Gamut Mapping to Preserve Spatial Luminance Variations

Raja Balasubramanian, Ricardo deQueiroz, Reiner Eschbach, and Wencheng Wu
Xerox Research & Technology
Webster, New York

Abstract

A spatial gamut mapping technique is proposed to overcome the shortcomings encountered with standard pointwise gamut mapping algorithms by preserving spatially local luminance variations in the original image. It does so by first processing the image through a standard pointwise gamut mapping algorithm. The difference between the original image luminance Y and gamut mapped image luminance Y' is calculated. A spatial filter is then applied to this difference signal, and added back to the gamut mapped signal Y'. The filtering operation can result in colors near the gamut boundary being placed outside the gamut, hence a second gamut mapping step is required to move these pixels back into the gamut. Finally, the in-gamut pixels are processed through a color correction function for the output device, and rendered to that device. Psychophysical experiments validate the superior performance of the proposed algorithm, which reduces many of the artifacts arising from standard pointwise techniques.

Introduction

Gamut mapping is an important problem in color management, and has been one of the most active areas of research in the Color Imaging Conference series. The optimal gamut mapping strategy for a given application depends on input and output gamuts, image content, user intent and preference. The design of the optimal technique thus involves a suitable trade-off among image attributes such as contrast, luminance detail, vividness, and smoothness. A plethora of gamut mapping algorithms has been proposed in the literature, optimized for different applications, and with different trade-offs. Excellent overviews and references to work in this area can be found in [1] and [2].

One might classify gamut mapping algorithms into three basic categories. The first category comprises device dependent algorithms, where-in the gamut mapping is a function of the input (usually CRT) and output (usually printer) gamuts. These algorithms are independent of input image content. The majority of well known gamut mapping algorithms fall in this category.[1,3-5]

The second category consists of image dependent algorithms, where-in the gamut mapping is a function of the input image statistics, and the output device gamut. These algorithms are generally expected to perform better since they adapt to image content.[2,6,7]


Figure 1. Mapping of black text and blue background from CRT gamut (dashed line) to print gamut (solid).

In both these categories, the gamut mapping is a pointwise operation from an input point to an output point in an appropriate (usually perceptual) 3D color space. One of the fundamental problems with such pointwise operations is that they do not take important spatial neighborhood effects into account. For example, consider an image composed on the CRT, with black text against a blue background. The text is easily distinguished against the background. However, when this is mapped to a printer's gamut with an algorithm that maps out-of-gamut colors to the nearest surface color, the CRT blue maps to a much darker blue in the printer's gamut. On the other hand, the CRT black maps to a lighter printer black. This is illustrated in Fig 1, where the dotted and solid gamuts represent the CRT and printer respectively, and the nearest point mapping is labeled GM1. As a result of this gamut mapping, much of the luminance distinction is lost between text and background, and the legibility of the text is diminished. A comparison of luminance profiles of the input and resulting printed images is shown in Fig. 2. Note that such a gamut mapping function is considered optimal when rendering large areas of black or blue in isolation. The problem arises when they are juxtaposed. One can alleviate this problem by adopting a different pointwise gamut mapping algorithm that preserves luminance (labeled GM2 in Fig 1). Now the visibility of the text will

greatly improve, but luminance preservation usually comes at the cost of significant loss in chroma, and this will likely be unacceptable in a different image. Hence, all pointwise algorithms are heavily constrained by such trade-offs, making it difficult to develop a common algorithm that achieves high quality for a large variety of images and gamuts.
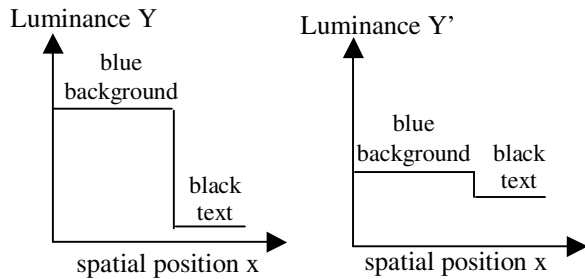

Figure 2. spatial luminance profile of the gamut mapping shown in Fig 1.

The third category of gamut mapping algorithms, which is the focus of this paper, consists of algorithms that take into account spatial characteristics in addition to color characteristics of the image. We believe not many algorithms exist in this class. With such algorithms, two pixels of the same color in an image can map to different colors in the output image, depending on the spatial characteristics in the neighborhood of the pixels. A few researchers have proposed techniques in this category. Meyer and Barth[8] used homomorphic filtering to separate low and high spatial frequency channels, and then to apply global dynamic range compression only to the low frequency channel. A potential problem with such approaches that separate spatial from color transformations is that they are susceptible to noise amplification. Kasson [9] proposed a blending of two gamut mapping algorithms, one preserving luminance and one preserving chrominance. The blending is a function of distance from gamut, and spatial frequency, with luminance being preserved at high frequencies, and chrominance preserved at low frequencies. McCann[10] used the principles of Retinex theory to develop an iterative gamut mapping which attempts to preserve ratios of colors at adjacent pixels.

Gamut Mapping with Spatial Feedback

In this paper, we describe a spatially dependent gamut mapping algorithm based on the principle that it is more important to preserve luminance at high spatial frequencies, while it is generally desirable to preserve chrominance at low spatial frequencies. Unlike some of the existing approaches mentioned above, in which spatial processing and pointwise gamut mapping occur in separate steps, our proposed method tightly couples the spatial and color transformations in a corrective feedback mechanism, resulting in a robust framework for gamut mapping.

In the following discussion, the term "luminance" is used generically to encompass the strict definitions of luminance (i.e. the Y component in XYZ) and lightness (i.e. the L* component in CIELAB). The chrominance components C_1 and C_2 are likewise a generic representation of opponent color signals. Gamut mapping operations take place in such a device independent luminance-chrominance space.

A block diagram of the proposed algorithm is shown in Fig 3. Let us define G_1 as a pointwise gamut clipping algorithm that emphasizes preservation of chroma over luminance. Let G_2 be another pointwise gamut clipping algorithm that emphasizes preservation of luminance over chroma. First G_1 is applied to the input colors, and an error image ΔY is computed between the luminances of the input and gamut mapped signals. A high-pass filter F is applied to the error image, resulting in image $\Delta Y'$. This image, which comprises only the high frequency errors introduced by gamut mapping, is then added back to the gamut mapped signal Y'. Note that this feedback step may move the pixel color (Y'' C_1' C_2') back out of the gamut, and hence, a second gamut mapping operation G_2 is applied to limit all colors to the gamut. Note that the algorithm preserves the characteristics of G_1 at low spatial frequencies, while using G_2 to preserve high frequency luminance variations in the original image that G_1 may have lost or altered. Hence the strengths of both algorithms are exploited in the appropriate spatial frequency bands, and the trade-offs that one must face with pointwise algorithms are now significantly mitigated. All the operations up to this point constitute the overall spatial gamut mapping algorithm, performed in a device independent luminance-chrominance space. The final step is to convert device independent color to device dependent color (i.e. CMYK) via a printer color correction transform. To reduce the overall computational complexity of the algorithm, G_1 can be implemented in a 3-D lookup table, and G_2 can be concatenated with the printer color correction transform from CIE color to CMYK.

The design of G_1 and G_2, and the spatial filter F can depend on many factors, including global and local image characteristics, device characteristics, rendering intent and preference. We will describe a simple initial implementation in this paper, recognizing that more research will be needed to further optimize the algorithm parameters. For this paper, the gamut mapping G_1 was chosen to map out-of-gamut colors to the nearest surface point of the same hue. This mapping generally favors preservation of chroma over luminance. For G_2, the cusp algorithm was chosen, where out-of-gamut colors are mapped to the surface in a direction towards a neutral point whose luminance is that of the cusp color.[1] (The cusp is defined as the point of maximum chroma in a given hue slice.) This algorithm tends to emphasize luminance over chroma preservation, especially for points close to the gamut surface. Both G_1 and G_2 leave in-gamut colors unaltered.

*Figure 3. block diagram of proposed spatial gamut mapping algorithm*

Figure 4. spatial luminance profile of black text and blue background of Figures 1 and 2 at the various stages in the proposed algorithm
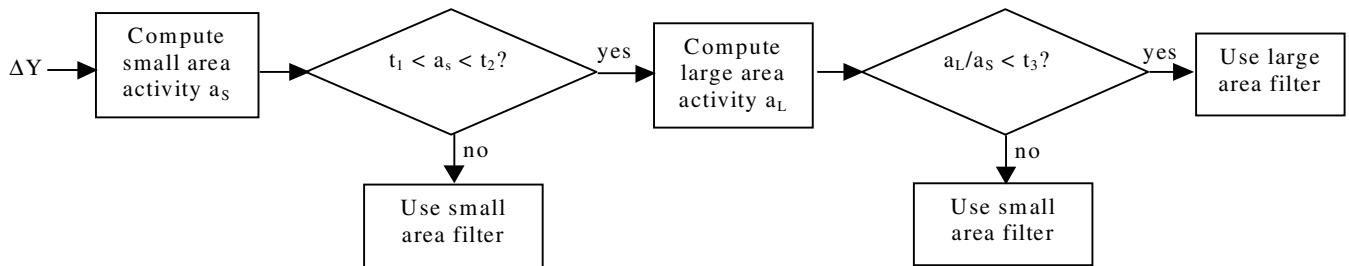

Figure 5. Flow diagram of adaptive filter size selection algorithm

While the chosen G_1 yields high-chroma reproductions, it is susceptible to the "lightning rod effect", where-in several image colors map to one point, especially near black and at the gamut cusp. In the proposed technique, if these image colors are from a high spatial frequency region, the filtered feedback will re-distribute their luminance values, and G_2 will retain luminance distinction, thus eliminating the problem.

Figure 4 demonstrates what happens at various points in the algorithm for the example of blue text on black background (Fig. 1, 2). Adding the filtered error $\Delta Y'$ to the gamut mapped luminance Y' yields the signal Y'', which retains the characteristics of the original input image Y near the edge while retaining the characteristics of the gamut mapped image Y' in smooth regions (see Fig 2). This signal, in combination with the chrominance signals

C_1' and C_2' must be remapped to the gamut surface with the transform G_2, to yield a luminance profile Y''', which may be somewhat different from Y''. In this example, Y'' contains luminances that are below the minimum luminance achievable by the printer, hence these values get clipped. However, even with this limitation, the algorithm restores the edge information that was diminished in the pointwise algorithm (shown as a dashed line in Fig 4). The extent and spatial footprint of the enhancement is dependent directly on the characteristics of the high-pass filter F. In this paper, we have chosen a simple linear filter whose operation at pixel i is given by:

$$\Delta Y'_i = \Delta Y_i - \frac{1}{N^2} \sum_{j \in S} \Delta Y_j \qquad (1)$$

where S is an NxN neighborhood around pixel *i*. With these characteristics the overall gamut mapping, with spatially filtered feedback, will approximately reproduce the variations in Y at high spatial frequencies, while reducing to the pointwise mapping G_1 at low spatial frequencies. As might be expected, the optimal filter size N depends on the type of image. For images with soft or noisy edges, e.g. in scanned pictorials, a large filter size was required for noticeable improvement. On the other hand, for images that have strong edges and low noise, e.g. computer generated business graphics, a large filter size gave rise to distinct halo effects around edges; hence a smaller filter size had to be used. From experimentation on a Xerox DocuColor12 xerographic CMYK printer with 600 dpi resolution, a filter size of N=15 was adopted for scanned pictorials, and N=3 was used for computer generated graphics. Clearly, the optimal N is resolution dependent.

Adaptive Filtering Algorithm

If the image type is known *a priori*, the algorithm can use this to switch between the small and large filter sizes. Since this information is not always known, a simple method was developed to adaptively determine filter size based on local spatial characteristics of the difference image ΔY. The adaptive algorithm was based on the observation that for graphics images, the local variance, or activity in the ΔY image tends to be either very small (i.e. flat areas) or very large (i.e. edges); while the local activity in scanned images tends to mostly take on intermediate values. To minimize the computational overhead, the absolute value of the high pass filter output, |ΔY'| was itself used as an approximate measure of local image activity. Figure 5 is a flow diagram of the proposed scheme. First, the activity measure a_S is computed at the small filter size (N=5), and checked to see if it lies within a certain range $[t_1, t_2]$. Activity measures outside this range are either very small or large, hence a small filter size is used, suitable for graphics images. If a_S lies within the given range, this suggests the characteristics of a scanned pictorial. However, an intermediate activity value can also indicate that the pixel is near, but not on a strong edge in a synthetic graphics image. To differentiate this signature from a noisy scan, a second decision step is used. The activity a_L is derived for a large filter size (N=15), and the ratio of the two activities a_L/a_S is calculated. This ratio will take on large values (>>1) if the pixel is near a strong edge (i.e. graphics), and will approach 1 if the image variance characteristics are invariant across the two filter sizes (i.e. a noisy scan). Comparing this ratio against a threshold t_3 provides the final decision on filter size. The various thresholds were obtained by trial and error to minimize the rate of misclassification.

Experimental Results

Psychophysical Evaluation of Spatial Gamut Mapping

A visual experiment was conducted to compare standard pointwise gamut mapping techniques with the proposed spatial mapping method for scanned pictorial imagery. The standard techniques were i) clipping to the nearest point on the gamut surface while preserving hue; and ii) nonlinear L* compression using the inverse-gamma-inverse (IGI) technique[4] followed by cusp clipping. These algorithms have been reported as successful pointwise techniques in previous experiments.[4] Two versions of the spatial gamut mapping method were tested: one without and one with IGI L* compression. The latter was applied as the very first step, i.e. before the filtered feedback in Fig 3. The operations G_1, G_2, and F were implemented as described in the previous section. CIELAB was used as the luminance-chrominance space. Since the images were known to be pictorials, the filter size was fixed to be N=15. We adopt the following symbols for the four algorithms: NP, IGI_CUSP, SGM, IGI_SGM.

Eighteen observers participated in two pairwise-comparison experiments: i) preference, where the observer was asked to select the most preferred image from a pair; and ii) reproduction, where the observer was asked to select the best reproduction from a pair with respect to an original reference image. Both original and gamut-mapped images were printed on a Xerox DocuColor12 printer. The gamut-mapped images were restricted to a smaller gamut of an inkjet printer. Five images were used, whose grayscale versions are shown in Fig 6. All images were viewed in a light booth with D50 illuminant. The printer was also calibrated for matching under D50.

*Figure 6 Grayscale versions of images used in psychophysical experiment*

Figure 7 Preference scales from visual experiment for (a) preference and (b) reproduction experiment. The algorithms are: 1. NP. 2. IGI_CUSP 3. SGM 4. IGI_SGM

Table 1 Standard deviation of preference scale values computed from 10 subsets of the psychophysical data

Algorithm	1	2	3	4
Preference	0.24	0.34	0.30	0.24
Reproduction	0.46	0.41	0.29	0.35

A standard Thurstone (Case V) analysis[11] was carried out on the pairwise-comparison data, and preference scales were generated for the two experiments. The scales are shown in Figure 7. The error bars were computed according to [12].

To gain insight on the variance in algorithm performance as a function of image content and observers, the following additional analysis was done. The observers were divided into expert and non-expert categories depending on their experience in color imaging. The collected experimental data was then divided into 10 subsets: 5 images x 2 observer categories. 10 preference scales were generated, one for each subset, and the mean and standard deviation of the 10 preference scores for each algorithm was computed. As expected, the mean scores closely tracked the scores shown in Fig 7 for the overall experiments. The standard deviations of the scores are shown in Table 1. Note that these values should be interpreted relative to each other, and not be directly related to the graphs in Fig. 7.

From Fig. 7, it is seen that IGI_SGM performed best in the preference experiment, and SGM performed best in the reproduction experiment. Generally, the spatial algorithms performed better than their pointwise counterparts, because they effectively retained detail and edge information in shadows and high-chroma regions that is often lost with standard techniques. L* compression resulted in improved performance in the preference experiments, presumably due to an increase in perceived overall lightness, colorfulness, and contrast of the images. However, this was not the case in the reproduction experiment, presumably because the color changes just mentioned would result in a less accurate match to the original image. Finally, from Table 1, we see that the spatial mapping algorithms that performed best also exhibited the smallest variance as a function of images and observers. This increases our confidence in these algorithms.

Adaptive Filtering Algorithm

To test the efficacy of the adaptive filtering algorithm, we processed the image shown in Fig 8(a), containing both pictorial and graphics content, and examined the filter size selected by the algorithm at each pixel. The two filter sizes were 15x15 and 3x3. Thresholds chosen by trial-and-error were t1=0.1, t2=15, t3=25. Fig 8b is a filter classification map, where-in black and white denote selection of large and small filters, respectively. The areas of the pictorial which benefit most from a large filter, namely the textured sweater and detail in the hair, are indeed classified in the "large filter" region. Note that portions of the girl's face are actually treated with a small filter. This is because these regions are actually within gamut, and the error image ΔY is flat (i.e. $\Delta Y = \Delta Y' = 0$). In this case, no filtering is needed; however for convenience, these pixels are grouped with the "small filter" category. Most of the graphics image is correctly treated with the small filter, except for certain edge regions. However, these misclassified regions are no larger than the support of the small filter (i.e. 3x3). Hence artifacts arising from erroneously using a large filter are highly localized, and generally not objectionable.
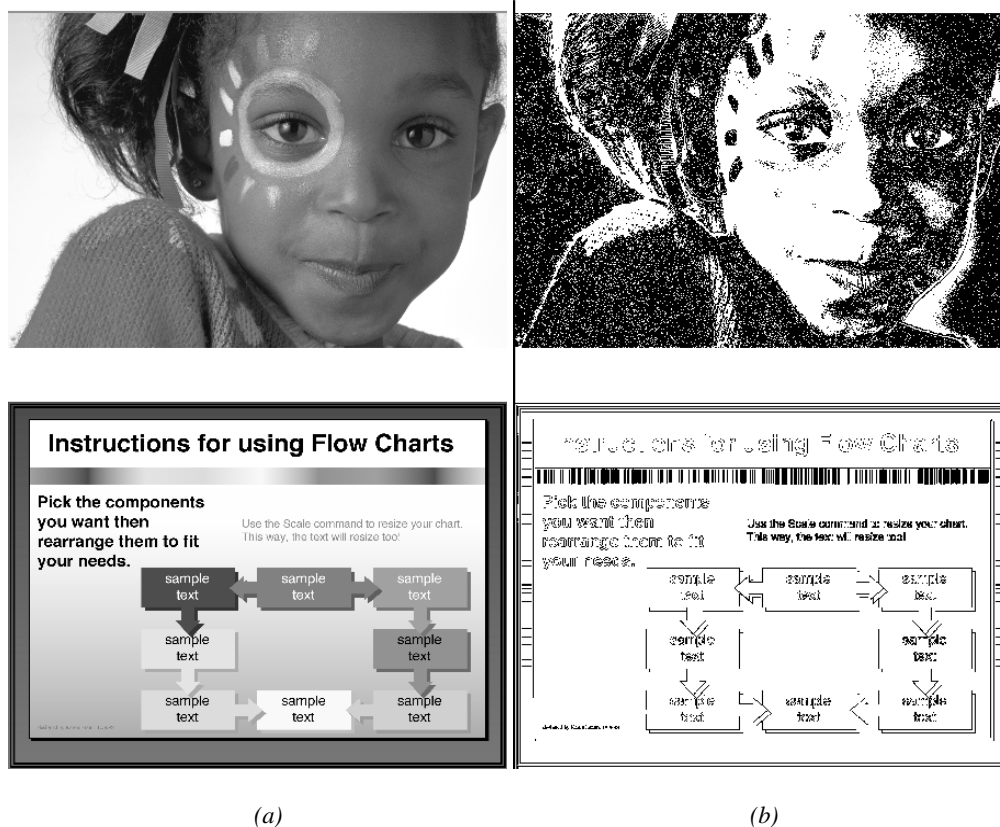
(a)   *(b)*

Figure 8(a) grayscale version of composite image containing pictorial and graphics content; (b) filter size classification map from adaptive filtering algorithm, wherein black and white correspond to large and small filter sizes respectively.

Summary

We have presented a unique gamut mapping algorithm that takes into account spatial characteristics of the image. This feature eliminates some of the compromises necessitated by standard pointwise algorithms. By closely coupling the spatial and color transformations in a corrective feedback mechanism, our approach does not suffer from the noise amplification problems that can arise when the two transformations are applied separably. Psychophysical experiments show that the new approach outperforms standard pointwise gamut mapping methods for pictorial images. We believe that with some extensions, the algorithm will be equally effective for business graphics images. The proposed adaptive filtering scheme is a promising step towards a more robust framework for spatial gamut mapping. Finally, while our method is likely to achieve results that are qualitatively similar to that obtained by Kasson and McCann, we believe our algorithm requires a simpler implementation and fewer computations than either of these approaches.

Acknowledgements

References

1. J. Morovic, *To Develop a Universal Gamut Mapping Algorithm*, Ph. D. Thesis, University of Derby, 1998.
2. G. Braun, *A paradigm for color gamut mapping of pictorial images*, Ph. D. Thesis, Rochester Institute of Technology, 1999.
3. N. Katoh, M. Ito, S. Ohno, "Three-dimensional gamut mapping using various color difference formulae and color spaces", *Journ. Electronic Imaging*, **8(4)**, pp. 365-379, 1999.
4. K. Braun, R. Balasubramanian, R. Eschbach, "Development and evaluation of six gamut mapping algorithms for pictorial images", *Proc. IS&T/SID's 7th Color imaging Conference*, pp. 144-148, 1999.
5. K. Braun, R. Balasubramanian, S. J. Harrington, "Gamut mapping techniques for business graphics", *Proc. IS&T/SID's 7th Color imaging Conference*, pp. 149-154, 1999.

6. C. Hung-Shing, M. Omamiuda, H. Kotera, "Adaptive gamut mapping method based on image-to-device", *Proc. IS&Ts NIP 15, 1999 International Conference on Digital Printing Technologies*, pp. 346-349, 1999.

7. G. J. Braun, M. D. Fairchild, "General-purpose gamut mapping algorithms: Evaluation of contrast-preserving rescaling functions for color gamut mapping", *Proc. IS&T/SID's 7th Color imaging Conference*, pp. 167-172, 1999.

8. J. Meyer, B. Barth, "Color gamut mapping for hard copy", *SID Digest 89*, 86-89, 1989.

9. J. M. Kasson, "Color image gamut-mapping system with chroma enhancement at human-insensitive spatial frequencies", US Patent 5450216, 1995.

10. J. McCann, "Lessons learned from Mondrians applied to real images and color gamuts", ", *Proc. IS&T/SID's 7th Color imaging Conference*, pp. 1-8, 1999.

11. L. L. Thurstone, "A law of comparative judgment", *Psych Rev.*, **34**, pp. 273-286, 1927.

12. K. M. Braun, M.D. Fairchild, P. J. Alessi, "Viewing environments for cross-media image comparison", *Color Research and Application*, **21(1)**, pp. 6-17, 1996.