

When Good Hues Go Bad

*Lawrence A. Taplin and Garrett M. Johnson
Munsell Color Science Laboratory
Center for Imaging Science, Rochester Institute of Technology
Rochester, NY*

Abstract

A psychophysical experiment was performed to determine the direction of least and most preferred color shifts for digital images. 360 images were subdivided into 18 hue regions. The pixels in each hue region were shifted eight directions in IPT color space, and observers were asked to classify which direction was least preferred, and which was most preferred. From these data preference maps were generated indicating the shifts in color space that should be avoided, if possible, when designing color reproduction systems.

Introduction

Traditional color reproduction systems rely heavily on metamerism in order to reproduce the wide variations of natural spectra with a limited number of color channels. Metamerism relies on the integration of human visual system responses from the trichromatic cone photoreceptors, essentially allowing an infinite number of spectral curves to integrate to identical cone responses. Thus metamerism allows us to capture a natural scene with a three-primary RGB camera, and reproduce that scene with a three or more primary printer.

Metamerism can also cause problems with color reproduction at both the input and output stages. If RGB camera sensitivities do not match the cone responses, or a linear transform of them such as CIE XYZ color-matching functions, then the captured colors might vary slightly from those perceived by a human observer. This phenomenon can be thought of as a form of observer metamerism. Similarly, it might be possible to create a print that matches an original scene in appearance under a specific illumination and viewing condition, but that no longer matches in appearance when that condition is changed. This phenomenon can be considered a form of illuminant metamerism.

The change in appearance results in color inconstancy for a color reproduction system. When dealing with changes in illumination, for both the input and output stages, this inconstancy often causes a chromatic shift for specific areas of color space. The extent and location of this shift is controlled both by the primaries of the imaging system and the color mixture model. When designing imaging systems it may be desirable to try and minimize or control the direction of this color shift, either through optimization of the sensors or primaries, or through the use of additional primaries, such as green and orange with a CMYK

printer.¹ If we cannot eliminate this color inconstancy, the question remains as to which direction in color space it should be directed when possible.

There has been considerable research into color preference for reproduction purposes. Bartleson²⁻⁴ and Hunt⁵ studied these problems in the 1950s, 60s, and 70s, specifically for memory colors such as sky-blue, grass-green, and Caucasian skin-tone for photographic systems. These topics have been addressed again more recently using digital imaging techniques.⁶⁻¹⁰ Memory colors can have an effect on overall color preference, such as a shift toward yellow for skin-tones and grass, or a shift towards cyan for sky. There is little research to indicate whether this preference holds true for objects that share a similar color-center, but are not associated with memory colors. For instance, does a “sky-blue” house exhibit the same behavior as an actual sky, or does the preference change? There is also very little information available to guide us when optimizing system primaries. These primaries cannot be scene-content dependent, and so they must behave similarly for all colors regardless of the object. Several questions remain that are addressed in this paper. If color inconstancy is inevitable, which color shifts should I avoid? Should different regions of color space be shifted in different directions?

Experimental

A psychophysical experiment was designed to address the preference, or more appropriately the dissatisfaction of color shifts for digital images. A large commercial stock-photo library of 19,000 PhotoCD images from Corel was utilized for this experiment. Each of the images was converted into the IPT color space¹¹ for examination and processing. The IPT space was chosen as it represents a color space designed specifically for linearity of the hue dimension and spacing of the Munsell re-notation data. The IPT space is an opponent color representation with achromatic, red-green, and yellow-blue channels. These channels were converted into hue angle for each image, and then segmented into 18 regions. The histogram of each hue region was calculated for all 19,000 images. From these histograms the images were selected at random such that there were 20 images with at least 5% of the total pixels inside each hue region. This yielded 360 images, representing the 18 hue regions. The 360 images contained many variations including portraits, art reproductions, natural scenes, wildlife, man-made objects and abstract photographs.

The images were processed by first creating a mask representing all pixels inside a specific hue region. The pixels inside the mask were then shifted in IPT space along the 4 axes, as well as the 4 diagonals. The magnitude of the shift was 0.08 in IPT units, which roughly correlates to 12 CIELAB units (a^* and b^*). The magnitude of the shift was modulated using a Gaussian distribution across hue angle to reduce visible artifacts at the edge of the hue region. This shift was designed to be clearly visible for most images. Figure 1 shows an original image and an example hue mask for 200 degrees. Images were selected so that at least 5% of the pixels in each image were shifted, although it was often more. This assured that different regions (such as foreground, background, and main subject) of the images were represented. Once the colors were shifted in IPT space, the images were converted back to sRGB for display.




Figure 1. Original image and hue region mask (right). White represents hues NOT shifted, black represents hues shifted, and gray represents hues shifted slightly.

In order to try and gain an understanding of preference and dissatisfaction for many color regions across many different image types, it was necessary to use a large number of images. This prohibited the use of many traditional psychophysical techniques such as paired comparison or rank ordering. Rather, a simple technique was utilized. Observers were presented with a single window that contained an original image, as well as eight images having a single hue region shifted in color space. The observers were instructed to select the reproduction that was the least preferred with the left mouse button, marking it with a red outline. Next, they were instructed to select the most preferred image with the right mouse button, marking it with a green outline. They were allowed to modify their selections before confirming their choice with the “next” button. This technique is similar to rank-ordering, but does not produce an interval scale of preference. Figure 2 shows a screen shot of the graphical user interface (GUI) used in this experiment.

The experiment was performed on a Lacie Electronblue 19” CRT display, set at 1600x1200 in sRGB mode. Observers were instructed to sit at a normal viewing distance, and viewed the display in a dim room.

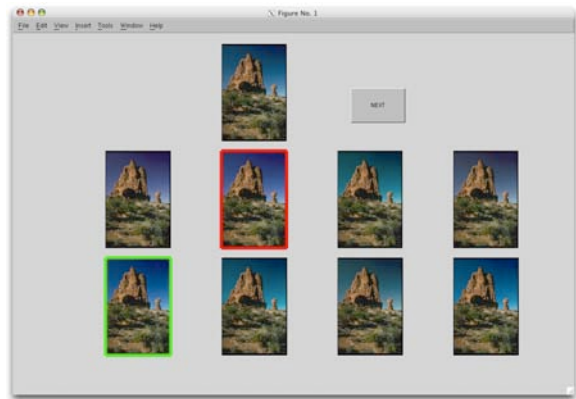


Figure 2. GUI used in psychophysical experiment. A red border indicated least preferred reproduction, a green border most preferred.

The examination of 360 images took approximately two hours for most observers. Five observers completed the experiment, for a total of 100 observations at each hue region.

Results and Discussion

The experimental method was not designed to create a scale of preference. Rather, we were interested in with determining the color shifts that should be avoided if at all possible. The collection of the most preferred color shift might also provide insight into this problem. The experimental data were first summarized by calculating the frequency, or percentage of times each shift direction was selected, for both the preference and dissatisfaction directions. The least preferred color shift should therefore have the highest percentage of selection, and likewise for the most preferred color. These data were collected for all 18 of the initial hue regions, resulting in average color shift preference maps for each color center.

An example of a single preference map for the hue region centered at 120 degrees in IPT space is shown in Figure 3. A preference map for all 20 images in a single hue region is shown, with an illustrated diagram of a single image. The red lines indicate the least preferred color shifts, while the green line indicates the most preferred. The eight color shift directions towards their respected hue angles are shown in the color circles, along with an example image shifted in those directions.

The extent of the red or green line away from the origin is proportional to the frequency of selection for that particular color shift. For the hue region centered at 120 degrees, as shown in Figure 3, it can be determined that the most dissatisfied color shift is towards 270 degrees, a purplish blue. It should be noted that because IPT is based on equal spacing of the Munsell data the unique hues of red, green, yellow, and blue do not necessarily lie on the cardinal axes of 0, 90, 180, and 270 degrees. Figure 3 also indicates that the neighboring color shifts towards 225 and 315 were often the least preferred as well. The relatively broad green map indicates that the most preferred shift was less apparent for this particular hue region. The most preferred color shift spanned the range from 0 to 180 degrees. The range of preference is probably due to the various image

dependencies between the 20 scenes in this hue region. It is interesting to note that for this particular hue region the least and most preferred color shift directions are mutually exclusive. This suggests that for this particular region of color space, it might be possible to determine a general direction or color inconstancy to avoid when designing reproduction systems. Figure 4 shows the radar preference plots for all 18 hue regions in IPT space. The preference plots should be read similarly to Figure 3, essentially relocating the plots to the origin to determine the directions of color shift.

preferred shifts appear to be in the greenish-yellow direction, and the most preferred are towards a more saturated red, or towards neutral in the cyan direction. For greenish colors between 120 and 200 degrees there are clear indications of a more shift toward a more bluish green, with the least preferred color shifting towards purple, resulting in an overall desaturation. When the images were shifted in the direction of the hue angle opposite the selection angle, a decrease in chroma occurs. However, in this experiment only zero and 180 degree selections angles have an exact opposite translation

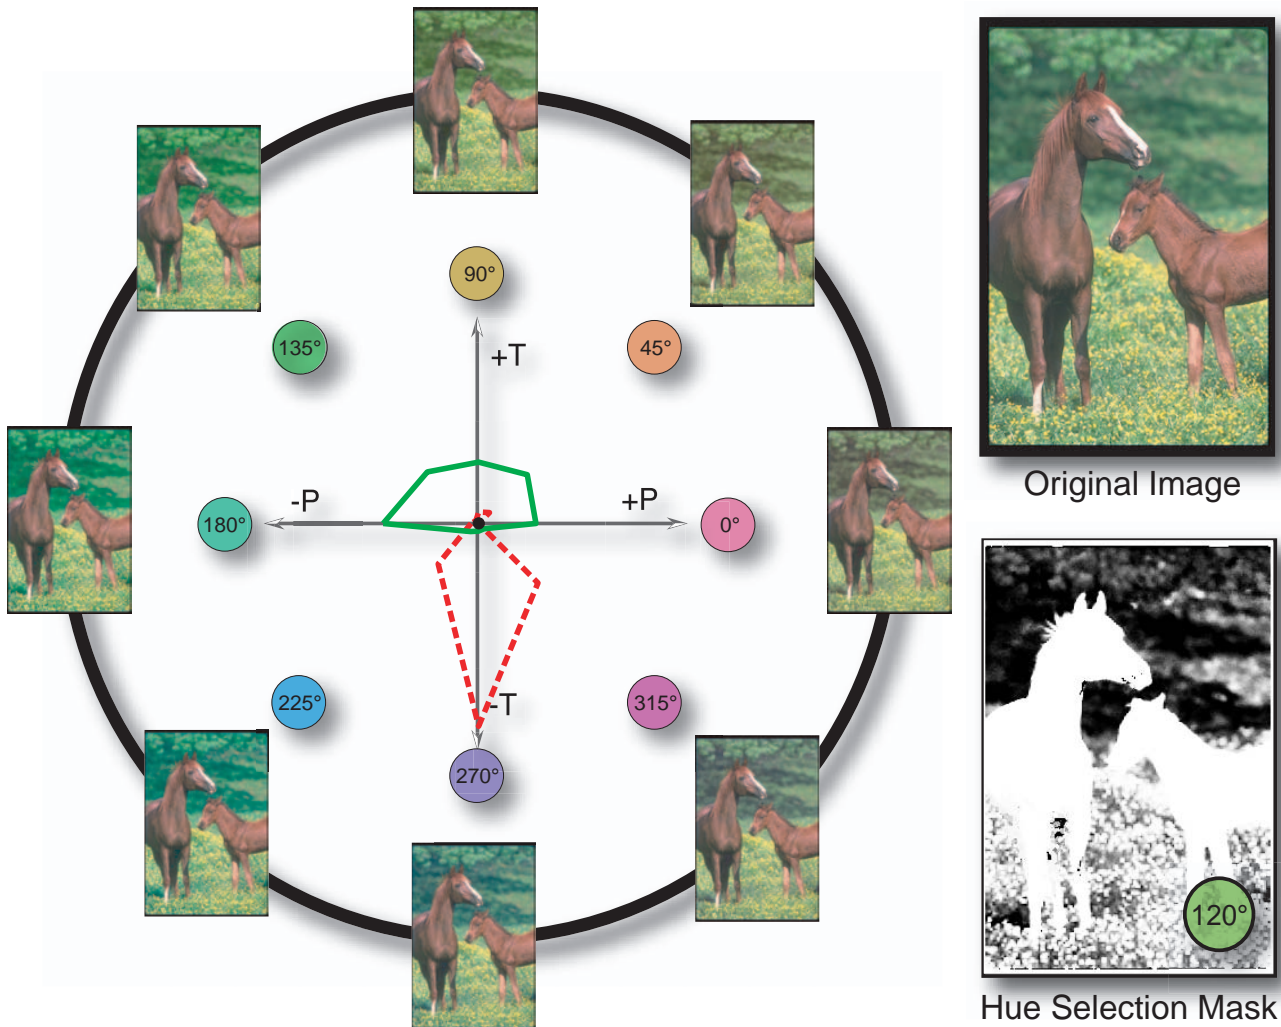


Figure 3. Preference Directions for Most (Green/Solid) and Least (Red/Dashed) Preferred Color Shifts for 18 Hue Regions.

The first striking feature from Figure 4 is the general orthogonality between the most and least preferred colors shifts. This indicates that there might be generalized axes in color space to avoid directional shifts for most colors, and general axes that are acceptable. These axes differ depending on the hue region. Other generalizations might be gleaned from analysis of Figure 4. For "red" hue regions between 320 and 40 degrees, the least

angle. It is therefore difficult to separate a preference for decreased chroma from one for a particular hue direction. There are weaker preferences in the 40-60 degree regions as well as the 220-240, as evidenced by the less sharp preference plots. The general directions still tend to be orthogonal in color space, but the preferences seem to be spread along these dimensions. These hue regions typically encompass skin tones and blue-sky. The lack of

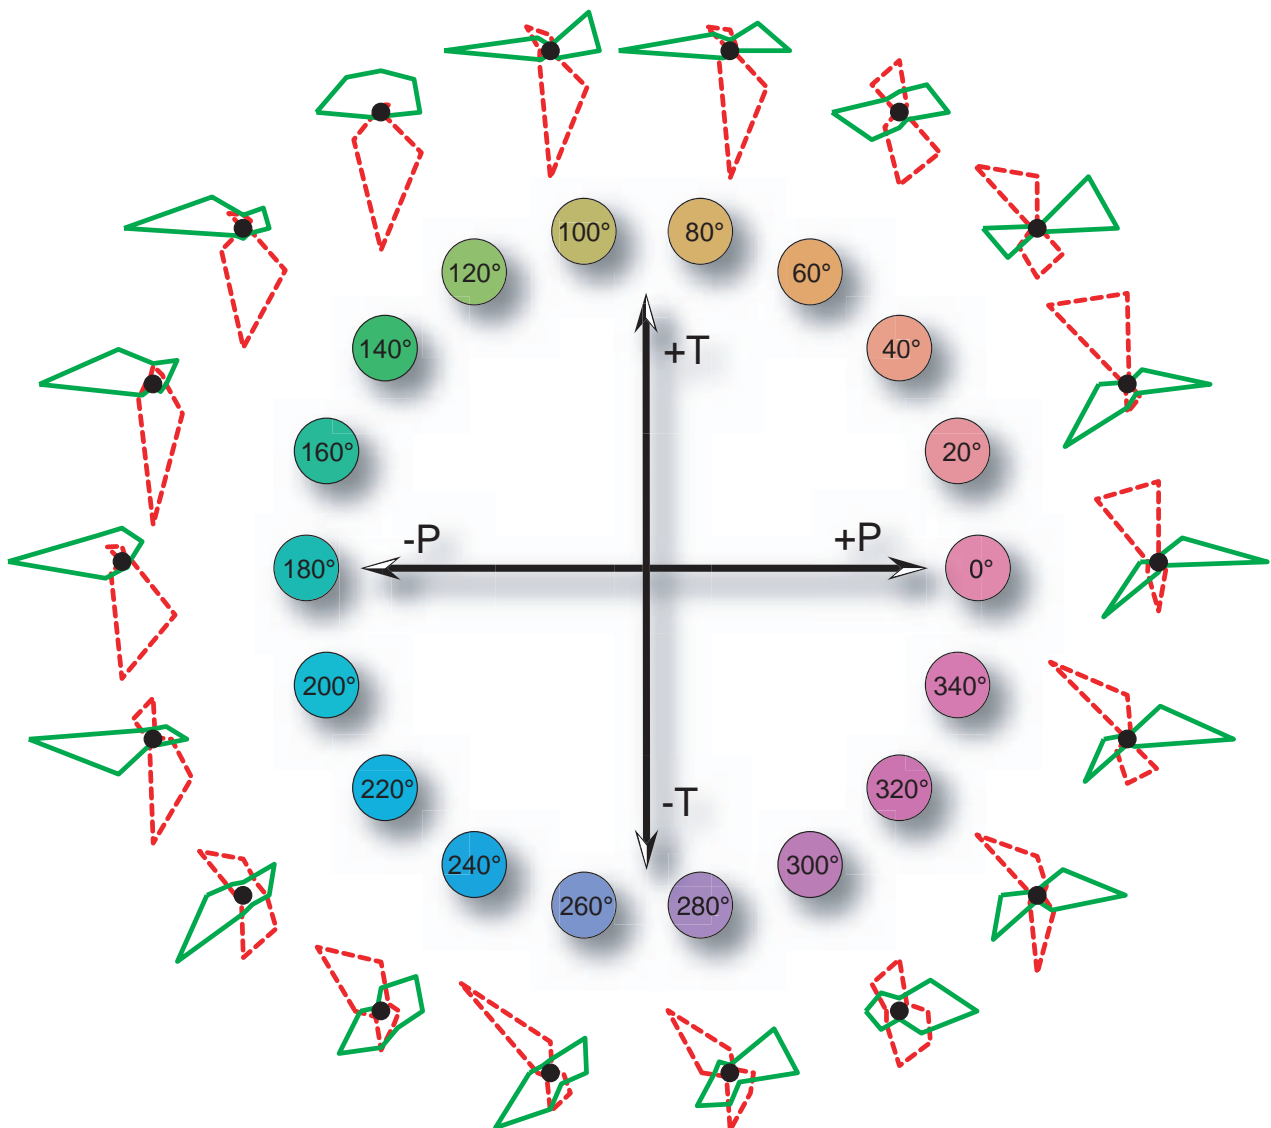


Figure 4. Preference Directions for Most (Green/Solid) and Least (Red/Dashed) Preferred Color Shifts for 18 Hue Regions.

clear preferences in either direction show that observers did not have an obvious color shift that was preferred or disliked. This might be because the photographer took care to create a pleasing rendition of both sky and skin tones.

Conclusions

This type of experiment provides a wealth of data for analysis. Future investigation will examine individual images at each hue region to determine any image dependencies. The images are also being analyzed with a color image difference metric¹² to determine if the least and most preferred images correspond to the least and most visible differences between each color shifted image and the original. The ultimate goal of this research is a complete preference mapping of color space. Knowing the direction of most and least preferred color shifts will

provide system engineers with a tool useful for ink design, sensor filter selection, and gamut mapping.

References

1. Yongda Chen, Roy S. Berns, Lawrence A. Taplin, and Francisco H. Imai, A multi-ink color-separation algorithm maximizing color constancy, Proc. IS&T 11th Color Imaging Conference, 277-281 (2003).
2. C. James Bartleson, Memory colors of familiar objects, Journal of Optical Society of America, 50 73-77 (1960).
3. C. James Bartleson, Color in memory in relation to photographic reproduction, Phtogr. Sci. Eng, 5 327-331 (1961).
4. C. James Bartleson and C.P. Bray, On the preferred reproduction of flesh, blue-sky, and green-grass colors, Phtogr. Sci. Eng, 6 19-25 (1962).

5. R.W.G. Hunt, I.T. Pitt, and L.M. Winter, The preferred reproduction of blue sky, green grass and Caucasian skin in colour photography, *Journal of Photographic Science*, 22 144-149 (1974).
6. S.N. Yendrikhovskij, F.J.J. Blommaert, and H. de Ridder, Color reproduction and the naturalness constraint, *Color Res. & Appl.*, 24(1) 52-67 (1999).
7. Karen Topfer and Robert Cookingham, The quantitative aspects of color rendering for memory colors, *Proc. IS&T PICS Conference*, 94-94 (2000).
8. Peter Bodrogi and Tunde Tarczali, Color memory for various sky, skin, and plant colours: effect of the image context, 26(4) 278-289 (2001).
9. Scot R. Fernandez and Mark D. Fairchild, Preferred color reproduction of images with unknown colorimetry, *Proc. IS&T 9th Color Imaging Conference*, 274-279 (2001).
10. Nathan Moroney and Ingeborg Tastl, Applications of a color-naming database, *Proc. IS&T PICS Conference*, 566-570 (2003).
11. Fritz Ebner and Mark D. Fairchild, Development and testing of a color space (IPT) with improved hue uniformity, *Proc. IS&T 6th Color Imaging Conference*, 8-13 (1998).
12. Garrett M. Johnson and Mark D. Fairchild, A top-down description of S-CIELAB and CIEDE2000, *Color Research and Application*, 28(6) (2003).

Biography

Lawrence A. Taplin is a scientist with the Munsell Color Science Laboratory at the Rochester Institute of Technology where he received a M. S. degree in Color Science. He also holds a B. S. in Computer Science from the University of Delaware. His research is focused on spectral imaging of museum artwork for digital archiving and reproduction.

Garrett M. Johnson is a scientist with the Munsell Color Science Laboratory. He received his Ph. D. in Imaging Science from RIT in 2003 for research on image differences, quality, and appearance. He holds B. S. and M. S. degrees in Imaging Science and Color Science, both from RIT.