Spatial Color-to-Grayscale Transform Preserving Chrominance Edge Information

Raja Bala and Reiner Eschbach
Xerox Corporation
Webster, New York, USA

Abstract

A color image sent to a monochrome output device must undergo a color-to-grayscale transformation. Such a transform typically retains the luminance channel or a derivative thereof. A problem with this approach is that the distinction between two different colors of similar luminance is lost. This loss can be particularly objectionable if the two colors are spatially adjacent. The paper describes a color-to-grayscale transformation technique that locally preserves distinction between adjacent colors by introducing high-frequency chrominance information into the luminance channel. This is accomplished by applying a spatial high-pass filter to the chrominance channels, weighting the output with a luminance-dependent term; and adding the result to the luminance channel. The outcome of this is that luminance variations are introduced into the image only in those regions containing high-frequency chrominance information. Regions with smoothly varying chrominance undergo little enhancement, and in particular, grayscale input is passed through without alteration. The spatially adaptive nature of the algorithm readily distinguishes it from standard approaches that apply global or pixelwise transformations. Also, preference experiments demonstrate the superior qualitative performance of the proposed approach over standard techniques.

Introduction

In certain applications, it is necessary to convert a color image to a grayscale representation. A common example is rendering color images to a monochrome device, or a color device in monochrome mode. Another example is a color scanner that exports grayscale images. Color is fundamentally a three-dimensional phenomenon, described by the perceptual attributes of lightness, chroma, and hue. The conversion from color to grayscale is therefore highly lossy, reducing 3-D information to a 1-D representation. In the most common approach, the lightness information is retained and the chroma and hue information is discarded. In Postscript, for example, the conversion assumes NTSC RGB primaries, and is given by:[1]

$$GRAY = 0.30\,R + 0.59\,G + 0.11\,B \tag{1}$$

Assuming that the R, G, and B are signals that are linear in luminance,* Equation (1) is essentially a calculation of luminance (or close approximation thereof), which is uniquely related to lightness. With this approach, colors with small lightness differences but large differences in hue or chroma will be indistiguishable in the grayscale output, even though they were quite distiguishable in the original color input. For example, according to Eqn (1), white, and pure yellow, which are easily distinguished from each other, map to luminances of 1.0 and 0.89 respectively. Two grayscale stimuli with these luminances are much more difficult to discern. This loss can be particularly objectionable if the two colors are spatially adjacent in a pictorial image.

Several alternative techniques have been proposed for addressing the aforementioned problems. One approach is to incorporate the Helmholtz-Kohlrausch (H-K) effect, which states that the perceived lightness of a stimulus increases as its chroma increases. In other words, a highly chromatic color at a given measured L* appears lighter than a less chromatic color at the same measured L*. It is possible to exploit this effect in converting color to grayscale. Fairchild and Pirrotta[2] defined a modified lightness measure, L**, to model the H-K effect, according to Eqn. (2):

$$L^{**} = L^{*} + 0.143\,C^{*} \quad . \tag{2}$$

The authors also propose a version of the model that varies with hue. Assuming a color can be represented in terms of its lightness, chroma, and hue, Eqn (2) can be used to compute "gray" as the L** value of the original input color. It can be readily seen from Eqn (2) that the distinction among certain combinations of color (e.g. blue and black) would be enhanced, while the distinction among other combinations (e.g. white and yellow) would be reduced. Thus incorporation of the H-K effect would be expected to improve the perceived lightness match between a color image and a grayscale reproduction, but offers no particular systematic improvement for the problem at hand, which is to preserve distinction among spatially adjacent colors.

Another class of techniques has been proposed that converts color to grayscale texture, in order to retain information from the original to the grayscale reproduction. A method by Harrington applies a different halftone screen to each component of color in the image, based on the frequency with which each color appears.[3] Another technique by Harrington and Taber map distinct colors to predefined pattern blocks.[4,5] Such techniques have useful applications, particularly in business graphics imagery. However, they do introduce texture in the image, which may be objectionable. For example, introduction of texture into smooth fleshtone regions in a portrait is likely to produce an undesirable appearance.

A recent approach by Bala and Braun[6] converts input colors to gray levels that are spaced according to their relative 3-D distances in color space. This technique is only applicable for business graphics containing a small number of colors, and as such does not apply to pictorial images, which is the focus of this paper.

Another technique[7] examines the relationship between a foreground and background color, and uses this to determine a correction to the corresponding foreground and background luminance levels. Again, this technique appears to be applicable mainly for graphics images comprising a small number of colors.

A Spatial Color-To-Grayscale Conversion Technique

At a previous Color Imaging Conference[8] we presented a gamut-mapping algorithm that attempts to preserve local spatial luminance relationships. In that algorithm, local luminance edge information lost from the gamut-mapping step was recovered by computing the difference between the original and gamut-mapped luminance image, processing through a spatial high-pass filter, and feeding this back additively to the gamut-mapped luminance.[8,9] The method described in this paper for converting color to grayscale is inspired by similar thinking. (After all, the color-to-grayscale transformation may be thought of as an extreme case of gamut-mapping.) We adopt a similar idea of feeding back information lost in the color to grayscale mapping into the system; however the feedback pays attention to high-frequency errors in chrominance rather than luminance.

Figure 1 is a block diagram of the new method. The input color image is assumed to be in a luminance-chrominance representation such as L*a*b*. If the image is in an RGB space a simple conversion produces such a luminance-chrominance representation. Without loss of generality, the rest of the description uses the L*a*b* representation.

The standard approach is described by the top path in Fig. 1 (i.e. $L_{out} = L_{in}$). The proposed technique introduces an additive correction term to L^*_{in} that accounts for spatial chrominance variations. The first step of this algorithm computes high-pass filtered versions of all three channels, denoted L_{hp}. a_{hp}, b_{hp}. The high-pass content from the two chrominance channels is then combined into a single signal c_{hp} that represents high-frequency chrominance information. An obvious candidate for the combination is the Euclidean metric:

$$c_{hp} = \sqrt{a_{hp}^{\ 2} + b_{hp}^{\ 2}} \tag{3}$$

Another alternative, used in our implementation, is the slightly computationally simpler 1-norm metric:

$$c_{hp} = |a_{hp}| + |b_{hp}| \tag{4}$$

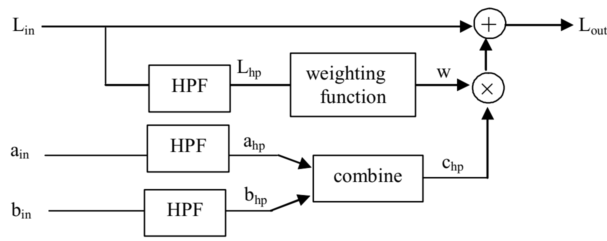

Figure 1. Block diagram of proposed algorithm. "HPF" denotes high-pass filter.

The signal c_{hp} is then multiplied by a modulating factor w that is a function of the high-pass signal L_{hp}. This factor is chosen with two criteria in mind:

1) The polarity of the chrominance correction has to be matched with the polarity of the luminance variation. To understand this, consider the scenario shown in Fig. 2 where an edge in an image is made up of variations in both luminance (a) and chrominance (c), but in opposite polarities. (This is a very common occurrence, e.g. at an edge between white and a colorful region.) The high-pass chrominance signal c_{hp} is shown in Fig 2 (e). Adding this directly to the luminance edge in Fig. 2(a) would not produce the desired edge enhancement. To avoid this, we compute the high-pass luminance signal L_{hp} shown in Fig 2(b), and assign the polarity of this signal (i.e. sign(L_{hp}) to that of the weight w, as shown in Fig. 2(f). In essence, the chrominance difference is always used only to increase the local luminance edge strength, never to weaken it.
2) The second criterion is to introduce the chrominance variation only in regions where the luminance variation is not sufficient to distinguish the local image variation. In order to achieve this, we reduce the amount of chrominance feedback when the luminance variation is large. This is a conservative strategy that avoids excessive edge enhancement in regions that already exhibit sufficient detail in the luminance channel.

To satisfy the aforementioned two criteria, the chosen weighting function is given by:

$$w = sign(L_{hp}) \times f(|L_{hp}|) \ , \tag{5}$$

where the sign( ) term addresses the first criterion, and the function f( ) addresses the second criterion. An example of f($|L_{hp}|$) is shown in Figure 3.

Parameters K, B1, and B2 control the amount of luminance edge enhancement as a function of the strength of the luminance edge in the original color image. Clearly, large values of K, B1, B2 result in more aggressive enhancement of high frequency chrominance information. As these parameters approach zero, the spatial technique converges to the standard technique. Note that this is only one simple example, and that many functions can be chosen that fulfill the desired characteristics.

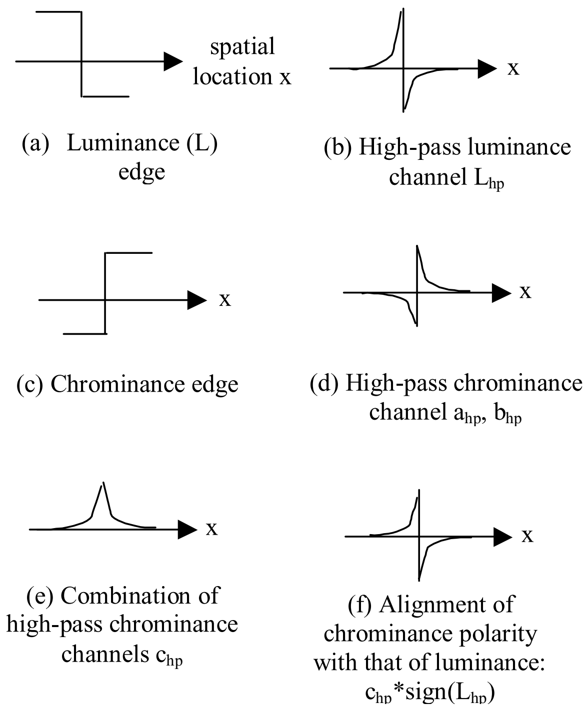

Figure 2. Interaction between luminance and chrominance edges

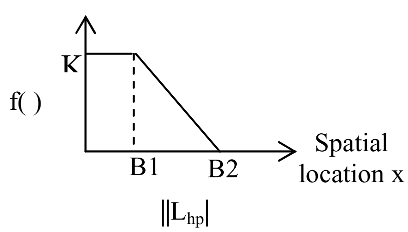

Figure 3. Example of function f( ) used to compute chrominance weighting factor

The output signal L_{out} in Fig. 1 may subsequently be passed through an L* compression function that maps the dynamic range of the input to that of the printer. The lightness values must finally be converted to device-dependent signals before sending to the printer; the latter are usually black (K) colorant amounts.

Experimental Simulations

Three color images were used in a preference experiment to evaluate the qualitative performance of the proposed algorithm. Two of these were pictorials, and one was an image of black text on white background, containing yellow highlights. Color prints were made by rendering the original sRGB color images on a color-characterized Xerox Phaser 7700 CMYK laser printer. Two grayscale reproductions were generated for each color original: the standard mapping that retains the L* component; and the proposed algorithm that generates a spatially corrected L* image. In both cases, the resulting L* image was converted to black (K) toner amount using a 1-D calibration function for the Phaser 7700, then printed. This simulates a true grayscale device.

The parameters of the weighting function f( ) were chosen as K=1, B1 = 15, B2 = 40 and were arrived at heuristically with some initial experimentation. The high-pass filter was chosen as the result of subtracting an NxN constant-weight average from the original image. That is, at pixel *i*, we have:

$$L_{hp}(i) = L(i) - \frac{1}{N^2}\sum_{j\in S} L(j) \quad , \qquad (6)$$

where *S* is an N×N neighborhood around pixel *i*. Analogous expressions hold for the chrominance channels a_{hp} and b_{hp}. This form of high-pass filter was based on successful results from previous gamut mapping work [8]. Given that the images were to be rendered at 600 dpi, a filter size of N=15 was used for the pictorial scenes, while for text and graphics content, a smaller size of N=5 was found to be more effective. These observations are similar to the findings in Ref. [8]. The same filter size was used for both luminance and chrominance channels. (Since the filtered luminance signal is used only to modulate the amplitude and polarity of the filtered chrominance signal, we had no particular motivation to use different filter sizes for luminance vs. chrominance.)

A visual paired-comparison experiment was conducted wherein each of the three color images were presented to observers, along with the standard and spatial grayscale reproductions. All images were presented as hardcopy prints, and viewed under nominal office lighting. For each image, observers were asked to select the grayscale reproduction that they felt best captured the appearance and intent of the original color image. There were six observers, which resulted in a total of 6×3 = 18 paired comparisons.

Results

Out of the 18 comparisons, 16 decisions were in favor of the spatial technique, while the remaining 2 were in favor of the standard approach. Both the latter decisions were by the same observer, who felt that the spatial technique resulted in somewhat excessive edge enhancement in certain image regions.

Figures 4 and 5 show a simulation of two of the images used in the experiment. In Fig. 4, the second and third caps are of red and green hues respectively. The colors on either side of the edge between the caps happen to be of similar luminance. Therefore, while the edge is readily distinguished in the original color image, much of this distinction is lost in the standard luminance mapping in Fig. 4(a). The spatial technique in Fig. 4(b) restores this distinction in the local vicinity of the edge. A similar observation can be made by examining the edge between the third (green) and fourth (pink) cap, and the edge between the fourth cap and the sky. On the other hand, the shadows are a result of mostly luminance, rather than chrominance variations; thus the proposed technique does not introduce any significant edge enhancement.

4(a)

4(b)

Figure 4. (a) standard and (b) spatial color-to-grayscale mappings

Figure 5 is a demonstration of how the proposed technique can be used effectively for computer-generated graphics imagery. The original color image comprises text with yellow highlights, which are essentially lost in the standard grayscale rendition. This is because yellow has a very high L* value, very similar to that of paper. This situation would actually worsen if the Helmholtz-Kohlrausch effect in Eqn (2) is included. Darkening the yellow areas, or rendering the yellow with a texture will likely reduce the legibility of the text. The proposed technique instead uses the strong yellow-white chrominance edges to produce outlines around the highlighted text. This serves to highlight the important regions in a non-obtrusive manner.

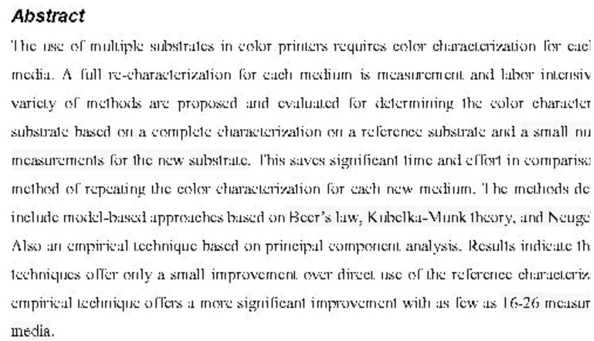

Abstract

The use of multiple substrates in color printers requires color characterization for eac
media. A full re-characterization for each medium is measurement and labor intensiv
variety of methods are proposed and evaluated for determining the color character
substrate based on a complete characterization on a reference substrate and a small nu
measurements for the new substrate. This saves significant time and effort in comparis
method of repeating the color characterization for each new medium. The methods de
include model-based approaches based on Beer's law, Kubelka-Munk theory, and Neuge
Also an empirical technique based on principal component analysis. Results indicate th
techniques offer only a small improvement over direct use of the reference characteriz
empirical technique offers a more significant improvement with as few as 16-26 measur
media.

5(a)

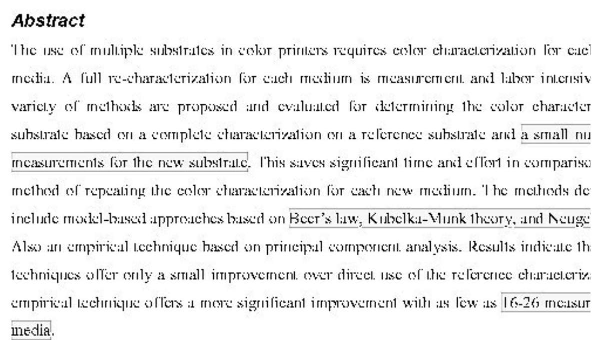

Abstract

The use of multiple substrates in color printers requires color characterization for eac
media. A full re-characterization for each medium is measurement and labor intensiv
variety of methods are proposed and evaluated for determining the color character
substrate based on a complete characterization on a reference substrate and a small nu
measurements for the new substrate. This saves significant time and effort in comparis
method of repeating the color characterization for each new medium. The methods de
include model-based approaches based on Beer's law, Kubelka-Munk theory, and Neuge
Also an empirical technique based on principal component analysis. Results indicate th
techniques offer only a small improvement over direct use of the reference characteriz
empirical technique offers a more significant improvement with as few as 16-26 measur
media.

5(b)

Figure 5. (a) standard and (b) spatial color-to-grayscale mappings

Conclusions

A novel spatial technique has been proposed for converting color images to grayscale. The main feature that distinguishes this algorithm from the standard L* preservation or any other pixel-wise technique is that the same color in the input image can map to different grayscale values in the output depending on the spatial surround. In general, the proposed algorithm helps restore chrominance detail that is otherwise lost in the standard grayscale (L*) image. Experiments strongly validate the improved visual results obtained from the technique. Care should be taken to avoid excessive enhancement. We have attempted to do this via the function in Fig. 3, as well as by maintaining unity filter gain (i.e. K=1). Also, it should be noted that the new technique introduces a computational overhead due to the spatial filtering in Fig. 1. However this can be kept to a modest level with the simple constant-weight filters described above.

Future work includes a more systematic optimization of the algorithm parameters (i.e. choice of luminance-chrominance space, filter size, K, B1, B2), and more extensive experimentation, including the effects of other image processing functions typically encountered, such as calibration, compression, resolution conversion, etc.

Acknowledgements

The authors wish to thank the reviewers for their valuable comments and suggestions.

References

* It is common to apply Eqn (1) to nonlinear (i.e. gamma-corrected) RGB signals. This approach produces a result that is different from luminance. However, the problem of lost discriminability remains, although the extent of the severity will be different in different parts of color space.

1. *PostScript Language Reference Manual*, 2nd ed., p. 304, Chapter 6: Rendering, Addison-Wesley, Reading, 1990.
2. M. D. Fairchild and E. Pirrotta, "Predicting the Lightness of Chromatic Object Colors using CIELAB," *Color Res. & Appl.* **16**, pp. 385-393, 1991.
3. S. J. Harrington, "Mapping of color images to black-and-white textured images," US Patent 5,153,576.
4. Y. Bai, S.J. Harrington, J. Taber, "Improved algorithmic mapping of color to texture," *Proc. SPIE Vol. 4300, Color Imaging: Device-Independent Color, Color Hardcopy, and Graphic Arts VI*, Reiner Eschbach, Gabriel G. Marcu, Eds., 4300-55, pp. 444-451, 2000.
5. S. J. Harrington and J. Taber, "Printing black and white reproducible color documents," U.S. Patent 5,701,401.
6. R. Bala, K. Braun, "Color-to-grayscale conversion to maintain discriminability", *Proc. SPIE Vol. 5293, Color Imaging IX: Processing, Hardcopy, & Applications*, Reiner Eschbach, Gabriel G. Marcu, Eds., pp. 196-202, 2004.
7. R. Pleva, M. Randall, "Color-to-monochrome conversion", US Patent 4,977,398.
8. R. Balasubramanian, R. deQueiroz, R. Eschbach, W. Wu, "Gamut mapping to preserve spatial luminance variations", IS&T and SID's 8th Color Imaging Conference, Nov 7-10, 2000.

9 R. Bala, R. DeQueiroz, R. Eschbach, W. Wu, "Gamut mapping to preserve spatial luminance variations", *Journ. Imaging Science & Technology*, Vol. 45 No. 5, pp. 436-443, Sept/Oct 2001.

Biography

Raja Bala received the Ph.D. degree from Purdue University in 1992 in Electrical Engineering. Since then, he has been with the Xerox Research and Technology Division in Webster, New York. He is a Principal Scientist, and leads a project on color science, device characterization and color management. Raja holds 30 patents and 40 publications in the field of color imaging. He is a member of IS&T.