based on the success of this method, we believe that face recognition can be used not only for *identifying* monotonically increasing colormaps, but also for *constructing* colormaps with predetermined luminance variation. Specifically, we use a face image to indicate the luminance equality of two colors, instead of the luminance monotonicity of a colormap.

3 BACKGROUND

Our approach to controlling the pattern of luminance variation within a colormap is derived from earlier approaches to solving a basic problem of color vision, namely, measuring the the *luminous efficiency function* $V(\lambda)$, which describes the sensitivity of the eye to light at different wavelengths [26]. Given $V(\lambda)$, luminance L is defined as:

$$L = K_m \int_\lambda L_{e\lambda} V(\lambda) d\lambda \qquad (1)$$

where $L_{e\lambda}$ is the spectral concentration of radiance, and K_m is a constant fixed at 683 lm/W. Often, the CIE 1924 standard photopic luminous efficiency function is used for $V(\lambda)$, but it should be noted that this function can vary somewhat across individuals [1, 2]. Luminance is a photometric quantity dependent on the amount of radiance reaching the eye, and the eye's sensitivity to each wavelength.

Many of the techniques traditionally used to measure the luminous efficiency function are based on matching paradigms. That is, the measurement consists of having participants vary the intensity of a given color until some subjective impression of that color (depending on the specific method) is judged to match a fixed, reference white. Repeated many times throughout the visible spectrum, a model of spectral sensitivity is acquired. Conversely, in the context of visualization, we would like to produce for the observer a predetermined pattern of luminance variation, using the colors in a given colormap, and as displayed on a given monitor, by taking into account the observer's pattern of color sensitivity. Measuring sensitivity to colors produced by CRTs for colormap generation is simpler than measuring spectral sensitivity to monochromatic sources, but we feel our approach can be informed by the characteristics of previous methods in photometry.

Existing psychophysical measurements tend to exploit either spatial aspects of the human visual system (direct heterochromatic brightness matching, minimally distinct border) or exploit the different temporal resolutions of the human visual system's chromatic and achromatic channels (flicker photometry, critical flicker frequency) [5, 24, 26]. The latter category tends to produce reliable and accurate results, but requires flicker frequencies of around 15 to 20 Hz, which are not easy to accurately produce on an average monitor. Moreover, flickering at these frequencies is somewhat annoying to observe for extended periods, and can induce photosensitive epileptic seizures in some individuals, raising important safety issues [14].

For practical purposes, a brightness matching task of some sort is more attractive, where brightness is defined as the dimension of subjective visual experience that correlates most closely with the physical intensity of light [5]. In direct heterochromatic photometry, a bipartite field is presented to an observer, where both sides are chromatic, but of different wavelengths. The observer is then asked to adjust the intensity of one side of the field until it matches the other side in brightness. This task turns out to be rather difficult and inaccurate results of this method are reported. However, the most fundamental flaw is that of additivity failure [5].

Additivity failure can be described as follows. Suppose that in two separate experiments two chromatic fields of saturated color are matched to an identical white field by varying the brightness of the chromatic fields. Then, in a third experiment, the two chromatic fields are optically superposed and the white field is doubled in intensity. If the user had been performing an accurate luminance match, then Equation 1 implies that the combined colors would equal the brighter white in luminance, and therefore the user would find that indeed they visually match. However, the combined chromatic field will generally appear *dimmer* than the doubled white field. Saturated colors appear to "glow" with a brightness out of proportion to their actual luminance; this is known as the Helmholtz-Kohlrausch effect [9, 26]. In general, the strength of this effect increases with the saturation of the color.

Related to direct heterochromatic brightness matching is a method known as minimally distinct border. Again, two fields are compared as above, but now the fields are precisely juxtaposed and one side is adjusted until the border between the two fields is minimally distinct according to the observer. The results of this method tend to be close to those obtained with heterochromatic flicker photometry. Moreover, this method does not suffer from additivity failure [5, 26]. Thus, while the minimally distinct border method may be generically known as a "brightness matching" method, its utility is derived from its apparent immunity to the Helmholtz-Kohlrausch effect, in that it produces measurements which do obey additivity. In this sense, we can consider minimally distinct border to be a *luminance* matching method. However, the method still relies upon observers matching patches with potentially widely differing chromaticities, between which there is always a boundary with some distinctness, even when luminances do match. Thus, minimally distinct boundary is a challenging task, so we expect this method to yield somewhat imprecise results.

Our goal is to create a task that is similar to the minimally distinct border method, in that it is robust against the Helmholtz-Kohlrausch effect, but that is easier and more intuitive for users to carry out. This would simultaneously make the task more practical and may improve the precision of the result. We propose a luminance matching technique based on a thresholded image of a human face. The reason for this suggestion is that humans appear to be very good at recognizing thresholded face images as long as there is an appropriate luminance difference between the representation of shadowed and illuminated surfaces [6]. Section 4 describes our method in detail, and Section 5 describes our validation study.

4 METHOD

Our method starts with a black-and-white thresholded image of a human face. After thresholding, the resulting binary image is black where shadows were cast and white where the face was lit directly. For the thresholded face to look natural, it is best to make the amount of ambient lighting very small, with the direct lighting coming from a single source, as with strong sunlight, or a bright lamp in an otherwise dark room. The test pattern used for our method

Figure 1: Double face image

is simply two copies of this image, placed side-by-side, with one reversed in black and white (Figure 1).

Next we replace black with a shade of gray, and white with a color. Depending on the luminance of the colored region in comparison to the luminance of the chosen gray, either the left or the right face will appear "positive" and the other face will be perceived as "negative" (the left face in Figure 1 is positive). Having two copies of the face image is important: as long as the luminances are unequal, one of the faces will be positive, and it will always "stand out" more than the negative. When the observer adjusts the intensity of either the gray or the color, it is easy to see when the contrast polarity between the gray and color changes, because it is signaled by a change in location of the positive face. In order to perform luminance matching, the intensity of either the gray or color is adjusted to be in the midst of the transition region in which neither face appears positive nor negative. The precision of this method, evaluated in Section 5, depends on the transition region being relatively small.

The rationale behind using face images is that humans are naturally good at recognizing faces, arguably because humans possess dedicated brain circuitry to process faces [4, 8, 11]. This is especially true under the assumption that light comes more from above than from below [17]. Previous work in psychology has shown that luminance is the most important determinant in perceiving thresholded images of faces [6]. Also, sensitivity to incorrect luminance levels within lit and shadowed regions of an image is stronger with faces than with other objects [23]. fMRI studies have localized the region in the brain which responds more strongly to positive faces than negative ones [12].

Changing the luminance of a color can be achieved by varying the lightness coordinate in HLS color space, the double hex-cone of hue, lightness, and saturation [10]. This is equivalent to linearly blending the color with varying amounts of either black or white in RGB space. This approach was chosen because due to device limitations, it is not possible to indefinitely increase a color's luminance without also decreasing its saturation, an immediate result of blending with white. Because we are using HLS space, in this paper we use "lightness" to refer to the L axis in HLS space, instead of the usual connotation of perceived whiteness or blackness of an opaque surface [26].

An illustration of our approach is given in Figure 2. Three different colors are varied in luminance from top to bottom by varying their HLS lightness, while the gray is fixed at RGB (0.5,0.5,0.5).

Positive faces are seen on the right in the top images, and on the left in the bottom images; the transition zone occurs somewhere in the middle. The precise location of the transition varies with the brightness of the color. On a printed page, the transition depends on the illumination spectrum, but this specific effect does not pertain to additive colors on a monitor.

5 USER STUDY

To show the effectiveness of the double face luminance matching method, a user study was conducted to compare this method with a variation of the minimally distinct boundary task. The minimally distinct boundary (MDB) method was chosen for the reasons outlined in Section 3. The MDB method was adapted to be used with a CRT display. A pilot study using the test pattern shown in Figure 3 confirmed that observers were less precise with this method than with the face method.

It could be argued that the reason for this imprecision is due to the difference in border length between the two fields. To eliminate this possibility, instead of using a test pattern with only four squares, we created a new test pattern by flipping the double face image upside down and rearranging connected regions, resulting in the minimally distinct boundary stimulus shown in Figure 4. This leaves unaffected both the border length between gray and color fields, as well as the amplitude spectrum of the resulting image. The main difference between the two test patterns is that one is readily recognized as a face whereas the other is not.

We conducted a user study to test our hypotheses that (1) the face-based luminance matching will produce the same result as the minimally distinct border method, and (2) that the face-based method will have higher precision. The test environment, including the CRT monitor, was kept constant across participants, but was not specifically controlled or calibrated across observers, except to verify that the monitor's black level ("brightness") was set correctly. The study was conducted using 12 participants, four females and eight males, ranging in age from 13 to 40 years, all with normal or corrected-to-normal vision.

Half the observers were presented first with a set of 36 face stimuli, followed by the same number of MDB stimuli. The order of tasks was reversed for the other half. Within each of the two tasks, 6 maximally saturated colors (red, yellow, green, cyan, blue and magenta) were presented 6 times in random order. For each color the initial lightness was set randomly to either black or white to avoid a bias due to starting position. The reference gray was set to RGB (0.5,0.5,0.5). Because the colors' lightnesses were varied to match the fixed gray, the work performed in the user study is similar to the first step in creating an isoluminant version of the rainbow colormap.

Before the experiment started, each observer was allowed a practice session to familiarize him/herself with the software, which consisted of a display of the square test stimulus which measured 14 cm on a side. Observers were positioned at a normal viewing distance from the monitor (roughly 40 cm). Changing the lightness of the color is achieved by clicking on the display and dragging the mouse left or right to increase or reduce the color's lightness. Moving the mouse all the way to the right will make the face on the right become positive, and vice versa. During this operation, the cursor on the display becomes invisible, to avoid giving the participant any cues of where they are on the lightness scale.

During each task (face and MDB), participants were allowed two one-minute breaks. The overall time needed to complete both tasks was between 30 and 40 minutes. The instructions for each task were presented in writing. For the double face task, the main instructions were:

> Depending on whether the gray shade or the color is

Figure 2: Double face stimuli examples to match cyan, magenta, and yellow to a fixed gray. The colors in each row have a common lightness in HLS color space, going evenly (top to bottom) from 0.1 to 0.9. The gray is at lightness 0.5. The brightness is matched when the face switches from the right to the left side.
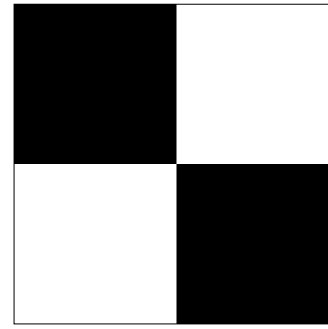

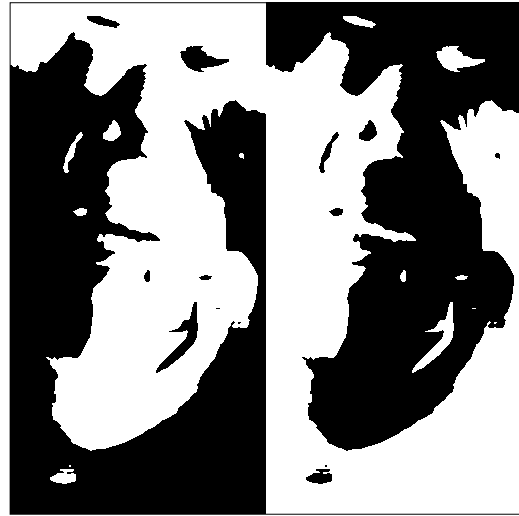Figure 3: Initial test pattern for minimally distinct boundary.


Figure 4: Test pattern used for minimally distinct boundary. Length of boundary between two colors is same as double face test image.

> brighter, the face will appear either on one side or the other. Make sure to explore the range of brightness enough to see the face on both sides. Then, try to find the point in the middle, where the over-all impression of the two faces are equally distinguishable or equally indistinguishable.

The main MDB task instructions were:

> Your goal is to make the brightness of the color and the brightness of the gray equal. If the color is much brighter or much darker than the gray, the boundary between them will be very distinct. Find the brightness in between so that the boundary is minimally distinct.

The data from the user study is summarized in Figure 5, which shows that two methods measured the same quantity, but the double face method had more precision. We performed statistical analysis to compare the two tasks, MDB and double face. In this context, "trial" refers to one luminance match performed by a participant in the study. The value produced by each trial is the HLS lightness at which the color was found to match the reference gray. We calculated the mean and standard deviation (SD) in lightness for each color (across all participants, all trials, and both tasks) to detect outliers. Four trials were removed because their lightness was above or below 3 SDs of the mean lightness.
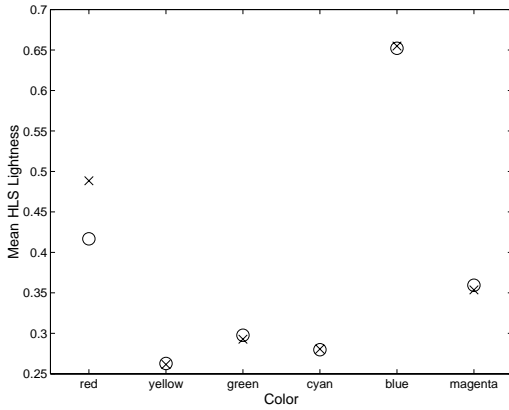
(a) Mean HLS Lightness for each color and task (x: MDB, o: face), averaged across participants.
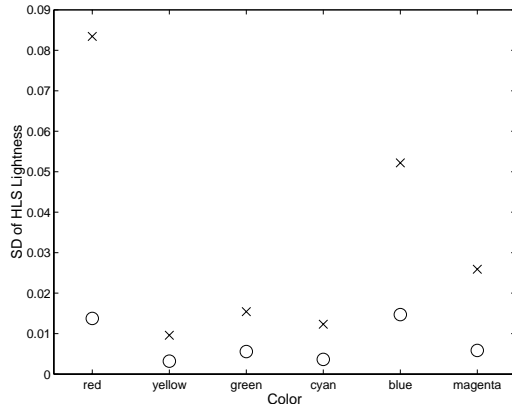
(b) SD of HLS Lightness for each color and task (x: MDB, o: face), averaged across participants.

Figure 5: Mean (a) and SD (b) of lightness values measured in user study. The mean values show the HLS lightness that the color had to be set to, in order to match gray 0.5.

The remainder of our analysis is based on the lightness means and SDs across six trials (or five trials, in the presence of an outlier), calculated for each color and each participant, for both the MDB and double face tasks. We analyzed the means to assess the similarity in measurement between the MDB and double face tasks. We then analyzed the standard deviations to assess the precision, or consistency, of responses.

A 2(task) × 6(color) repeated-measures analysis of variance (ANOVA) was performed on the mean lightness values. The analysis indicated no effect of task (p = 0.29), that is, there was no significant difference in mean lightness measured with the MDB versus the double face methods. Individual paired t-tests for each color (across trials and participants) also confirmed that there was no significant difference between the measurement results of the MDB and double face tasks (red: p > .10; green: p > .25: all other colors: p > .5).

We found that the double face task was superior to the MDB task in the precision of responses as assessed through the standard deviation calculated over the six trials, for each color and participant. In other words, participants were more consistent in the lightness values they chose with the face method than the MDB method, as can be seen in Figure 5(b). A 2(task) × 6(color) repeated-measures ANOVA was performed on the SD values for each color. The anal-

ysis indicated a significant effect of task, $F(1,11) = 20.27$, $p < .001$ and color, $F(5,55) = 8.85$, $p < .001$. Individual paired t-tests for each color confirmed a significant task difference for each color (red: $t(11) = 3.29$, $p < .01$; yellow: $t(11) = 6.13$, $p < .001$; green: $t(11) = 5.48$, $p < .001$; cyan: $t(11) = 3.89$, $p < .01$; blue: $t(11) = 2.29$, $p < .05$; magenta: $t(11) = 4.38$, $p < .001$).

In order to get a rough sense of how long it takes to perform the MDB and double face tasks, the response time was recorded for each trial. This was measured from the first mouse click (to start adjusting the color lightness) until the "Next" button was clicked (to start the next trial). The mean response time was 19.9 seconds for the MDB task and 18.5 seconds for the double face task. The increased precision of the double face method does not seem to come at any cost in increased response time. However, we feel further statistical analysis on response times is not warranted because participants were not given explicit instructions concerning their speed of performance.

The analysis above confirms that the double face task and the MDB task are equivalent in measured output. Although the literature supports the claim that MDB does not suffer additivity failure, it would be reassuring if there was a way of directly verifying whether the lightness levels determined with the double face task were in fact additive, for the reasons discussed in Section 3. Properly testing this would involve a new user study. We have devised a test of additivity, and we have performed an informal study. The additivity test described here is entirely self-contained, in that it requires no knowledge of the monitor's chromaticities or gamma. It uses a version of the face image (Figure 1) in which the shadow color is formed by alternating black and *gray* lines; call the lightness of the gray A. Our additivity test is comprised of six luminance matches:

1. Find A such that the alternating line pattern matches pure RGB blue (0,0,1). Blue is the darkest of the primaries.

2. Find the shade of pure red that matches the alternating line pattern with gray A. Call this (L_r,0,0).

3. Find the shade of pure green that matches the alternating line pattern with gray A. Call this (0,L_g,0).

4. Find *solid* grays A_{rg}, A_{gb}, A_{rb} that match (L_r,L_g,0), (0,L_g,1), and (L_r,0,1), respectively.

Since the gray formed by alternate gray A and black lines should have half the luminance of solid gray (A,A,A) (this being the standard mechanism of visual gamma measurement [15]), the pure blue (0,0,1), the red shade (L_r,0,0), and the green shade (0,L_g,0) should all have half the luminance of (A,A,A), from the matches performed in the first three steps. If we assume that the luminances of independent primaries add linearly, then (L_r,0,1), (0,L_g,1), and (L_r,L_g,0) should all have the same luminance as (A,A,A), so A, A_{rg}, A_{gb}, and A_{rb} should be equal. For five participants we show these values in Table 1. The values deviate from equality by at most 5%. Considering that all colors formed on a monitor are the additive result of these RGB primaries, we believe the three pair-wise tests of step 4 above should ensure additivity between any two colors created on a CRT. In addition, because the primaries are the most saturated colors possible on a CRT, we believe this is a useful test of robustness against additivity failure induced by the Helmholtz-Kohlrausch effect.

6 COLORMAP GENERATION

The task described in Section 5 is only the first step in making an isoluminant colormap. The six principal hues in RGB space were matched with a fixed gray, but we want to produce the continuous

A	A_{rg}	A_{gb}	A_{rb}
0.552	0.557	0.557	0.551
0.593	0.611	0.586	0.614
0.598	0.605	0.600	0.600
0.599	0.606	0.587	0.593
0.613	0.622	0.600	0.604

Table 1: Data from informal additivity test on five participants.

Figure 6: Double face image for gamma measurement

range of hues in between, by interpolation. If we had the luxury of matching a great many control points along the colormap, then the choice of colorspace in which to do color interpolation would not significantly matter. However, the trade-off we encounter, if we aim to perform as few matches as possible, is that we must know the gamma of the display device.

Luminance matching based on the image of the double face can be employed to measure the gamma of the monitor. Specifically, the black in Figure 1 is replaced by a constant gray value which can be adjusted, and white is replaced by alternating black and white scanlines, as seen in Figure 6. This relies on the same principle used in existing gamma measurement images and applets, namely that a gray value created by alternating black and white lines has intensity half that of white, regardless of gamma [15].

Knowing the monitor gamma γ, we can perform interpolation in what is essentially gamma-corrected RGB space. Suppose we have two RGB colors $c_0 = (r_0, g_0, b_0)$ and $c_1 = (r_1, g_1, b_1)$ which have been determined to have equal luminance. These could be, for instance, two of the colors determined as part of our user study. The interpolation between them can be parameterized by $f \in [0.0, 1.0]$, and is calculated by:

$$c_f = \begin{pmatrix} ((1-f)r_0{}^\gamma + fr_1{}^\gamma)^{1/\gamma} \\ ((1-f)g_0{}^\gamma + fg_1{}^\gamma)^{1/\gamma} \\ ((1-f)b_0{}^\gamma + fb_1{}^\gamma)^{1/\gamma} \end{pmatrix} \quad (2)$$

This has the effect of converting RGB component levels to intensity, linearly interpolating, and converting back to RGB component levels.

If we use the data generated by our user study, we can average over all participants and all trials to produce six points along an isoluminant rainbow colormap. These values, and the resulting colormap, are shown in Figure 7.

The methods described thus far can also be applied to the

(a) Isoluminant colormap created by user study

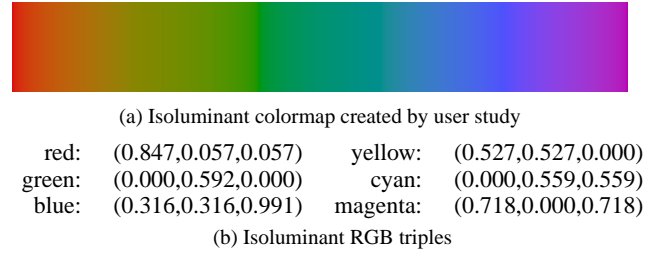| red:	(0.847,0.057,0.057)	yellow:	(0.527,0.527,0.000)
green:	(0.000,0.592,0.000)	cyan:	(0.000,0.559,0.559)
blue:	(0.316,0.316,0.991)	magenta:	(0.718,0.000,0.718)

(b) Isoluminant RGB triples

Figure 7: Isoluminant colormap (a) generated by averaging double face luminance matching data across participants (b), using evenly spaced control points, starting and ending with red. The gamma used for interpolation (2.7) was estimated using the image in Figure 6.

problem of generating colormaps which monotonically increase in luminance, while also varying in hue. Such a colormap combines perceptual benefits from both grayscale and isoluminant colormaps [25]. Instead of adjusting colors (in HLS space) to match luminance with a fixed gray, we can specify a different gray level for each colormap control point. Equation 2 is again used to interpolate in a way that controls luminance, but now luminance is linearly increasing between control points. The sequence of luminances chosen for the control points can increase linearly, or according to a power law that accounts for the non-linearity of brightness perception [26]. Figure 8 shows a colormap produced by one of the authors, by sampling the standard rainbow colors (going from magenta through blue and green to red), and matching against lightness increasing linearly from 0.0 to 1.0.

Figure 8: Monotonically increasing luminance colormap.

The properties of these colormaps can be demonstrated with the help of the Craik-O'Brian-Cornsweet illusion, shown in Figure 9. The gray region in the center of the circle should appear brighter than the gray at the outer edge of the circle, because of how local edge brightness contrast tends to propagate over neighboring regions [16]. The effect is somewhat weaker with the monotonically increasing colormap, but is eliminated with the isoluminant colormap. Although the strength of these effects vary with the method of printing or display, and with the observer, this is an example of how isoluminant colormaps can be preferable for interpreting image values.

7 DISCUSSION AND FUTURE WORK

We have shown that a simple perceptual test, observing the double face image, allows a user to quickly create a luminance match between two colors. As compared to luminance matching using the minimally distinct border technique, the double face method is equivalent in measured result, but more precise, and no slower. Given that the monitors we generally use for creating and displaying visualizations are not calibrated, this test provides a convenient means of creating colormaps with any pre-determined pattern of luminance variation (such as constant, or increasing). We believe the success of our method is due to the brain's special ability to detect and interpret images of human faces. Because of the simplicity of the method, we feel these results should be simple to reproduce. Also, since each color match takes about 20 seconds, a color map
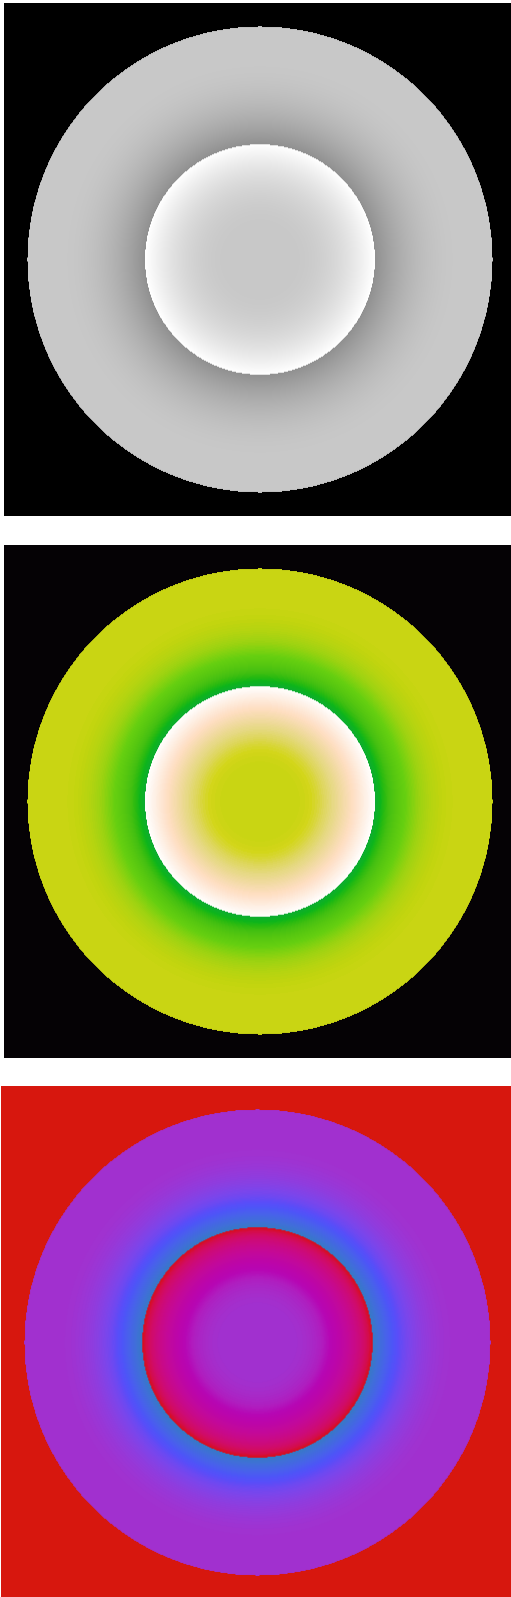
with 8 control points can be created in under ten minutes (using two matches per point).

The participants were asked a few questions after completing the experiment, and their answers may help explain the difference in precision between the two tasks. When asked which task was "easier", all but one participant identified the double face task. Answers about why it was easier generally related to face perception: "It was easy to tell when it didn't look like a face", "I could watch the eyes", "I was looking at her teeth", "I could just see the transition–one face pops out". Also, two participants mentioned that they appreciated how the side on which the positive face appears determines the direction in which to move the mouse to better match the luminance. This element of directional guidance, made possible by having two copies of the face image, is not present in the MDB task, since boundary distinctness increases with both increasing and decreasing color luminance relative to the fixed gray.

One interesting result of the user study is a demonstration of the variations in individual's color vision, even though all of the participants believed they were normal trichromats. For instance, generating "isoluminant" colormaps from the levels determined by participant #9 and participant #10 (based on the double face data), results in the two colormaps seen in Figure 10
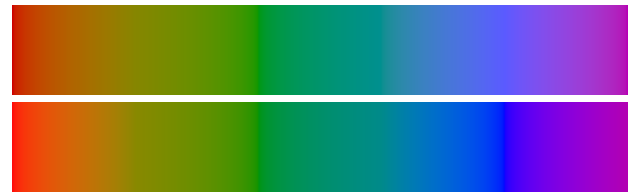

Figure 10: "Isoluminant" colormaps generated from the double face luminance matching data from participants #9 (top) and #10 (bottom).
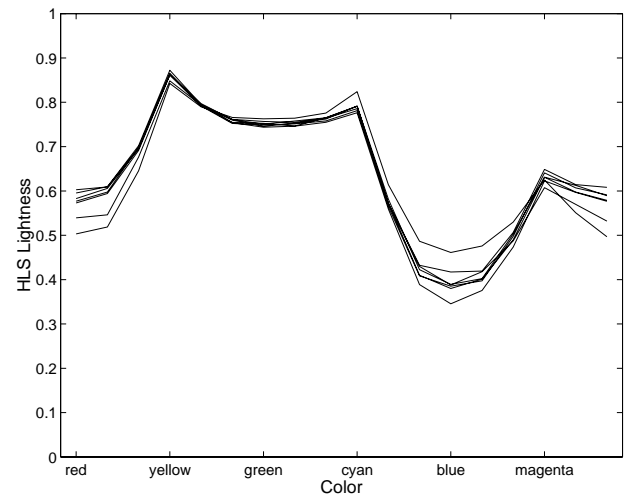

Figure 11: Plots of lightness of 18 colors along rainbow colormap, as measured by seven different participants.

Another way to see the variation in individual's color vision is to perform the converse of the experiment in the user study, varying the gray level in order to match the luminance of a fixed and fully saturated color. Using the double face image, seven participants did this luminance matching on 18 evenly spaced colors on the standard RGB rainbow colormap. Besides illustrating just how non-isoluminant the rainbow colormap is, the results plotted in Fig-

Figure 9: Craik-O'Brian-Cornsweet pattern, displayed (from top to bottom) in grayscale, with the monotonically increasing luminance colormap of Figure 8, and with the isoluminant colormap of Figure 7(a). The standard illusion (at top, in gray), may be stronger when viewed on a monitor.

ure 11 show that the most variation tends to occur in reds and blues. Considering that perceptual issues do play a significant role in colormap creation, we believe this raises the interesting possibility of user-specific colormaps.

As seen in Figure 5(a), the color red produced by far the largest difference between the MDB and double face methods, though not a statistically significant one. It may be that the difference was highest in red because red was the most saturated at the HLS lightness matching gray 0.5. Therefore, we would like to test our main hypotheses with a second experiment in which the lightness of gray is varied, instead of the color, so that all colors stay maximally saturated. A basic problem in perceptual colormaps not addressed in this paper, but which we are planning to pursue, is making the colormaps *perceptually uniform*, that is, making equal changes in the position along the colormap domain correspond to equal perceptual differences in the associated colors. While the CIELAB or CIELUV colorimetric spaces were intended to address this problem, they are only approximately uniform, in that the ability to discriminate between neighboring colors can vary significantly, depending on the position and direction within the color space [26]. In any case, these color spaces are not available on an uncalibrated monitor, and we have shown that it is possible to make perceptual measurements in an entirely device-dependent color space such as RGB.

8 ACKNOWLEDGMENTS

REFERENCES

[1] CIE Proceedings 1924. Cambridge University Press, Cambridge, 1926.

[2] CIE Proceedings 1951. Vol. 1, Sec. 4; Vol. 3, Page 37, Bureau Central de la CIE, Paris, 1951.

[3] Lawrence D. Bergman, Bernice E. Rogowitz, and Lloyd A. Treinish. A Rule-based Tool for Assisting Colormap Selection. In *Proceedings IEEE Visualization 1995*, pages 118–125. IEEE, October 1995.

[4] I Biederman and P Kalocsai. Neural and Psychological Analysis of Object and Face Recognition. In *Face Recognition: From Theory to Applications*, pages 3–25. Springer-Verlag, 1998.

[5] Robert M Boynton. *Human color vision*. Holt, Rinehart and Winston, New York, 1979.

[6] Patrick Cavanagh and Yvan G Leclerc. Shape from Shadows. *Journal of Experimental Psychology*, 15(1):3–27, 1989.

[7] William B Cowan. An Inexpensive Scheme for Calibration of a Colour Monitor in Terms of CIE Standard Coordinates. *ACM Computer Graphics*, 17(3):315–321, jul 1983.

[8] H D Ellis. Introduction to Aspects of Face Processing: Ten Questions in Need of Answers. In H Ellis, M Jeeves, F Newcombe, and A Young, editors, *Aspects of Face Processing*, pages 3–13, Dordrecht, 1986. Nijhoff.

[9] Mark D Fairchild. *Color Appearance Models*. Addison-Wesley, Reading, MA, 1998.

[10] J. Foley, A. van Dam, S. Feiner, and J. Hughes. *Computer Graphics Principles and Practice*, pages 592–595, 722–731. Addison-Wesley, second edition, 1990.

[11] Isabel Gauthier, Michael J Tarr, Adam W Anderson, Pawel Skudlarski, and John C Gore. Activation of the Middle Fusiform 'Face Area' Increases with Expertise in Recognizing Novel Objects. *Nature Neuroscience*, 2(6):568–573, June 1999.

[12] N George, R J Dolan, G R Fink, G C Baylis, C Russell, and J Driver. Contrast Polarity and Face Recognition in the Human Fusiform Gyrus. *Nature Neuroscience*, 2(6), June 1999.

[13] Andrew S Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann, San Fransisco, CA, 1995.

[14] S Ishida, Y Yamashita, T Matsuishi, M Ohshima, T Ohshima, K Kato, and H Maeda. Photosensitive Seizures Provoked While Viewing "Pocket Monster", a Made-for-Television Animation Program in Japan. *Epilepsia*, 39:1340–1344, June 1998.

[15] Ricardo J Motta. Visual Characterization of Color CRTs. In Ricardo J Motta and Hapet A Berberian, editors, *Device-Independent Color Imaging and Imaging Systems Integration, Proceedings SPIE*, volume 1909, pages 212–221, 1993.

[16] Stephen E Palmer. *Vision Science: Photons to Phenomenology*. The MIT Press, Cambridge, Massachusetts, 1999.

[17] V S Ramachandran. Perception of Shape from Shading. *Nature*, 331(6152):163–166, 1988.

[18] Penny Rheingans. Color, Change, and Control for Quantitative Data Display. In *Proceedings Visualization 1992*, pages 252–258. IEEE, October 1992.

[19] Penny Rheingans. Task-Based Color Scale Design. In *Proceedings Applied Image and Pattern Recognition*. SPIE, October 1999.

[20] P K Robertson and J F O'Callaghan. The Generation of Color Sequences for Univariate and Bivariate Mapping. *IEEE Computer Graphics and Applications*, 6(2):24–32, February 1986.

[21] B E Rogowitz and L A Treinish. How Not to Lie with Visualization. *Computers in Physics*, 10(3):268–273, 1996.

[22] Bernice E. Rogowitz and Alan D. Kalvin. The 'Which Blair Project:' A Quick Visual Method for Evaluating Perceptual Color Maps. In *Proceedings of IEEE Visualization*, 2001.

[23] S Subramanian and I Biederman. Does Contrast Reversal Affect Object Identification? *Investigative Opthalmology and Visual Science*, 38(998), 1997.

[24] T Troscianko and I Low. A Technique for Presenting Isoluminant Stimuli Using a Microcomputer. *Spatial Vision*, 1:197–202, 1985.

[25] Colin Ware. Color Sequences for Univariate maps: Theory, Experiments, and Principles. *IEEE Computer Graphics and Applications*, 8(5):41–49, September 1988.

[26] G Wyszecki and W S Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae*. John Wiley and Sons, New York, 2nd edition, 1982.