Gamut Mapping
Squeezing the Most out of Your Color System

M. Wolski, J. P. Allebach and C. A. Bouman
Purdue University, West Lafayette, Indiana

Abstract

Gamut mapping has become an important research topic in the last few years. It is an important component of device-independent-color. In this paper we review some common gamut mapping techniques and present a new technique in which the type of mapping performed is based on the position in color space.

Introduction

The advent of low-cost personal computers capable of capturing, displaying, and printing color images has raised considerable interest in device independent color. One concern in device independent color is that different devices are capable of producing different ranges of colors. The range of colors associated with a device is known as its gamut. In a paper on device independent color, Fairchild has stated that gamut mapping is "perhaps the most important element in transforming images across media"[1].

In general, similar devices have similar gamuts. Dissimilar devices can have quite different gamuts. For example, emissive devices (*e.g.* computer monitors) attain their highest chroma at a relatively high lightnesses compared to reflective devices (*e.g.* color printers)[2]. The variability of color gamuts among various devices leads to an important question. "How can we optimally alter a color image so that it fits within a specified output gamut?". This paper presents a first step toward a general solution to that problem.

The gamut can be represented in any of a number of spaces. Some examples are CIE XYZ, the luminance-chromaticity space Yxy, uniform color spaces L*a*b* and L*u*v*, or color appearance spaces[3,4,5,6,7]. In most cases, the underlying space is 3-dimensional and can be described in terms of cylindrical coordinates. The vertical axis represents some form of lightness, the angle around the axis represents hue, and the distance outward from the axis represents either chroma or saturation.

In this paper, we investigate gamut mapping in L*u*v* space. For simplicity, we chose L*u*v* over the color appearance spaces[5,6,7]. We chose L*u*v* over L*a*b* because chromaticity can be preserved in L*u*v* by keeping saturation constant. This is not quite true in L*a*b* ([3] p.168).

Gamut mapping is a fairly new topic in the literature. Stone, Cowan, and Beaty[8] investigated both clipping and compression techniques in XYZ space. Gentile, Walowitt, and Allebach[9] compared several methods of gamut mismatch compensation in L*u*v* space. Generally, they preferred clipping in chroma while keeping lightness and hue constant. Pariser[10] performed a study similar to that of Gentile but with hard copy images. He found that, depending on the input image, either of 2 techniques was preferred. The first preserves lightness and hue while clipping chroma. The second preserves hue while clipping lightness and chroma toward the (50,0,0) point of L*a*b* space. MacDonald[11] investigated various mappings in the Hunt[6] color appearance space. His preferred method is also simultaneous compression of lightness and chroma toward a mid-gamut point. Hoshino and Berns[12] looked at lightness mappings in the Hunt color appearance space. They introduced the concept of "soft compression" in which a cut-off point is defined on the axis of interest. Compression takes place only for values above the cut-off.

Our target device gamuts represent a typical CRT monitor and a typical ink jet printer. However, all our output is observed on the monitor so the gamut of the printer is modified (clipped) to fit within that of the monitor. All input images are taken from a photo-CD [13,14,15] and have gamuts exceeding those of both the monitor and printer.

We used a calibrated display to view and compare images. For the target monitor gamut, we compared our technique to scaled versions of the original image and to clipping in RGB (the default photo-CD mapping technique). For the printer gamut, we compared our technique to strategies which compress saturation only. In all cases, our method improved image quality at least slightly (as judged by ourselves). Moreover, in many cases, our method markedly improved image quality. This was true for both the monitor and printer target gamuts.

Based on the preferred results of both Pariser[10] and MacDonald[11], we investigated mapping methods which preserve hue while changing lightness and saturation. We experimented with several techniques before settling on the one described in this paper. Our first experiments altered saturation only. Our findings agreed with the previously mentioned studies - the method works well with some images and poorly with others. A common problem with this method is a noticeable loss of saturation. This is especially true for dark reds mapped to a printer gamut. Another problem is that some bright chromatic pixels get mapped to white, giving the unwanted appearance of specular highlights.

A second mapping method which we tried was changing lightness only. This method also met with mixed results, depending on the image. The biggest problem with this method was a speckled appearance due to a large increase in lightness of out-of-gamut dark chromatic pixels. Another problem was loss of contrast due to compression in lightness.

A third method which we investigated was simultaneous clipping of lightness and saturation toward the center of the target gamut. The major problem with this method is loss of bright highlights and desaturation of bright chromatic colors.

Mapping Strategy

From our experiments, we concluded that different parts of the gamut volume could benefit most from different mapping techniques. There are two conflicting objectives for a lightness-saturation mapping: preservation of original image contrast and preservation of original colorfulness. To preserve contrast, we would like to retain original lightness levels, or at least original lightness differences. To preserve colorfulness, we would like to retain original saturation. We found that the relative importance of these two objectives changes, depending on position in LHS color space.

The mapping strategy is divided into two parts (Figure 1). Both parts maintain a constant hue while changing lightness and saturation. The first part is image dependent, but can be implemented with simple data structures describing the gamut boundary. It maps the input image to fit within the smallest "cylinder" containing the target gamut. This cylinder extends from the lowest to the highest attainable device lightness and has radii (saturations) which are constant with varying lightness but which vary with hue angle. At each hue angle the cylinder radius is equal to the maximum attainable device saturation at that hue angle.

Part I of the mapping strategy is illustrated in Figure 2. Saturations are mapped using soft compression with a 95% cut-off. Lightnesses are mapped using a combination of shifting and soft compression. Following a small lightness shift, light pixels are soft compressed down to the upper cylinder boundary, using a cut-off of 75% the distance from 50 to the upper boundary. Dark pixels are clipped to the lower cylinder boundary.

We tried to implement an automatic shift calculation based on the number of out-of-gamut pixels or on the total out-of-gamut volume for each shift position. However we had little success with this strategy. We found that we obtained the best results by basing the shift solely on the target gamut. For the monitor target gamut, we shift the image down slightly in lightness (5 units) while for the printer target gamut we shift it up slightly (2 units). We only considered shifting and scaling for the lightness mapping. This was not because these are necessarily better than other techniques, but because they are straightforward to implement and we wanted to concentrate on other aspects of the gamut mapping problem. It may be possible to optimize other lightness mapping techniques based on image quality metrics[16, 17].

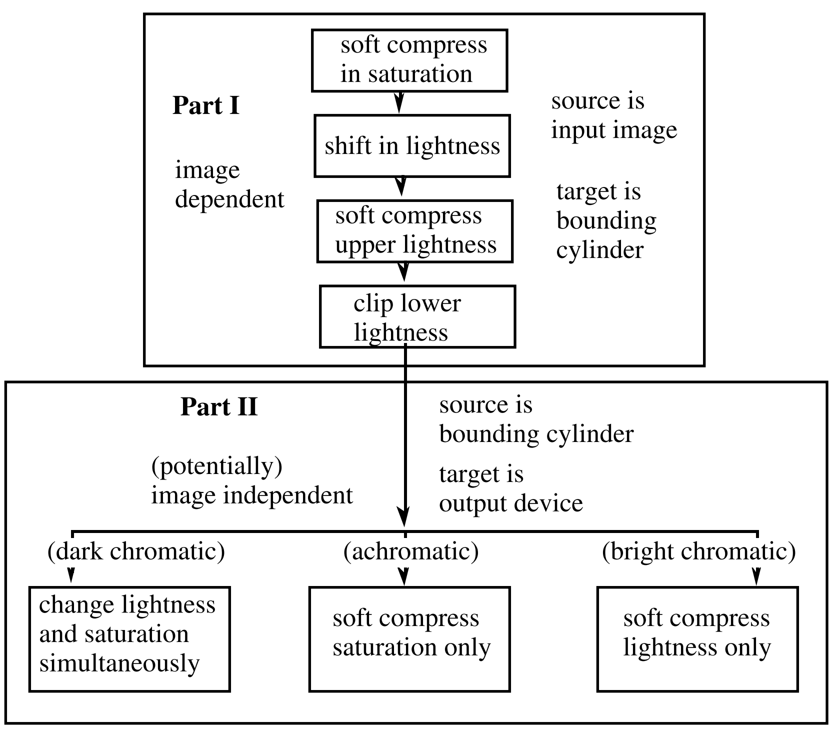

Figure 1. Overview of our gamut mapping technique.

We note that our lightness mapping does not explicitly include global lightness scaling. There are several reasons for this. First, soft compression takes care of selectively scaling lightness. Second, we found it important to attempt to maintain image contrast. Global down-scaling reduces contrast while global shifting does not. Third, scaling does not help reduce the number of out-of-gamut pixels which have lightnesses below the minimum target lightness. An upward shift can reduce this number.

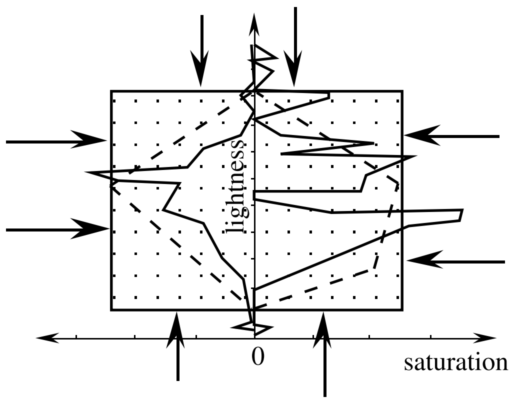

Figure 2. Illustration of Part I of the mapping technique. The dashed lines represent the target gamut. The shaded area is the bounding cylinder. The solid lines represent the input gamut. Pixels with saturations exceeding the cylinder boundary are soft compressed using a cut-off of 95%. The entire image is then shifted up 2 units in lightness (for the printer target). Pixels with lightnesses exceeding the upper limit of the bounding cylinder are soft compressed using a cut-off of 75%. Pixels with lightnesses below the lower limit of the bounding cylinder are clipped.

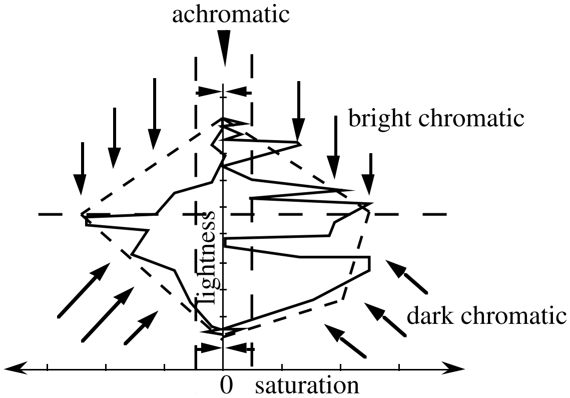

Figure 3. Illustration of Part II of the mapping technique. The dashed lines represent the target gamut. The shaded area is the bounding cylinder. The solid lines represent the input gamut. Achromatic pixels are mapped by changing saturation only. Bright chromatic pixels are mapped by changing lightness only. Dark chromatic pixels are mapped by simultaneously changing lightness and saturation.

Part II of the mapping strategy is illustrated in Figure 3. In this part, pixels within the cylinder boundary are mapped to the target gamut boundary. The input range is divided into 3 sections: a (near) achromatic section which we take to be a regular cylinder of small radius about the central axis; an upper chromatic section which is the area outside the achromatic cylinder having lightnesses greater than a middle lightness (which we take to be the lightness of greatest saturation); and a lower chromatic section which is the complement of the upper chromatic section outside the achromatic cylinder. A different type of mapping is performed in each of these sections.

We map out-of-gamut pixels in the achromatic section by compressing in saturation while keeping lightness constant. This tends to preserve contrast of achromatic pixels. We map upper half chromatic pixels by compressing in lightness while keeping saturation constant. This tends to preserve colorfulness.

Our mapping strategy for the lower chromatic section is image dependent. Pixels in this section are always mapped by simultaneously changing lightness and saturation. However the ratio of lightness changes to saturation changes had to be parameterized. For most images, a rate of 2 lightness units for each saturation unit worked well. However, for images containing dark reds, a rate of 10 lightness units for each saturation unit gave much better results. Keeping the ration 2/1 helps retain image contrast (*i.e.* it reduces any hazy appearance). Setting the ratio to 10/1 retains colorfulness (*i.e.* prevents chromatic colors from being mapped to black). We note that keeping saturation constant while changing lightness does not work well in the dark chromatic section because the image tends to take on a speckled appearance due to large changes in lightness.

Besides the lightness/saturation change ratio of the dark chromatic section, we found that we also had to parameterize the achromatic cylinder radius. For most images, a radius of anywhere from 0.3 to 1.0 saturation units worked well. However, for images containing bright sky scenes, a radius of 1.0 worked much better in that it retained the contrast of cloud formations. At least 2 of the images (#18 and #22) required a radius of 0.3 in order to avoid noticeable loss of saturation. We experimented with a linearly varying achromatic radius, narrow at low lightnesses and wider at high lightnesses. This did not give any improvement.

Observations and Comments on the Mapping of the Sample Images

In this section we compare the effectiveness of our mapping technique to scaling and RGB clipping for the monitor target and to saturation compression for the printer target, over a few sample images. For each image we state the best choices of parameters (achromatic radius and slope of simultaneous lightness/saturation change in dark chromatic section) and comment on the results.

image #12 (couple on beach):

best achromatic radius:1.0

best lightness/saturation change ratio: 10

For the monitor target, our method is superior to RGB clipping. The clipping method loses all detail in the clouds. We prefer our mapped image to the scaled image because it retains colorfulness and detail but is brighter than the scaled version. For the printer target, our method is superior to saturation only mapping because it retains the colorfulness of reds.

image #18 (woman in black dress near metal sculpture):

best achromatic radius: 0.3

best lightness/saturation change ratio: 2

For the monitor target, our method is preferred to both the other methods. The RGB clipping produces some anomalies in the flesh tones. The scaled version is darker than our version. For the printer target, our method is superior to saturation only mapping, which produces undesirable highlights in flesh tones. Note that this is one image which works well only with the small achromatic radius. Flesh tones are altered (as in the saturation-only) mapping if the radius is made larger. The lightness/saturation change ratio must be kept small to prevent a hazy appearance. The value of 2 is a trade-off between lower values, which give better contrast, and higher values which retain more colorfulness and detail in tree leaves.

image #22 (red barn):

best achromatic radius: 0.3

best lightness/saturation change ratio: 2

For the monitor target, our method is superior. The scaled version is darker and retains less contrast than our version, although it retains texture in the roof shingles which is somewhat lost by our method. The RGB clipped version changes the hue in the roof and also appears to have less contrast than our version. For the printer target, our method is superior mainly because of the unacceptable highlights in the specular roof reflections in the saturation only mapping. This image required an achromatic radius of 0.3 in order to remove the undesirable roof highlights.

image #23 (macaws):

best achromatic radius: 1.0

best lightness/saturation change ratio: 2

For the monitor target, our method is preferred. It retains better contrast and colorfulness than either the RGB clipped or scaled versions. For the printer target, our method is superior. It retains all colorfulness and loses only some detail. The saturation only mapping loses colorfulness of bright yellow (the yellow bird becomes a mixture of yellow and white).

Conclusions

This work formulates a repeatable methodology for studying gamut mapping techniques. There remain many open questions in gamut mapping. An important one is whether it is possible to completely automate the process. At least one author thinks that it is not[1]. We are not so sure. Our work shows that different areas of the color space have different preferred mapping directions. Further refinements of our technique may show that it is indeed possible to define an acceptable mapping strategy that requires no human intervention.

References

1. Fairchild, M.D., “Some hidden requirements for device-independent color imaging,” *SID 94 Digest*, 865-868, 1994
2. McCann J., “Color sensations in complex images,” *Proc. IS&T/SID Color Imaging Conference: Transforms and Transportability of Color*, **1**, 16-23, 1993.
3. Wyszecki, G. and Stiles, W.S., *Color Science*, Wiley, 1982
4. Fairchild, M.D. and Berns, R.S., “Image color-appearance specification through extension of CIELAB,” *Col. Res. and Appl.*, **18**, 178-190, 1993.
5. Guth, S.L., “Unified model for human color perception and visual adaptation II,” *Human Vision, Visual Processing, and Digital Display IV*, SPIE vol. **1913**, 440-448, 1993.
6. Hunt, R.W.G., “Revised colour-appearance model for related and unrelated colours,” *Col. Res. and Appl.*, **16**, 146-165, 1991.
7. Nayatani, Y., K., Takahama, K., Sobagaki, H. and Hashimoto K,, “Color-appearance model and chromatic adaptation transform,” *Col. Res. and Appl.*, **15**, 210-221, 1990.
8. Stone, M.C., Cowan, W.B., and Beaty, J.C., “Color gamut mapping and the printing of color images” *ACM Trans. on Graphics*, **7**, 249-292, 1988.
9. Gentile, R.S., Walowitt, E., and Allebach, J.P., “A comparison of techniques for color gamut mismatch compensation,” *J. Imaging Tech.*, **16**, 176-181, 1990.
10. Pariser, E.G.., “An investigation of color gamut reduction techniques,” *IS&T Symposium on Electronic Prepress Technology - Color Printing*, 105-107, 1991.
11. MacDonald, L.W., “Gamut mapping in perceptual color space,” *Proc. IS&T/SID Color Imaging Conference: Transforms and Transportability of Color*, **1**, 193-196, 1993.
12. Hoshino, T. and Berns, R.S., “Color gamut mapping techniques for color hard copy images,” *Device Independent Color Imaging and Imaging Systems Integration*, SPIE vol. **1909**, 152-164, 1993.
13. Kodak Photo-CD Photo Sampler, Distributed by Apple Computer for Eastman Kodak Company, c. 1991.
14. Eastman Kodak Company, *A Planning Guide for Developers - Photo CD Products*, Rochester NY, 1992.
15. “Basic Parameter Values for the HDTV Standard for Studio and International Programme Exchange,” CCIR Recommendation 709, 1990.
16. Hultgren, B.O., “Subjective quality factor revisited,” *Human Vision and Electronic Imaging: Models, Methods and Applications*, SPIE vol. **1249**, 12-22, 1990.
17. Roufs, J.A.J, Koselka, V.J.F., and van Tongeren, A.A.M., “Global brightness contrast and the effect on perceptual image quality,” *Human Vision, Visual Processing, and Digital Display V*, SPIE vol. **2179**, 80-88, 1994.