https://doi.org/10.2352/ISSN.2169-2629.2017.32.163

Exploiting Wide-gamut Displays

Greg Ward, Hyunjin Yoo, Afsoon Soudi & Tara Akhavan
Irystec, Inc.

Abstract

We present a hybrid color mapping (HCM) designed to preserve a selected region in chromaticity space while exploiting the larger gamut of the intended target display. Our method is based on the hypothesis that outside a certain set of critical colors, people are less particular and prefer a more saturated appearance if available. This hypothesis was borne out in the subject study we conducted. While other practitioners have implemented similar gamut-mapping techniques, our definition of preserved colors is more flexible and not based purely on saturation level. We employ an exponent or "acceleration factor" in our mapping to better preserve neighboring colors for a more natural appearance. Our method further avoids contrast and luminance changes, works between arbitrary gamuts, and is a fully invertible one-to-one mapping between color volumes.

Introduction

The introduction of the sRGB color space in 1996 greatly simplified the color reproduction landscape, especially over the Internet where profiles were often lost or ignored by popular browsers. For better or worse, sRGB has since become the de facto standard for color encoding, and most still images and video content arrive with its somewhat limited gamut. While we hope and expect this situation to change with new content, we must face that fact that legacy content will in all likelihood remain sRGB with its Rec. 709 primaries.

Meanwhile, consumer television and mobile displays continue to improve and offer wider color gamut. In particular, OLED pixels and RGB LED backlights make it possible and even economical to provide very narrow color primaries, yielding a much wider (if still triangular) gamut. Cinema projectors are also trending towards laser light sources, with their nearly monospectral primaries and wider dynamic range capabilities.

Given a limited input gamut (sRGB) and a wide gamut (OLED or laser) display, what is a color scientist to do? The "right" thing to do from a purist's perspective is simply reproduce the sRGB gamut as faithfully as possible, which should not pose a problem if it fits entirely within the target device gamut. However, this is most probably *not* the right thing to do from the perspective of the average consumer, who paid a little extra for a more vivid display and wants to see it that way. One very simple approach is to feed the sRGB images to the display and just see what it does. This is called the *same drive signal* (SDS) approach, which proved disastrous in the first mass marketed, wide gamut televisions. Since that time, manufacturers have backed off to the point where "wide gamut" often means "just measurably larger than sRGB" in order to avoid a garish appearance.

What we propose in this paper, and indeed others have implemented in one form or another, is a compromise between SDS and a colorimetric, or *true color mapping* (TCM), which is called a h*ybrid color mapping* or HCM. Our method for exploiting wide gamut displays lies somewhere between the saturation-based approach of most HCM implementations and the more sophisticated perceptual optimizations introduced in recent years.

Background

Color gamut mapping is a mature field, with entire books dedicated to the subject. For an excellent overview, we recommend Morovic's work [1]. Much effort has been devoted historically to the problem of reducing real-world colors to the limited gamut of a particular display or hardcopy device. We concern ourselves with the inverse problem of filling the wider color gamut of an LED or laser primary display from the smaller gamut that is more commonly provided, typically sRGB. This has become a more relevant goal with the advent of wide-gamut consumer OLED displays and laser projectors.

Laird et al. [2] evaluated an HCM method similar to the one we propose, comparing it to a more sophisticated method that adaptively adjusts chroma and lightness. Although the latter was preferred in their experiments, the HCM mixing function they employed was a basic linear blend keyed on saturation, which provided little control over the set of colors to preserve.

Our region-preserving approach is similar in spirit to that of Song et al. [3], who protect skin colors in their gamut extension algorithm. The main difference is that we protect earth tones as well, and work in (u′,v′) chromaticity space rather than HSV. This dimensionality reduction simplifies our method considerably and avoids contrast changes or the need to backfill color regions to maintain continuity.

The manipulation of contrast has been noted as problematic by Zamir et al. [4], who subsequently chose to work in CIE *ab* coordinates to circumvent issues with their original gamut extension algorithm [5]. A further modification to their optimization method was proposed by Vazquez-Corral and Bertalmío [6] to enable general gamut mapping (expansion and compression together), which is given automatically by an HCM method.

We start with our motivation for developing an HCM technique based on a preserved chromaticity region. We then describe our experimental design and the results of our subject study. We end with a brief discussion of future directions.

Gamut Mapping Method

We understand that there is no *correct* way to expand a smaller sRGB gamut to occupy a larger OLED or laser primary gamut. We assume the information about the larger gamut has been lost in the capture or creation of the sRGB input we are given, thus we cannot deduce the correct representation to fully utilize our display's color range. Rather than maintaining the smaller sRGB gamut on the more capable display, we wish to expand into a larger gamut in a perceptually preferred manner.

We seek a mapping that is straightforward, while achieving the following goals:

1) Unsaturated colors in the critical region of color space, i.e., earth- and flesh-tones, must be untouched (i.e, colorimetric).
2) The most saturated colors possible in sRGB should map to the most saturated colors in the destination gamut, achieving an *injective function* (one-to-one mapping) between gamut volumes.
3) Luminance and the associated contrast should be preserved.

We start by defining a region in color space where our mapping will be strictly colorimetric, and assume this is wholly contained within both source and destination gamuts. We call this the *sacred region*, which we define as a point in CIE (u′,v′) color space and a radial function surrounding it. For our experiments, we selected a central position of (u′,v′)=(0.217,0.483) with a constant radius of 0.051 based on empirical measurements of natural tones. (This center might be further tuned or adjusted, and a more sophisticated radial function employed in future.)

Our injective gamut mapping function is defined as follows. For colors falling inside the defined sacred region, we map values colorimetrically (TCM), reproducing them as closely as possible to the original sRGB values on our target display. We call this linear 3x3 mapping matrix $\mathbf{M_d}$. Thus:

$$RGB_d^T = \mathbf{M_d}\, RGB_i^T$$

where:

RGB_i = linearized input values in CCIR-709 primaries

RGB_d = linear colorimetric display drive values

The white point may be transformed as well by the above matrix to match the source white point to that of the display. We also map our linearized input colors to CIE XYZ using the matrix $\mathbf{M_x}$ then to (u′,v′) using the following standard formulae:

$$XYZ_i^T = \mathbf{M_x}\, RGB_i^T$$

$$u' = 4X/(X + 15Y + 3Z)$$

$$v' = 9Y/(X + 15Y + 3Z)$$

$$\text{where: } \mathbf{M_x} = \begin{bmatrix} 0.497 & 0.339 & 0.164 \\ 0.256 & 0.678 & 0.066 \\ 0.023 & 0.113 & 0.864 \end{bmatrix}$$

(The $\mathbf{M_x}$ matrix deliberately leaves off the D65 white point conversion, since the viewer is adapted to display white and we do not want to move the center of the sacred region.)

For input colors outside the sacred region, we interpolate between the colorimetric mapping above and an SDS mapping that sends the original RGB_i values to the display, applying linearity ("gamma") correction to each channel as needed.

We refer the reader to Figure 1. The sacred region is shown in green, and the red line drawn from the center to the sRGB gamut boundary represents our approximation to constant hue. The distance *a* is how far the input color is from the edge of the sacred region in (u′,v′) coordinates. The distance *b* is the distance from the edge of the sacred region to the sRGB gamut boundary along that hue line. The value *d* is the ratio of *a/b*. Our linear drive value is then computed as:

$$RGB_o = (1 - d^2) \cdot RGB_d + d^2 \cdot RGB_i$$

We found a power of *d* was preferred over the more commonly used linear interpolant in our pilot study, although the results were not overly sensitive to the acceleration factor. This differs from previous blending factors for HCM, which apply a linear ramp keyed on saturation rather than distance between a sacred region and the gamut boundary. The power function provides functional continuity and better preserves "almost sacred" colors.

The effect of this mapping on a regular array of (u′,v′) chromaticity coordinates is shown in Figure 2, where we map from sRGB to a particular set of AMOLED primaries. Note that there is little to no motion in the central portion defined as our sacred region. Even in the more extreme case of the laser primaries shown in Figure 3, neutral colors are mapped colorimetrically. However, more saturated colors are expanded out towards the enlarged gamut boundary, even rotating hue as necessary to reach the primary corners. Our hypothesis is that observers are less sensitive to color shifts at the extremes, so long as general relationships between color values are maintained. Interpolating between colorimetric and direct drive signal mappings maximizes use of the destination gamut without distorting local relationships. The third dimension (luminance) is not visualized, as it does not affect our mapping. Values that were clipped to the gamut boundary in sRGB will be clipped in the same way in the destination gamut; this is an intended consequence of the HCM method.

Figure 4 shows (to the extent possible) the color shifts we see when expanding from an sRGB to laser primary color space using our method. Unsaturated colors match between the original and our display, while saturated colors become more saturated and may shift in hue towards the target device primaries.

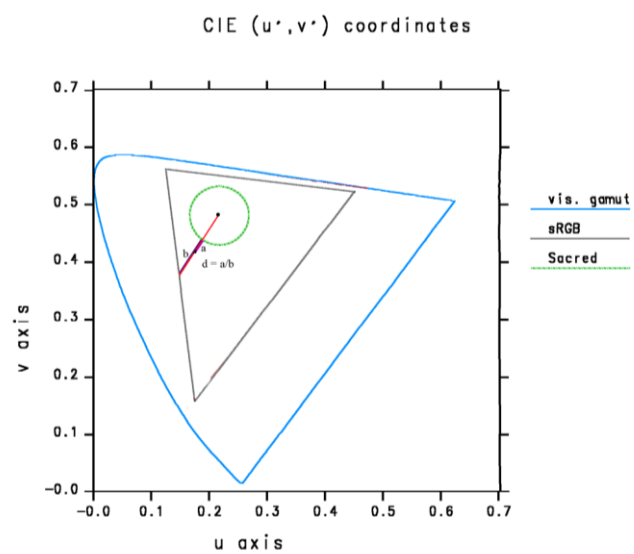

Figure 1. Sacred region (green) with line drawn from center through input color to sRGB gamut boundary.

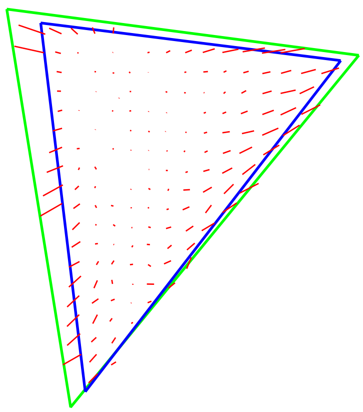

Figure 2. Mapping from an sRGB gamut to AMOLED primaries. The red lines represent color motions for our HCM method.

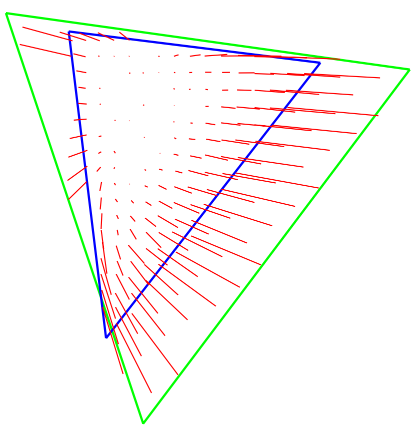

Figure 3. Mapping from an sRGB gamut to laser primaries used in subject study with example color motions.

Experimental Validation

We subjectively evaluated the performance of our gamut-mapping model using the pairwise comparison approach introduced in [7]. The experiment was set up in a dark room with a laser projector (PicoP by MicroVision Inc.) having a wide gamut color space shown in Figure 3. We used 10 images processed by 3 different color models, our proposed HCM gamut mapping, colorimetric or true color mapping - TCM, and original image – SDS (same drive signal). We asked 20 naïve observers to compare the presented result. Our observers were asked to pick their preferred image of the pair. For each observer, total 30 pairs of images were displayed using the laser projector, 10 pairs for TCM:HCM, 10 pairs for HCM:SDS, and 10 pairs for SDS:TCM. The observers were instructed to select one of the two displayed images as their preferred image based on the overall feeling of the color and skin tones.

Figure 4. Proposed color mapping applied to sRGB input (top) and shown using laser display on bottom. Intense colors become more intense, and some shift slightly in hue, especially in deep blue where primaries do not align.

Our observers consist of 7 females and 13 males from the age of 20 to 58. On average, the whole experiment took about 10 minutes for each observer.

Figure 5 shows gamut mapping results (HCM) with original images (SDS) and colorimetric mapping (TCM). A few of the images include well-known actors whose skin tones may be familiar to the observers. Our gamut mapping result keeps the face and skin color as in the colorimetric reference, but represents other areas more vividly, such as the colorful clothes in the image Wedding (1st row, left), the tiger balloon in the image Girl (4th row, right), and the red pant of a standing boy in the image Family (5th row, right).

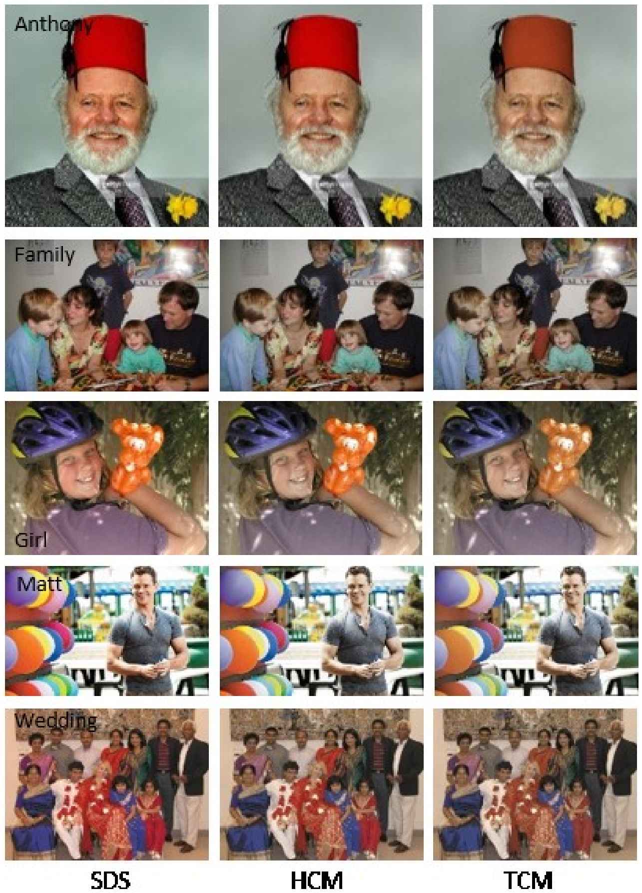

Figure 5. Gamut mapping examples with original images and colorimetric references. HCM - our proposed gamut mapping, SDS - original image, and TCM – colorimetric or true color mapping.

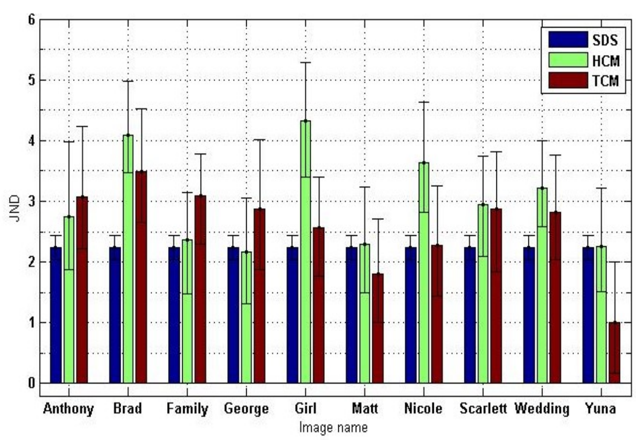

Figure 6. Subjective evaluation results of pairwise comparison representing as JND values for each 10 images including error bars which denote 95% confidence intervals calculated by bootstrapping. HCM: our proposed gamut mapping, SDS: original image, and TCM: colorimetric or true color mapping.

We used the pairwise comparison method with just-noticeable-difference (JND) evaluation in our experiment. This approach has been used recently for subjective evaluation in the literature [7-9]. We used the Bayesian method of Silverstein and Farrell [10], which maximizes the probability that the pairwise comparison result accounts for the experiment under the Thurstone Case V assumptions. During an optimization procedure, a quality value for each image is calculated to maximize the probability, modeled by the binomial distribution. Since we have 3 conditions for comparison (HCM, TCM, SDS), this Bayesian approach is suitable, as it is robust to unanimous answers and common when a large number of conditions are compared.

Figure 6 shows the result of the subjective evaluation calculating the JND values as defined in [7]. The absolute JND values are not meaningful by themselves, since only relative difference can be used for discriminating choices. A method with higher JND is preferred over methods with smaller JND values, where 1 JND corresponds to 75% discrimination threshold. The Figure 6 represents each JND value for each scene, rather than the average value, because JND is a relative value that can be also meaningful when compared with others. In the Figure 6, we also represent the confidence intervals with 95% probability for each JND. To calculate the confidence intervals we used a numerical method, known as bootstrapping which allows estimation of the sampling distribution of almost any statistic using random sampling method [18]. We generated 500 ramdom sampling, then computed 2.5^{th} and 97.5^{th} percentiles for each JND point. The reason why JND values of SDS are same is that both JND and confidence intervals for JND are relative values. So we need a reference point to calculate them. In our method we just choose SDS as the reference point. For 7 of the images, our proposed mapping is the most preferred method, with JND differences of 0.03 ~ 1.8 between it and the second most preferred method. For 3 of the images (Anthony, Family, George), HCM is not the most preferred method, losing by JND differences of 0.31 ~ 0.71.

Conclusion

We presented a simple yet flexible hybrid color mapping method for extending sRGB content on a wide gamut display, and showed this method was preferred over colorimetric (TCM) and same drive signal (SDS) methods on average in images containing natural tones with saturated colors. The key to our method is assigning a sacred region of chromaticity space where colors are mapped precisely, and a function for moving towards SDS at the gamut boundaries.

In future work, we hope to refine our description of this sacred chromaticity region with a customized radial function that better defines the boundary of viewer-critical colors. We have noted that there are still certain flesh tones, especially in ruddy-faced individuals, where further saturation is not preferred. Under some lighting, these tones may cross over our sacred circle, where a radial function would avoid such issues.

Finally, we would like to compare our method to more sophisticated, perceptually-based optimization methods.

References

[1] Ján Morovič, *Color Gamut Mapping*, Wiley Publishing, 2008.

[2] J. Laird, R. Muijs, J. Kuang, “Development and Evaluation of Gamut Extension Algorithms,” *Color Research & Application*, 34:6 (2008).

[3] G. Song, X. Meng, H. Li, Y. Han, “Skin Color Region Protect Algorithm for Color Gamut Extension,” *Journal of Information & Computational Science*, 11:6 (2014), pp. 1909–1916.

[4] S.W. Zamir, J. Vazquez-Corral, M. Bertalmío, “Gamut Extension for Cinema: Psychophysical Evaluation of the State of the Art, and a New Algorithm,” *Proc. SPIE Human Vison and Electronic Imaging XX*, 2015.

[5] S.W. Zamir, J. Vazquez-Corral, M. Bertalmío, “Gamut Mapping in Cinematography Through Perceptually-Based Contrast Modification,” *IEEE Journal of Selected Topics in Signal Processing*, 8:3, June 2014.

[6] J. Vazquez-Corral, M. Bertalmío, “Perceptually inspired gamut mapping between any gamuts with any intersection,” *AIC Midterm Meeting*, 2015.

[7] G. Eilertsen, R. Wanat, R. K. Mantiuk, and J. Unger, “Evaluation of tone mapping operators for hdr-video,” Computer Graphics Forum, vol. 32, no. 7, pp. 275-284, Wiley Online Library, 2013.

[8] R. Wanat, and R. K. Mantiuk, “Simulating and compensating changes in appearance between day and night vision,” ACM Trans. Graph., vol. 33, no. 4, 147-1, 2014.

[9] M. Rezagholizadeh, T. Akhavan, A. Soudi, H. Kaufmann, and J. J. Clark, “A Retargeting Approach for Mesopic Vision: Simulation and Compensation,” Journal of Imaging Science and Technology, 2015.

[10] D. A. Silverstein, and J. E. Farrell, “Efficient method for paired comparison,” Journal of Electronic Imaging, vol. 10, no. 2, pp. 394-398, 2001.

[11] Getty Images, Image-Anthony, from http://www.gettyimages.ca/detail/news-photo/anthony-hopkins-wearing-a-red-fez-unveils-a-statue-of-tommy-news-photo/79953291.

[12] Associated Newspapers Ltd, August 28, 2008, Image-Brad, from http://www.dailymail.co.uk/tvshowbiz/article-1050071/Brad-Pitts-shining-moment-Venice-interrupted-hes-asked-present-gift-George-Clooney.html#ixzz453n84rM6.

[13] NAVER Corp., Image-George, from http://movie.naver.com/movie/bi/mi/photoView.nhn?code=81834&imageNid=6261978

[14] FANPOP, Inc., Image-Yuna, from http://www.fanpop.com/clubs/yuna-kim/images/9910994/title/scheherazade-yuna-kim-08-09-season-free-skating-long-program-photo.

[15] Emma, April 21, 2012, Image-Scarlett, from http://vivanorada.blogspot.ca/2012/04/you-got-style-scarlett-johansson.html

[16] Just Jared Inc., July 13, 2011, Image-Nicole, from http://www.justjared.com/photo-gallery/2560483/nicole-kidman-keith-urban-snow-flower-06/#ixzz453u8lJIE

[17] Bett Watts, December 13, 2011, Image-Matt, from http://www.gq.com/story/matt-damon-gq-january-2012-cover-story-article.

[18] H. Varian, “Bootstrap tutorial.,” Mathematica Journal, vol. 9, no. 4, pp. 768-775, 2005.