Development and Evaluation of Six Gamut-Mapping Algorithms for Pictorial Images

Karen M. Braun, Raja Balasubramanian and Reiner Eschbach
Digital Imaging Technology Center, Xerox Corporation
Rochester, New York, USA

Abstract

Six techniques for mapping the colors of an image into the gamut of printable colors were compared. Six pictorial scenes were used in two psychophysical experiments, one to test accurate reproduction and one to test preferred reproduction. A new contrast-enhancing algorithm was found to give more favorable reproductions than several gamut-mapping techniques described in the literature. This algorithm performs luminance compression by applying an inverted power function to images in a linear *RGB* color space: $1 - (1 - RGB)^{\gamma}$. Remaining out-of-gamut pixels are clipped to the gamut surface in the direction of a central point on the neutral axis.

Other algorithms that performed well were those that clip out-of-gamut colors to the surface of the gamut, and do not affect colors within the gamut. These algorithms can sometimes result in undesirable artifacts for certain images, including contouring and loss of shadow detail. However, observers did not object to the loss of shadow detail if the colorfulness of the image was maintained or increased.

Also, the results of a matching experiment (original present) and a preference experiment gave quite different results. Clipping algorithms did well in the matching experiments, while contrast boosting algorithms did best in the preference matching. The preferred techniques did well in both experiments.

Introduction

The goal of color reproduction is to deliver matching or preferred images on different devices. Output devices, including photographic film, ink jet printers, xerographic copiers, and CRT displays, are limited in the range of colors they can produce. A good gamut-mapping algorithm achieves the best compromise among image contrast, shadow and highlight detail, vividness, and smoothness of transitions. There are many ways to map out-of-gamut colors into the printer gamut. A simple method is to replace them with the nearest in-gamut colors, and leave in-gamut colors alone. This technique, called *clipping*, often leads to unwanted artifacts such as the apparent flattening of curved surfaces, and loss of detail information due to a “many-to-one” type mapping. Other algorithms compress both in- and out-of-gamut pixels so that smooth transitions in images are preserved. If done incorrectly, this can lead to desaturation or loss of contrast in the images, which many people find undesirable. This study generated new and better techniques for gamut mapping, and provided insight into the best compromises among these color quality attributes for development of future techniques.

A psychophysical study was conducted to test the quality of six algorithms for gamut mapping. Four of these algorithms consisted of transforming the data in some way followed by clipping to the surface of the gamut; two of these four transformations acted on CIELAB values and two on *RGB*. The final two algorithms consisted only of clipping. The best results were obtained using a power function on an inverted *RGB* image, then inverting again. This algorithm boosts the mid-tone contrast slope to better match the perceived contrast in the original.

Algorithms Studied in this Work

A two-step process was used to apply gamut-mapping algorithms. First, colors in the original image were processed through various transformations intended to compress input lightnesses to better fit within the output gamut. The lightness mapping occurs in this first step, in order to match the dynamic range of the input and output devices. Chroma was essentially clipped using this paradigm. Several techniques for gamut mapping were tested informally and rejected for inclusion in the psychophysical experiments, because they had obvious failure modes. Chroma compression while keeping lightness constant is often used in practice, but severe desaturation of the images was observed, particularly in light yellow and dark blue regions of color space.

Six algorithms were tested for gamut mapping, using knowledge of the literature and observations on images. The six algorithms are described in detail. The white points of the gamuts were equivalent, with L^*_{paper} equal to 100. Therefore, no gamut compression needed to be done at the light end of the gamut.

Nearest-Point Clipping

This technique involves mapping every out-of-gamut color to the closest CIELAB point on the surface of the destination gamut.[1] Colors within the destination gamut remain unchanged. No pre-processing was performed for this algorithm. Figure 1 shows an example of this type of clipping.

Centroid Clipping

Straight centroid clipping involves mapping out-of-gamut colors to the surface in the direction towards a fixed point along the neutral axis (Fig. 1).[2]

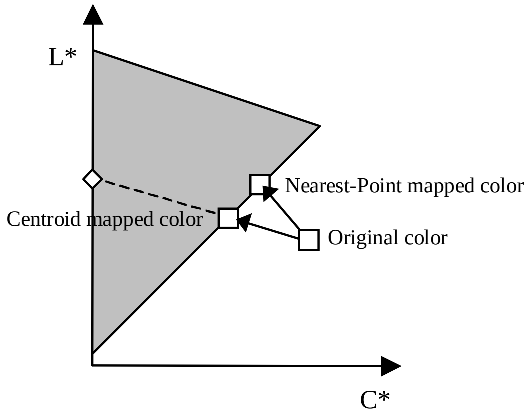

Figure 1. Two clipping algorithms, one that maps to the closest point on the surface of the destination gamut and another that maps toward a centroid on the neutral axis.

Chroma is reduced by this technique compared to nearest-point clipping, as seen in Fig. 1. This was mitigated by using a centroid region instead of a centroid point, such that light colors were mapped towards a darker point on the neutral axis, and darker colors were mapped towards a lighter point on the axis. A centroid range of 10 *L** units was used in this experiment.

*L** Compression (followed by Nearest-Point Clipping)

This algorithm involves linearly scaling the *L** values to match the dynamic ranges, by the function in Eq. (1).[1]

$$L*_{out} = \frac{100 - L*_{out,\min}}{100} * L*_{in} + L*_{out,\min} \tag{1}$$

where $L*_{in}$, $L*_{out}$, and $L*_{out,min}$ are the lightnesses of the input pixel and output pixel, and the minimum lightness of the output device, respectively. Lightness scaling has been used widely throughout the literature.[3-7] This pre-processing step was followed by nearest-point clipping.

Weighted *L** Compression (followed by Centroid Clipping)

Weighted *L** compression is a variation on the *L** compression algorithm, in which the amount of compression depended on chroma as well as *L**. In general, strong *L** compression is required in the dark neutral regions to maintain shadow detail. However, the same *L** compression applied to high-chroma colors followed by clipping results in loss of chroma. The new technique gradually blends from linear *L** compression along the neutral axis to a less aggressive *L** compression function as the chroma, *C**, increases. The intent was to find an acceptable trade-off between maintaining shadow detail and maintaining chroma. For the experiments described in here, Eq. (2) was used.

$$L*_{out} = weight*L' + (1- weight)*L*_{in} \tag{2}$$

where *L'* and *weight* are as shown in Fig. 2. This is related to the GCUSP method described by Morovic et al.[7]

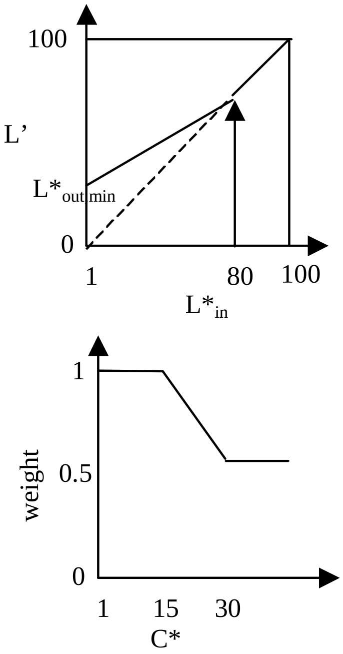

Figure 2. Weighting functions used to calculate destination L for weighted L* compression technique.*

Inverse-Power-Inverse (IPI) (followed by Centroid Clipping)

The inverse-power-inverse (IPI) technique maps colors in such a way as to try to preserve the appearance of contrast in an image. It operates in a linear colorimetric *RGB* space. In this experiment the following specification was used:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 2.944 & -1.461 & -0.457 \\ -1.095 & 2.026 & 0.036 \\ 0.078 & -0.272 & 1.4552 \end{bmatrix} * \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{3}$$

The technique requires information about the mismatch of the black values between the source gamut black and the darkest black of the destination device. If the input gamut black is unknown, an *L** of 0 can be used. The mapping function is given by Eq. (4) where gamma is calculated using Eqs. (5)-(7). The transformation in (4) is performed individually on the *R, G,* and *B* signals (where $0 \le R,G,B \le 1$) .

$$RGB_{new} = 1 - (1 - RGB_{orig})^{\gamma} \tag{4}$$

The gamma value, γ, is calculated in such a way as to map 95% of the input luminance range (1 – Y) to 95% of the output luminance range (1 – Y').

$$\gamma = \log(1 - Y_{95}') / \log(1 - Y_{95}) \qquad (5)$$

where,

$$Y_{95} = (1 - Y_{min}) * (1 - 0.95) + Y_{min} \qquad (6)$$

$$Y_{95}' = (1 - Y_{min}') * (1 - 0.95) + Y_{min}' \qquad (7)$$

A 95%-to-95% mapping was tested in the described experiments. The percentage value used may be optimized through further experimentation.[8]

Implementation of this algorithm was very straightforward, involving one-dimensional transformations directly to the individual *RGB* signals. The effect of these calculations on the image was to lighten and increase the chroma of dark colors, and lighten and decrease the chroma of light colors. The effect on the image lightness is one of the most valuable aspects of this technique. In Fig. 3, the *RGB* values of a neutral gray ramp ($R = G = B$) were converted to CIELAB values, and the source and destination L^* values were plotted. It can be seen that the IPI operation preserves or increases the slope of the source-to- destination lightness relationship. Thus mid-tone contrast is preserved or enhanced. The resulting effect, after clipping, is similar to the sigmodial tone-reproduction curves applied in Braun and Fairchild.[10]

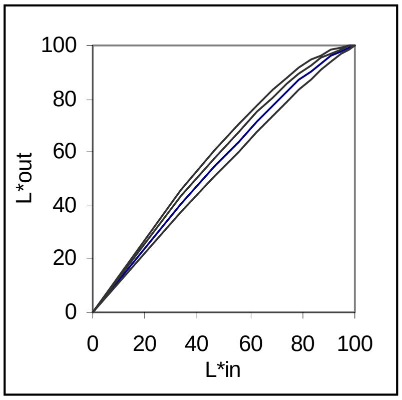

Figure 3. Effect on lightness, L, of IPI technique for gamma values of 1.25, 1.5, 1.75, and 2.0.*

One potential problem with applying this or any other algorithm in *RGB* space is the possibility of shifting the hue from the source to the destination color. Hue shifts only occurred for relatively high chroma colors, and in this region, observers are likely to be less sensitive to hue shifts. The hue-preserving IPI algorithm, to be discussed next, attempted to alleviate this problem by insuring that hue was preserved.

Hue-preserving IPI (followed by Centroid Clipping)

This algorithm maps colors to a smaller gamut while preserving perceived contrast and hue in the images. It requires information about the mismatch of the black values between possible inputs and the darkest black of the destination printer. It is similar to a technique described by Kasson for correcting midtones in *RGB* while preserving chromaticity.[10] Colorimetric *RGB* values of an image are multiplied by the ratio of destination luminance to source luminance. Destination luminance is determined by inverting, applying a power, and inverting source luminance values of each color in the image, as described in the previous sub-section.

The mapping function is given by Eq. (8) where γ is calculated using Eqs. (5)-(7). The transformation in (8) is performed individually on the *R*, *G*, and *B* signals (where $0 \le R,G,B \le 1$).

$$\begin{bmatrix} R_{new} \\ G_{new} \\ B_{new} \end{bmatrix} = \alpha * \begin{bmatrix} R_{orig} \\ G_{orig} \\ B_{orig} \end{bmatrix} \qquad (8)$$

where α is the ratio of Y' to Y, and Y' is calculated in Eq. (9).

$$Y' = 1 - (1 - Y)^{\gamma} \qquad (9)$$

As in the IPI algorithm, a 95%-to-95% mapping was tested in the psychophysical experiments described. The percentage value used may be optimized through further experimentation. The effect of these calculations on the image is to lighten and increase the chroma of colors while preserving hue.

It can be shown that the resulting effect on chroma is $C_{new} = \alpha^{1/3} C_{orig}$. The output minimum L^* is larger than the input minimum L^* for dynamic range compression, so $\alpha > 1$, and thus $\alpha^{1/3}$ is greater than 1. Therefore chroma is increased through the operations. (This is in contrast with the IPI method, which decreases chroma for higher-lightness colors.) The effect on the image lightness is identical to what was seen from the IPI technique, and shown in Fig. 3.

Gamma was 1.58 for both the IPI and hue-preserving IPI techniques. This was based on the source and destination gamuts used in the experiment and discussed in the Experiment section.

Experiment

The six gamut-mapping algorithms under investigation were tested in two psychophysical experiments. Part 1 was intended to simulate a printing environment where the original is not present for comparison. Part 2 simulated a copying environment and observers compared gamut-mapped images to an original, uncompressed image. Before describing the details of the specific psychophysical experiments, image selection and preparation (the same for both parts) will be discussed.

Original images were compared to reproductions that had been mapped to fall within a simulated gamut on the same printer. Source and destination gamuts were both for hardcopy devices. This eliminated any differences in the look and feel of the prints and reduced problems associated with the calibration error of the printer. The printer used in this experiment was the Xerox 5760 MajestiK printer, with an average CIELAB color reproduction error (ΔE^*_{ab}) of 4.04, with a maximum color difference of 9.83. Six scenes were mapped using centroid clipping to fall within the entire gamut of this printer and these became the originals to which reproductions would be compared. L^*_{min} was 4.8 for the source gamut and 22.8 for the resulting destination gamut.

Six scenes were used and included a wide range of content spanning various characteristics of real images: chromatic and non-chromatic content; various skin tones; natural and synthetic content; memory and non-memory colors. They are referred to as Fruit, Lighthouse, Macaws, Sungirl, Trees, and Wetgirl. No graphics images were included in this study.

The original images were clipped using centroid clipping with a range of 16 L^* units to fall within the input gamut. These clipped images were used as the originals, and subsequent processing included this clipping. Reproductions were made by passing the originals through LUTs corresponding to the various gamut-mapping algorithms. The images were labeled with randomized codes so that observers could not identify which algorithm they were viewing. They recorded only the codes, which were later cross-referenced with the algorithm name. The original images contained anywhere from 8% to 73% out-of-gamut pixels.

Data Analysis and Results

In this research, the comparative-judgment method of data reduction was used.[11] This method uses the idea that, in order to rank the stimuli, observers are actually comparing each reproduction with every other one. With this assumption, an analysis similar to that of paired-comparison data can be used. The proportion of choice, p, for each gamut-mapping algorithm was calculated by dividing the mean choice, M_C, by one less than the number of stimuli, (n - 1). The mean choice is the total number of stimuli, n, minus the mean rank, M_R. Then the z-score was calculated for each proportion. z is the standard normal deviate for the given proportion. Equations (10) and (11) show the necessary equations.

$$M_C = n - M_R \tag{10}$$

$$p = M_C/(n-1) \tag{11}$$

Mean rank, M_R, is calculated by averaging the ranks assigned to a given algorithm. n is the number of stimuli; in these experiments, $n = 6$. The resulting z-score values represented an interval scale of each algorithm's quality. These values were calculated for each scene then averaged.

For part 1, a ranking experiment was performed in which thirty-eight observers were asked to order the gamut-compressed reproductions from the one they most preferred to the one they least preferred. Images were mounted on a wall where the lighting was relatively uniform across the images. The light source was a D50 simulator. Six observers claimed to be very experienced, 20 reported moderate experience, 7 reported little experience, and 5 reported no experience judging color images.

For part 2 of this study, a ranking technique was again employed. In this experiment, observers were given an original image and were asked to order the six gamut-compressed reproductions. They were instructed to assume that the original image was the one they were trying to duplicate and asked to rank the reproductions from the one they would be most satisfied with to the one they would be least satisfied with. Unlike the previous part, observers were allowed to physically rearrange the images to decide on their ranked order. The lighting was again a D50 simulator. Thirty-one observers performed this experiment. Seven observers considered themselves to be very experienced in judging the quality of color images, 9 reported moderate experience, 13 had little experience, and 2 had no experience. The interval scale analysis results are shown in Fig. 4.

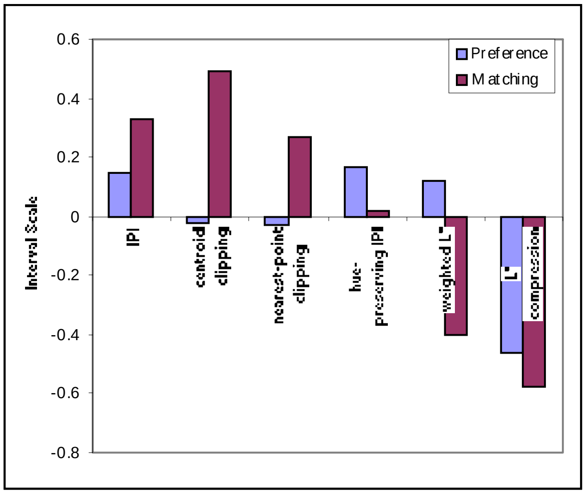

Figure 4. Interval scale results of Parts 1 and 2.

In part 1, L^* compression technique is least favorable by far. The IPI and hue-preserving IPI techniques performed better than the clipping algorithms, but this depended on image content. Since originals were not available in this part of the experiment, it is unclear whether the clipping techniques performed worse than the IPI techniques because (a) the IPI techniques lightened the images (which may have been too dark to begin with) or (b) the clipping techniques caused some unfavorable artifacts. This dependence of image content and quality of the original will be discussed in more detail. The IPI techniques and the weighted L^* compression were approximately equal to each other and better than the clipping methods. Figure 4 was the result of averaging the interval scales for all five scenes (using the second Lighthouse original).

In part 2, centroid clipping gave the most preferred results on average, followed closely by IPI and the nearest-point clipping. Because this was a side-by-side comparison, it is not surprising that clipping algorithms did so well, given that they are colorimetric matches to originals except in out-of-gamut regions. This may account for the relatively lower ranks of the weighted L^* compression technique and the hue-preserving IPI technique.

The success of the clipping algorithms appeared to be bimodal, with large numbers of ranks 4, 5, and 6. The number of times each of these three techniques was ranked in the top three is shown in Table 1. The IPI technique seems to be more stable than the clipping techniques, resulting in more consistently good images. In several images, the clipping algorithms showed significant artifacts.

Table 1. The number of times each of the top three techniques was given a ranking of 1-3 and 4-6 in Part 2.

Top Algorithms	*1-3*	*4-6*
Inverse-Power-Inverse	133	53
Centroid clipping	111	75
Nearest-point clipping	107	79

Summary

The two parts of the experiment, with and without an original present, resulted in different algorithms being chosen as best. With no original, the algorithms that lightened the image were preferred, including hue-preserving IPI, IPI, and weighted L^* compression. This was probably because the originals from which the images were mapped were darker than observers would have preferred. These technique may in fact have been enhancing the original. When an original was present for comparison (part 2), clipping algorithms were favored, as well as the IPI technique. Linear L^* compression was not preferred in either experiment as it gave reproductions a “washed out” appearance.

IPI with centroid clipping gave good, consistent results in both experiments. Nearest-point clipping produced more chromatic images, but was also more prone to artifacts than centroid clipping.

As with any gamut-mapping technique, it is expected that the described technique will give improved results if the statistics of the given image are taken into account. For example, if the image colors are completely within the destination gamut, then no IPI correction should be done before clipping, equivalent to setting gamma equal to 1.0. Improved results would likely be seen if gamma were made a function of the darkest image color instead of the darkest expected color. However image-dependent algorithms are much more computationally intensive since they can not be integrated into a generalized look-up table.

References

1. R. S. Gentile, E. Walowit, and J. P. Allebach, A Comparison of Techniques for Color Gamut Mismatch Compensation, *SPSE/SPIE Symposium on Electronic Imaging,* Los Angeles, pp. 176-81, 1989.
2. E. G. Pariser, An Investigation of Color Gamut Reduction Techniques, *IS&T 2nd Symposium on Electronic Publishing*, pp. 105-107, 1991.
3. M. C. Stone and W. E. Wallace, Gamut Mapping Computer Generated Imagery, *Graphics Interface ’91*, pp. 32-39, 1991.
4. J. A. S. Viggiano and J. Wang, A Comparison of Algorithms for Color Matching Between Media of Differing Luminance Ranges, *TAGA Proceedings* **2**, pp. 959-974, 1992.
5. L. MacDonald and J. Morovic, Assessing the Effects of Gamut Compression in the Reproduction of Fine Art Paintings, *Proc. of the 4th IS&T Color Imaging Conference,* pp. 194-200, 1995.
6. E. D. Montag and M. D. Fairchild, Psychophysical Evaluation of Gamut Mapping Techniques Using Simple Rendered Images and Artificial Gamut Boundaries, *IEEE Trans. Image Processing* **6**, pp. 977–989, 1997.
7. J. Morovic and M. R. Luo, Verification of Gamut Mapping Algorithms in CIECAM97s Using Various Printed Media, Proceedings of 6th IS&T/SID Color Imaging Conference, Scottsdale, pp. 53-61 (1998).
8. A patent application has been filed by the authors with the U.S. Patent Office on inverse-power-inverse technique for gamut mapping.
9. G. J. Braun and M. D. Fairchild, Image Lightness Rescaling using Sigmoidal Contrast Enhancement Functions, *Proc. of IS&T/SPIE Color Imaging: Device-Independent Color, Color Hardcopy, and Graphic Arts IV* **3648**, 96-107, 1999.
10. J. M. Kasson, Efficient, Chromaticity-Preserving Midtone Correction for RGB Images, *Proc. of the 2nd IS&T/SID Color Imaging Conference*, pp. 78-84, 1994.
11. C. J. Bartleson and F. Grum, *Optical Radiation Measurements, Vol. 5: Visual Measurements,* Academic Press, New York, 1984.