Gamut Mapping Computer Generated Imagery

Maureen C. Stone
Xerox PARC
3333 Coyote Hill Rd.
Palo Alto, CA 94304, USA

William E. Wallace
University of Waterloo
Waterloo, Ontario
Canada N2L 3G1

Abstract

Much computer graphics research involves the design of images on color monitors. Reproducing these monitor images as prints or slides is difficult because color is generated and viewed differently on different media. Device independent color standards provide a way to represent individual colors in a form that is independent of any particular method for creating the color. However, these representations are based on standards established for identifying single colors viewed on a neutral background or quantifying small differences between similar colors. When reproducing images or illustrations, the relationship between colors is often more important than the fidelity of individual colors. Furthermore, different devices can reproduce only a subset of the visible color space, the device *gamut*, so out-of-gamut colors must be transformed to reproduceable colors. The process of maintaining the overall appearance of an image while mapping the colors to fit the gamut of the target device is called *gamut mapping*. Our approach is to non-linearly adjust the image colors in lightness to control the dynamic range, and in chroma to bring overly saturated colors inside the target gamut. We perform our calculations in the CIE L*a*b* color space. This paper will describe the benefits and limitations of this approach.

1 Introduction

Printing computer generated images designed for monitor viewing in such a way that the image looks "right" when printed is difficult. Printers differ from each other and differ from monitors both in the specification of colors and in the gamut available. One approach to digital image reproduction is to define and adjust the colors relative to each other in a device-independent, perceptually based color space such as CIE L*a*b* [6, pp 166–169]. A function called a *calibration* maps the device color space to and from the device independent color space. The calibration specifies the *gamut* for the device. Typical device gamuts for a monitor, printer and transparency film are shown in figure 1.

Once identified, out-of-gamut colors must be modified to lie inside the target device gamut. There are many ways to make these modifications, and it is difficult to define quantitatively the best method because the parameters that affect color image appearance are not fully understood. The methods in this paper are derived from the following graphic arts principles: 1) Adjust the lightness to make detail visible throughout the tone range of the image. 2) Adjust the color without changing the hue names. 3) Maintain or slightly enhance the overall saturation of the image relative to the output gamut. These principles are encoded as geometric transformations in CIE L*a*b*.

2 Gamut Mapping

A function acting on a color space to transform colors from one gamut to another is called a *gamut map*. A simple example is a *tone reproduction curve* that maps input lightness values to output values. This particular gamut map affects only the lightness of an image and depends only on lightness. More generally, lightness, hue and saturation can all be modified as a function of the original lightness, hue and saturation. That is, a gamut map is a function $(h, s, l)' = f(h, s, l)$ or, similarly, $P' = f(P)$ where P is a point in any color space. The purpose of a gamut map is to enhance the appearance or to preserve the appearance of an image when the image is to be reproduced on a different device.

By preserving the appearance of an image, we mean that when the image is moved from one device to another, the image "looks the same" relative to the normal viewing conditions for the given devices. There are no simple, quantitative metrics for color appearance in complex images. The following are some guidelines. The lightness scale should be mapped to preserve detail in both the light and dark regions. Colors that were distinct in the original images should remain distinct. Changes in hue are usually more noticeable than changes in saturation. For example, if the background color was previously a

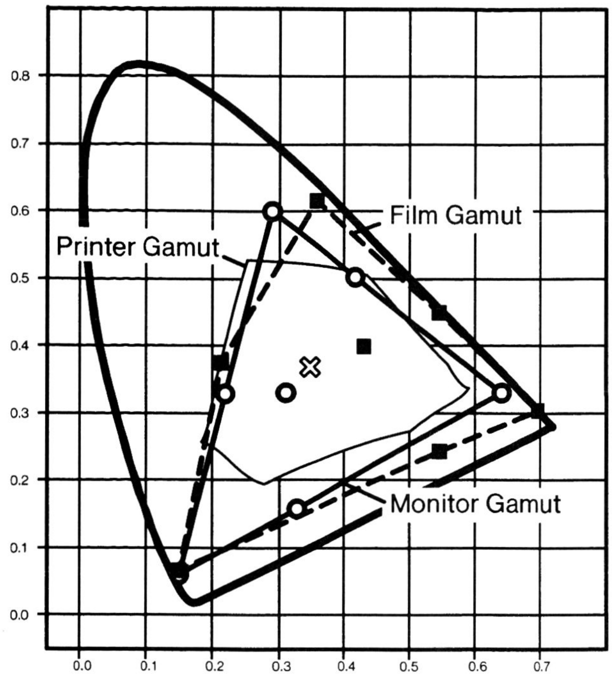

Figure 1: Typical printer, film and monitor gamuts plotted on the CIE 1931 chromaticity diagram. Each gamut has a typical white point marked with a corresponding token.

saturated blue then it is less noticeable to change it to slightly less saturated blue than to change it to a saturated blue-purple.

Mathematically, we can describe these principles by examining the magnitude of the directional derivatives of $f(P)$. If the magnitude is zero or almost zero, then some region of the color space is all being mapped to a single color. This removes any color gradations in that area. If the magnitude is instead very large (or $f(P)$ is discontinuous at P) then contouring occurs in a region surrounding P. The contouring occurs because the colors are effectively quantized; a minimal change in P in adjacent areas produces a large, noticeable, change in the color. The components of the directional derivative of $f(P)$ with respect to a vector V should have the same sign as the components of V. This ensures that increasing any of lightness, hue and saturation in the original image also increases it in the final image.

The perception of hue is more difficult to define mathematically. One interpretation of the CIE L*a*b* color space defines lines of constant hue which can be used to constrain the derivative of $f(P)$. However, the perception of hue is not well defined by these lines over the entire space, especially for highly saturated colors that lie near the edges of typical monitor gamuts. Exploring more exact specifications of hue for use in digital color reproduction is an area for future research.

A gamut map can be defined to map a particular image to a device, or it can map all colors from one device onto another device. If an image does not fill the entire device gamut, an image gamut map may transform the colors less than a device-to-device gamut map would. An image gamut map provides the flexibility to tune the color reproduction to produce the best reproduction for a specific image. However, a gamut map must be generated for each image as it is reproduced. A device-to-device gamut map is created only once. When images are being compared (or simply printed on the same page), it may be important to use the same gamut map for the set of images to get comparable results. This can be either the device-to-device gamut map or a special gamut map for the image set.

2.1 Mapping Lightness

In absolute terms, monitors are much more dim than printed pages under normal viewing conditions. The darkest black printable may be lighter (under some given illumination) than the brightest light given off by a monitor. Therefore, it is not the absolute lightness but some relative measure that needs to be considered. In the L*a*b* system, lightness is measured relative to the brightest achromatic color. When transformed to L* values, the point corresponding to "white" is always L*=100. However, the black point is in different positions relative to the white point on different devices. Typical values for a monitor black are L*=2 or 3. For a printer black, we have seen values as high as L*=25.

Currently, we are linearly scaling the L* axis and translating it so that the L* value for black on input is equal to or slightly less than the minimum L* value that is black on output. Using a value less than the true black value means that the darkest colors are projected to a point on the surface of the destination gamut. This will produce images with improved contrast compared to exactly matching the black values, at the cost of detail in the dark regions. Images that are not very dark tend not to lose much detail using this method. With very dark images, the black point should be matched exactly so that minimum detail is lost. Note that a print gamut is much more narrow around the black point than the monitor gamut, so some compression of the colors in the dark areas is inevitable.

A linear map to scale and translate the monitor L* to the printer range for L* is adequate if the image already appears satisfactory on the monitor and the monitor calibration is accurate [3]. However, there are many other ways to map the lightness axis. A tone reproduction curve, encoded as a lookup table from input pixel to

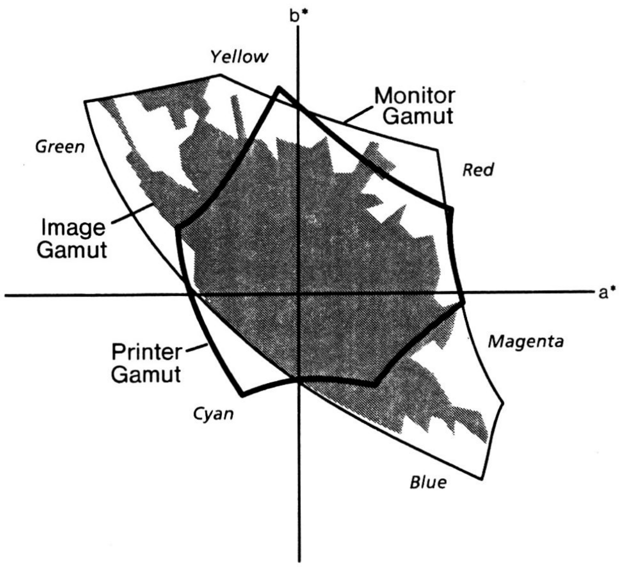

Figure 2: Typical monitor, printer and computer generated image gamuts plotted on the a*b* plane.

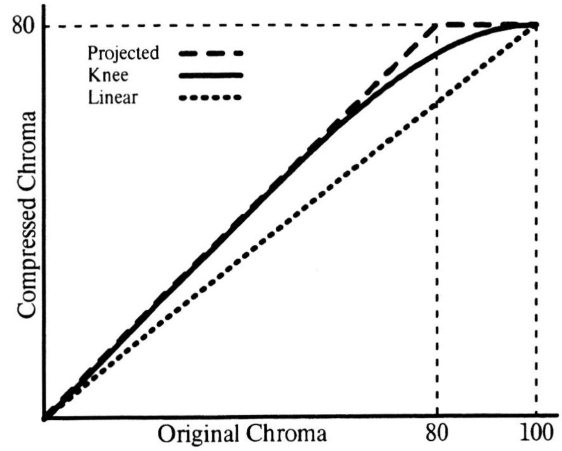

Figure 3: Output chroma as a function of input chroma for three different mappings.

output brightness, can be used to globally modify lightness or contrast. Other authors have suggested a local function that both adapts the lightness scale and provides edge enhancement [4]. Whatever method is chosen, correctly mapping lightness is the most critical part of gamut mapping.

2.2 Mapping Chroma

One common problem when reproducing images, especially computer generated images, is that different devices have different maximum saturations for colors. Figure 2 shows typical computer graphics monitor, printer and image gamuts mapped on the a*, b* plane. The gray filled region represents the colors used in the image while the two curves represent the boundaries of the monitor and printer gamuts. In this figure, the lightness axis is pointing out of the plane of the paper. Colors on the L* axis are neutral gray. Colors increase in saturation moving away from the L* axis. Interpret a*, b* as hue (angle) and chroma (radial distance from the L* axis). Compressing chroma along lines of constant hue provides a method for mapping such an image gamut into a print gamut without changing the hue, within the limits of the hue definition for the uniform color space.

The simplest form of chroma compression projects all out-of-gamut values to the surface of the target gamut. While this maintains the maximum image saturation, it can cause shading artifacts because multiple colors are mapped to the same point (some of the directional derivatives of the gamut map are 0). Another approach is to map all the values linearly along lines of constant hue. This maintains the shading, but can cause an objectionable decrease in saturation, especially if there is a mix of very saturated and pastel colors along the same hue line. The mapping necessary to bring the saturated colors into gamut maps the pastel colors towards gray so that the color disappears. This produces achromatic regions in the picture where there was previously color.

The approach we have found most satisfactory so far is a function that is tangent to a linear function near the gray axis, and then approaches a horizontal (projection) line near the maximum output saturation. We call this a *knee function.* The linear function need not have a slope of 1. For some images, we found that boosting the saturation (slope > 1) in the pastel regions produced a more pleasing picture. These three mappings are shown graphically in figure 3.

The effect of these functions on a simple computer graphics image is shown in figures 4-7. The scene is a simple, three-sided room with two spherical bumps on the floor. One wall is bright pink (R:1.0, G:0.2, B:1.0) and the other is bright blue (R:0.2, G:0.4, B:1.0). The floor is light gray (R:0.8, G:0.8 B:0.8). The color of one sphere is a desaturated version of the pink wall color (R:0.65, G:0.5, B:0.65) and the other is bright yellow (R:1.0, G:0.9, B:0.3). There is an invisible, white, point light source near the corner of the room. All surfaces are matte; there are no shadows or highlights. On the monitor image, the walls each appear to have a circular spot of pure color near the light source shading towards a dark color in the corners. Similarly, the spheres show their brightest color near the light. The pink sphere shades to black on the side away from the light. The illustration was printed on a 254 spots/inch thermal sublimation printer that nominally prints 256 lightness levels for

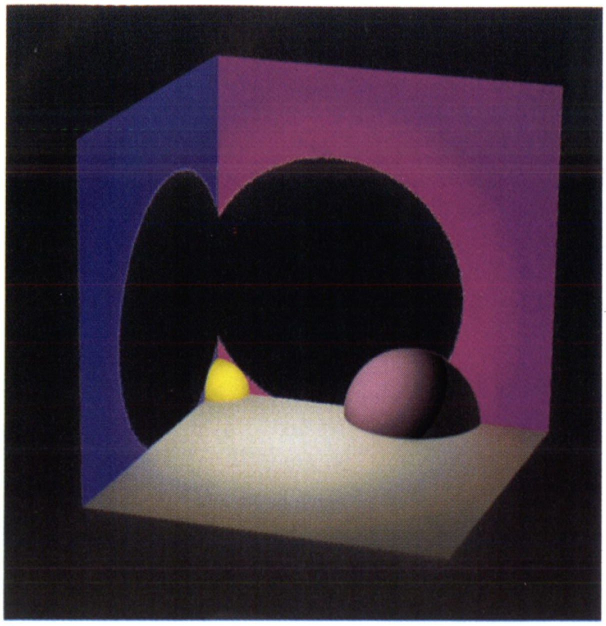

Figure 4: In-gamut colors are duplicated, out-of-gamut colors are black.

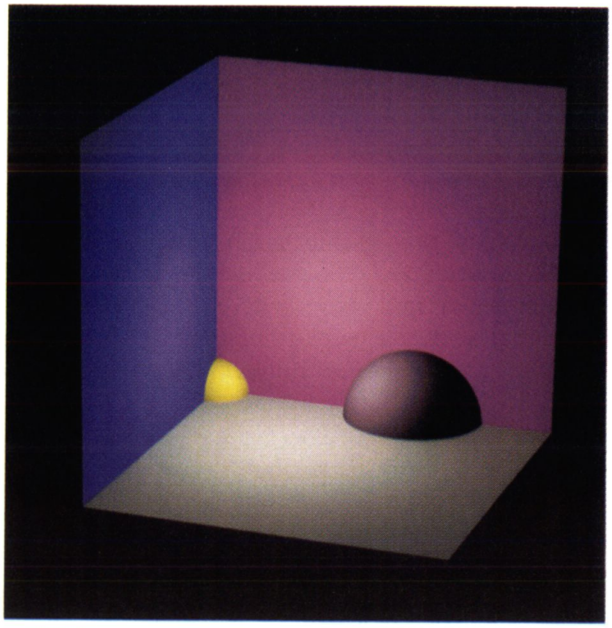

Figure 6: Out-of-gamut colors are linearly scaled along lines of constant hue.

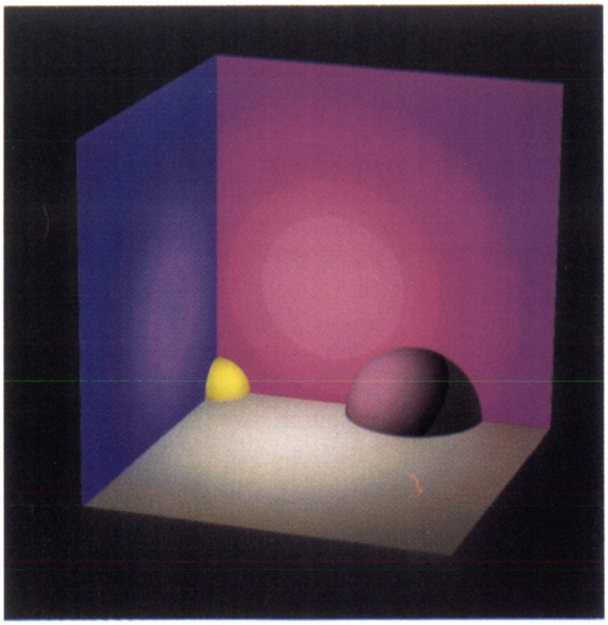

Figure 5: Out-of-gamut colors are projected to the gamut surface.

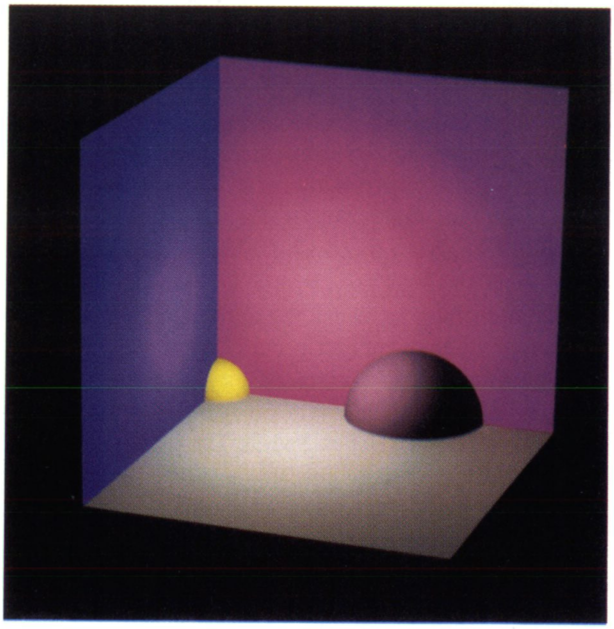

Figure 7: A knee function is applied to the out-of-gamut colors along lines of constant hue.

each of cyan, magenta and yellow. The image was then scanned and halftoned for printing in the proceedings. The printer gamut is quite similar to the offset printing gamut so color changes and other artifacts from the reproduction process should (we hope) be minimal.

In each of the illustrations, the colors have been converted to L*a*b* using the monitor white point (R:1.0, G:1.0, B:1.0) as the reference white. The L* component of the colors have been scaled to fit the range of the printer gamut. Both gamuts have a maximum L* value of 100 (by definition). The monitor minimum, measured in a completely dark room, is 2.5 and the printer minimum, measured in a graphic arts lighting booth, is 25. This is a fairly light value for black. The corresponding value for offset printing is around 10.

In figure 4, in-gamut colors are duplicated and out-of-gamut colors are printed as black. There are no out-of-gamut values on the yellow ball or the gray floor. There is a narrow line of black on the pink sphere where the printer gamut abruptly narrows near black. In figure 5, out-of-gamut colors have been projected to the nearest color of the same lightness and hue. The sudden change in the color cause contour lines. Note also that the brightest sections of the walls are pastel rather than vivid. This is because the saturated pink and blue colors on a monitor are much lighter than the corresponding hues on the printer. The only printer colors of the same hue and lightness are very desaturated.

Figures 6 and 7 demonstrate linear and knee transformations of chroma. The shading in the linear case is smoother, but the colors are less saturated. Note particularly the pastel pink ball which is nearly gray in the linearly scaled figure but retains its pink color in the knee version. Neither has noticeable contour lines because both are smooth transformations. For typical, highly saturated computer graphics images, the knee transformation seems to give a more satisfactory result than linear scaling.

To demonstrate the effect of gamut mapping on a more realistic, complex computer generated image, we will use the cover of the book "Graphics Gems" [2], part of which is shown in figures 8-12. The image was designed and produced by Thad Beier at Pacific Data Images, Sunnyvale, CA. The gamut of this image fills a large part of the monitor gamut, as shown in figure 2. Large areas of the gems image are outside the printer gamut. We generated the color separations for the book cover using the techniques described in this paper to compress the monitor gamut to the offset printing gamut.

Figures 8-11 correspond to figures 4-7 and were produced the same way. Figure 8 reproduces the out-of-gamut colors as black. Figure 9 contains artifacts where the colors change abruptly due to the projection. For example, look at the interior of the large, blue gem or in the dark regions of the bag. Figure 10 is clearly less saturated than figure 11 although some people prefer the overall appearance of the linear compression because the detail and shading seems more natural. The knee compression provides a compromise between shading detail and saturation. Compared to the monitor original, all print reproductions appear less saturated. Furthermore, some of the areas that look like large highlights are artifacts of maintaining lightness over saturation, as discussed previously. However, the print is a pleasing image that satisfied the book editor.

2.3 Sampled Gamut Maps

Many gamut maps are expensive to calculate and each pixel in an image must be mapped. Therefore, it is important for performance reasons to have a precomputed representation of the gamut map. The particular choice we made is multi-linear interpolation on a rectangular grid. The gamut mapping function is applied to each element in the grid, forming a three-dimensional table. Intermediate values are interpolated. Another choice is to use interpolation on a tetrahedral grid. Both of these forms of interpolation are discussed in chapter 11 of Akin's book [1].

The use of an intermediate table produces some smoothing at the gamut boundaries for any gamut mapping algorithm because the edge of the gamut is not guaranteed to lie on grid points. Consider a point just outside the gamut boundary. The cell containing this point has some corners inside and some outside of the gamut. Since we are using multi-linear interpolation to define the values inside of the cell, the value of all points inside of the cell are related by a smooth function. There is no way to define a sharp boundary inside of the cell. The extent of the smoothing produced by this approximation depends on the sampling density used. The illustrations in this paper used a 14 x 14 x 14 table for the linear and knee mappings. However, the illustrations showing black where the colors fell outside the gamut and those with out-of-gamut colors projected to the gamut surface had to be generated without a gamut mapping table to fully demonstrate an abrupt change of color at the gamut surface.

3 Results

The approach to gamut mapping presented here is quite general and provides good results for a wide range of images. However, this use of the CIE L*a*b* space is not one intended by the original designers of the system. Both CIE L*a*b* and CIE L*u*v* were designed to be

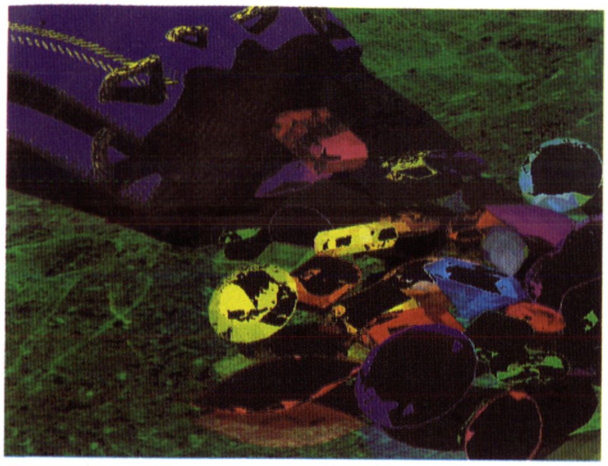

Figure 8: In-gamut colors are duplicated, out-of-gamut colors are black.

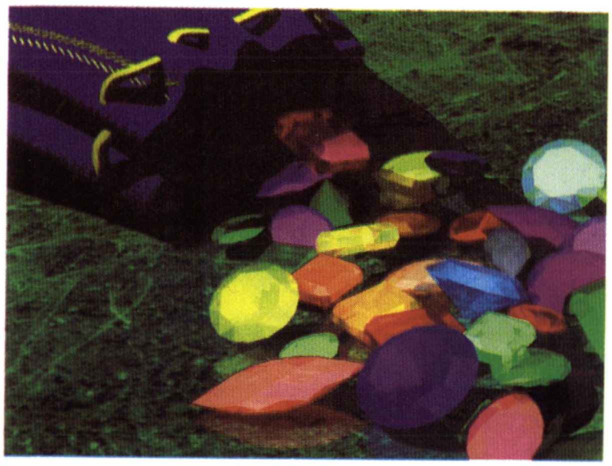

Figure 9: Out-of-gamut colors are projected to the gamut surface.

accurate for small color differences, or small, single digit values of ΔE [6]. Gamut mapping can involve color changes more than ten times larger. In our work, we have found examples where the use of CIE L*a*b* is inadequate, even in an engineering sense. Two significant problems are the assumption that radial lines are lines of constant hue and mapping CIE L*a*b* values computed relative to significantly different reference whites.

3.0.1 Lines of constant hue

In using the CIE L*a*b* color space, we have assumed that colors along radial lines vary only in saturation. While this is only an approximation to true lines of constant hue, we were interested to see whether if in practical situations this simple assumption was adequate. In practice, we have found that there is a problem for cer-

Figure 10: Out-of-gamut colors are linearly scaled along lines of constant hue.

Figure 11: A knee function is applied to the out-of-gamut colors along lines of constant hue.

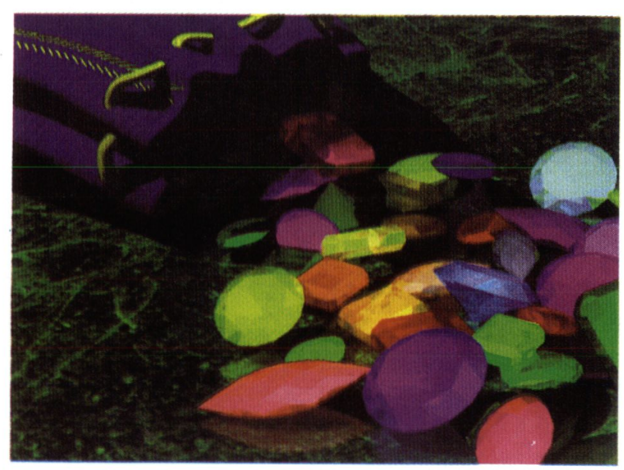

Figure 12: Gems image reproduced using a color temperature of 9300K as reference white when converting to L*a*b*.

tain radial lines in the saturated blue/purple region of the space near the color of the blue phosphor on the monitor. These lines change color from purple, near the gray axis, to blue, near the edge of the color space. In practice, this problem arises when an image with a saturated blue is measured on a monitor, desaturated to bring it into the print gamut and then printed. The color that appeared blue on the monitor appears purple on the print.

3.0.2 White point differences

The white point of an image displayed on a monitor is usually quite different from the white point on a printed page viewed under normal lighting conditions. Typical monitor white points, defined as (R:1.0, G:1.0, B:1.0), range from 6500K to 9300K, where 6500K is the NTSC broadcast standard. Printer white is defined as white paper under a standard lamp. The ANSI standard light for viewing prints has a color temperature of 5000K, and incandescent indoor lighting can be as low as 2000K. While some graphic arts monitors are set with a white point near 5000K, most computer graphics monitors are operated at a much higher color temperature to improve the brightness.

It is almost always desirable to map the monitor white point to the printer white, but the best way of doing this is not yet understood. One common solution is to generate the L*a*b* values for the print relative to the printer white and the monitor values relative to the monitor white point. We have discovered that this approach, especially when coupled with the problem of defining lines of constant hue described in the previous section, can produce disturbing hue shifts in the blue-purple region of the color space. Figure 12 shows this effect. The dark blue colors in the image, such as the colors in the bag and the large gem in the foreground, shift from blue to purple depending on the white point of the monitor. Figures 8–11 were mapped using a color temperature of 5000K for the monitor's reference white. In figure 12, a color temperature of 9300K was used, causing the bright blue colors to appear bright purple. These colors appeared blue on the monitor (though not the same shade of blue) at either white point setting.

4 Conclusions

The goal of image reproduction is to provide the best appearance for the image given a particular medium. Each different medium has its own set of standard viewing conditions, though in practical systems the viewing conditions vary significantly from user to user. Gamut mapping provides a general purpose mechanism for adjusting image colors in a device independent manner to best suit a particular medium.

The most difficult problem in this type of work is accurately evaluating the results of different gamut mapping algorithms, both numerically and perceptually. Analysis of a function can only be undertaken once a metric is established that defines the goodness or badness of the results. Using ΔE units in CIE L*a*b* space provides one type of metric, but does not cover many perceptual issues that affect image quality. Analytic methods for evaluating image appearance are still a basic research issue.

Within the context of computer generated images, we can imagine algorithms that do not depend strictly on the transformation of independent pixel colors in some color space, but that could use information from the original description of the picture to better define the best gamut mappings. For illustrations, such descriptions often have structure in the design sense, i.e., here is a palette of related colors. For 3D computer graphics, a scene is composed of objects and lights which are rendered according to some algorithm. Adapting the scene colors to the target device gamut before rendering would guarantee that there were no shading artifacts in the mapped image.

Finally, color fidelity is only one criterion for evaluating the quality of a reproduction. In the Graphics Gems example, the mapping chosen for the cover was similar to the one in figure 12. The facets of some gems such as the oblong, pink gem to the left of the large blue stone tended to disappear when the monitor white point was set to D5000. The designer chose to sacrifice color fidelity for better detail. While device independent color specifications and careful gamut mapping can automatically produce good results for a wide range of images, there will always be cases when the user will want to adjust the final result. Finally, for images with colors that are severely out-of-gamut, the only acceptable solution may be to redesign the picture.

5 Acknowledgments

Our special thanks to Rick Beach who has encouraged this work both as lab manager of the Xerox PARC Electronic Documents Laboratory and as Editor-in-Chief of ACM/SIGGRAPH. Andrew Glassner provided the simple test image used in section 2.2. Much of the original work on gamut mapping [5] was performed jointly with William Cowan and John Beatty, University of Waterloo, whose ideas are reflected here also. This work has been sponsored in part by ACM/SIGGRAPH.

References

[1] J. E. Akin, *Finite Element Analysis for Undergraduates*, Academic Press, 1986.

[2] A. Glassner, editor, *Graphics Gems*, Academic Press, San Diego, 1990.

[3] M. J. Lamming, W. S. Rhodes, "A Simple Method for Improved Color Printing of Monitor Images," *ACM Transactions on Graphics*, **9**(4), 1990, 345-375.

[4] J. Meyer, B. Barth, "Color Gamut Matching for Hard Copy," in the proceedings of the Optical Society of America meeting on *Applied Color Vision*, July 12-14, 1989, 138-143.

[5] M. C. Stone, W. B. Cowan, J. C. Beatty, "Color Gamut Mapping and the Printing of Digital Color Images," *ACM Transactions on Graphics*, **7**(4), 1988, 249-292.

[6] G. Wyszecki and W.S. Styles, *Color Science*, Second edition, John Wiley & Sons, New York, 1982.