

 Open access • Journal Article • DOI:10.1111/J.1467-8659.2005.00867.X

Re-coloring images for gamuts of lower dimension — [Source link](#)

Robert Geist, Karl Rasche

Institutions: Clemson University

Published on: 01 Jan 2005 - Computer Graphics Forum (Blackwell Publishing, Inc)

Topics: Grayscale

Related papers:

- [Color2Gray: salience-preserving color removal](#)
- [Decolorize: Fast, contrast enhancing, color to grayscale conversion](#)
- [Detail preserving reproduction of color images for monochromats and dichromats](#)
- [Apparent Greyscale: A Simple and Fast Conversion to Perceptually Accurate Images and Video](#)
- [Computerized simulation of color appearance for dichromats](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/re-coloring-images-for-gamuts-of-lower-dimension-i4dvuhtx5u>

Re-coloring Images for Gamuts of Lower Dimension

Karl Rasche, Robert Geist, and James Westall[†]

Clemson University, Clemson, South Carolina USA 29634-0974

Abstract

Color images have a gamut that typically spans three dimensions. Nevertheless, several important applications, such as the creation of grayscale images for printing and the re-coloring of images for color-deficient viewers, require a reduction of gamut dimension. This paper describes a technique for preserving visual detail while reducing gamut dimension. The technique is derived by focusing on the problem of converting color images to grayscale. A straightforward extension is then provided that allows re-coloring images for color-deficient viewers. Care is taken so that the resulting images remain within the available gamut and visual artifacts are not introduced.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision]: Enhancement-Filtering

1. Introduction

“Black and white” televisions and monitors have essentially disappeared from the commercial landscape, but grayscale printing devices abound. In its grant proposal guide [Nat04], the US National Science Foundation warns investigators that “For cost and technical reasons, the Foundation cannot, at this time, reproduce proposals containing color. PIs, therefore, generally should not rely on colorized objects to make their arguments.” When preparing documents for publication in any venue, authors must often convert their color images to grayscale to comply with publisher restrictions. The standard solution to this conversion problem is simple and robust: linearly map pixel luminance values to the available grayscale. Although this procedure produces satisfactory results in many cases, in others it does not. Authors do not always have complete control over image design (e.g. photographs), and so a significant portion of the image information content may appear in the chrominance variations. In Figure 1 we show some examples of images of this type. When the color images (left column) are mapped to grayscale based on luminance (middle column), significant information is lost. An alternative mapping to grayscale (right column) can preserve some of the information.

The purpose of this paper is to suggest a method for con-

verting color images to grayscale in a way that, to the extent possible, preserves image information. We contend that there are two components to information preservation:

- **Preserving contrasts:** Colors that are readily distinguishable in the original image should be represented by gray values that are also readily distinguishable. In terms of a perceptual color space (e.g. CIELAB), we would like the distance between any pair of image colors to be proportional to the distance between their gray values.
- **Maintaining luminance consistency:** Shadows, highlights, and color gradients provide depth cues through luminance variations. Grayscale images in which these cues exhibit luminance reversals, e.g. lightened shadows or darkened highlights, can be disorienting. These effects can be ameliorated if luminance gradients within narrowly defined chrominance bands are maintained during the translation to grayscale.

We incorporate these components into an objective function that allows us to cast the problem in the form of a constrained, multi-variate optimization. We then derive a sequence of linear programming (LP) problems whose solutions converge to the desired optimum. A straightforward extension of the method will allow us to re-color images in a way that preserves information for color-deficient viewers.

[†] Dept. of Comp. Sci., email:{rkarl,rmg,westall}@cs.clemson.edu




Figure 1: Color to Grayscale: (a) Original image, (b) mapping luminance to grayscale, (c) proposed, alternative mapping.

1.1. Related Work

There are several related problems that have been addressed extensively in the literature. The goal of color quantization [Hec82] is to preserve image detail while reducing the number of colors in the image from n to m , where usually $n \gg m$. Color quantization algorithms have two components, palette selection and mapping. Clever palette selection is the major component and critical to the success of these algorithms. Mapping image colors to the selected palette is usually a simple, three-dimensional error diffusion [FS75]. In our problem, the palette is fixed (grayscale), and so the principal benefit of any color quantization algorithm is lost. Simple diffusion mappings to grayscale are almost always inadequate.

Gamut mapping attempts to preserve the appearance of an image when it is displayed on two or more devices that have different available color gamuts. Stone *et al.* [SCB88] describe several principles to be followed to create pleasing gamut maps. Two are directly applicable to producing grayscale images. First, the gray axis of the source gamut should be transformed so that it is aligned with that of the destination gamut. For a color-to-grayscale transformation, this implies that we should map the gray axis of the color image (luminance) to grayscale. This is precisely the standard mapping, upon which we are trying to improve. The second principal is that luminance contrast should be maximized. For grayscale images, this implies that we should normalize the results over the target gray range. This normalization will be an integral part of our method.

An important instance of gamut mapping is the tone mapping problem of displaying high dynamic range images on low dynamic range displays [TR93]. The objective is to preserve an abundance of visual information within the constraints of a limited gamut, and our approach to grayscale conversion shares this goal. Nevertheless, tone mapping is generally concerned with compression of the gamut range, whereas we are interested in compression of gamut dimensionality.

Stollnitz *et al.* [SOS98] examine the problem of printing a color image with an arbitrary number of inks. For the case of a single ink, their solution is also to use luminance, i.e., the standard mapping.

Several methods have been proposed for solving the general problem of reducing an n -dimensional set of data to m dimensions where $m < n$. Principal Component Analysis [Joi02] (PCA) is one such method. A set of n orthogonal vectors is constructed, where each points in the direction of the highest remaining variation in the original data set. The m vectors that capture the most variation form a basis of the lower dimensional space. Choosing the axis of maximum variation in the color image data does provides a high contrast grayscale image. Nevertheless, significant smaller detail usually appears in the second principal component,

and such is lost to this technique. Further, it is difficult to incorporate constraints on luminance consistency.

Multi-dimensional Scaling [CC94], or MDS, is an alternative to factor analysis for detecting meaningful underlying dimensions in a set of multi-dimensional data points. It most often takes the form of minimizing a quadratic function of the set of all distances between pairs of points in the multi-dimensional data. Because the input involves all observed distances, the technique does not often scale well to large data sets. For the special case of reduction to a single dimension (uni-dimensional scaling), Hubert *et al.* [HAM02] provide an overview of several approaches. In general, these do not scale well to large data sets, and the addition of external constraints can be problematic. Nevertheless, as shown in Section 2, our own approach can be characterized as a special case of MDS.

Locally linear embedding [RS00] (LLE) takes a slightly different approach to dimension reduction in that it assumes each data point can be constructed as a linear combination of its neighbors. The choice of neighborhood is crucial for satisfactory results. As is the case for PCA, it is unclear how to incorporate constraints into LLE.

ISOMAP [TdSL00] is a popular technique that is similar to LLE in that it requires a neighborhood about each high dimensional point. Here the neighborhood is used to create a weighted graph that approximates the structure of the data. Traditional MDS techniques are then used to create a lower dimensional configuration from the weighted graph.

In the domains of medical imaging or remote sensing, multiple high-dimensional images from a single source are often obtained, e.g., a CT scan and an MRI scan of the same region of a patient. The goal of image fusion is to create a single image that captures the salient features of the entire original data set and thereby allows visualization of data correlations. There have been several methods proposed for image fusion. They include PCA [SOM*87], linear combinations of the data that maximize contrast objective functions [HB96], analysis using wavelets [MLM95], and neural network schemes with self-organizing maps [Man96]. Luminance consistency is not typically a goal of image fusion methods, and so these methods are not easily constrained.

Socolinsky [Soc00] and Socolinsky and Wolff [SW02] construct grayscale images to match local contrast in color images. They regard contrast as a gradient and then recover grayscale by solving a Poisson equation. They specifically avoid introducing global contrast to avoid artifacts when computing sub-images or sequences of images. As a result, their method has difficulty with certain classes of images, such as the three color example in Figure 2. Here the contrast between adjacent red and gray blocks matches the contrast between adjacent gray and green blocks. A grayscale transformation based solely on local contrast sends red and green to identical gray values. Use of global contrast information will separate the three values.



Figure 2: Why global contrast is necessary: *original color (left), transformation using local contrast (center), and transformation using global contrast (right)*

Gooch *et al.* [GOTG05] describe a technique for converting color to grayscale in which they iteratively adjust the gray value at each pixel to minimize an objective function. To resolve the ordering of gray values arising from original colors of differing hues, they provide a user-controlled parameter.

Photographers often encounter difficulties when trying to record color images on monochromatic film. The traditional solution to this problem is to use a “contrast filter” over the lens [Ada81]. These filters come in a variety of hues, each of which attenuates the complementary hue of the filter. For example, to increase the contrast between clouds and sky, a yellow filter is often used. This attenuates the blue sky, producing light clouds on a dark background. In general, the choice of filter is left to the eye of the artist.

1.2. Color Vision Deficiencies

Deficiencies in color vision arise from differences in pigmentation of optical photoreceptors [Wan95]. Normal vision is characterized by three distinct pigmentations of the photoreceptors (cones), which collectively allow reception of long, medium, and short wavelengths of the visible spectrum. Anomalous trichromatopia is a condition in which the pigment in one cone is not sufficiently distinct from the others. The viewer still has three distinct spectral sensitivities, but the separation is reduced. Dichromatopia is a condition in which the viewer has only two distinct pigments in the cones. For both dichromatopia and anomalous trichromatopia, there are three subclassifications based on which cone has the abnormal pigmentation. Deficiencies in cones sensitive to long, medium, and short wavelengths are referred to as protanopic, deutanopic, and tritanopic, respectively. Protanopic and deutanopic deficiencies, the most common forms of color-deficient vision, are characterized by difficulty distinguishing between red and green tones. Tritanopic deficiencies are associated with confusion between blue and yellow tones. Monochromatism is another form of deficient color vision, but it is quite rare.

Significant work has been done in simulating color deficient vision [MG88, Kon90, BVM97]. From studies of subjects who have normal vision in one eye and color-deficient vision in the other, it has been determined that there are specific wavelengths that are perceived identically by those with normal vision and those with vision deficiencies. For protanopes and deutanopes, these are 575nm (yellow) and 475nm (blue). For tritanopes, these are 660nm (red-orange)

and 485nm (cyan). From knowledge of these invariants, color-deficient vision can be simulated. Colors in RGB space are transformed into an LMS (long, medium, short) color space that is based on cone response. Cone deficiencies are then simulated by modifying the response of the deficient cone. Meyer and Greenberg [MG88] find lines of “confused colors” by intersecting lines of chrominance. Kondo [Kon90] uses a series of linear transforms governed by a parameter specifying the degree of deficiency in each particular component. Brettel *et al.* [BVM97] project colors in the cone response space onto a pair of planes that intersect along a neutral color line and represent the gamut of the deficient viewer.

Given the ability to simulate color-deficient views, it may be possible to re-color images in such a way that confused detail is restored for color-deficient observers. Reinhard *et al.* [RAGS01] describe a strategy for matching color statistics between two images and then transferring the color distribution from one image to another. Although this can be useful for compressing a full color image into a color-deficient gamut, it does not attempt to preserve the perceived distances between contrasting colors in the original image.

Walraven and Alferdinck [WA97] describe a color palette editor that is coupled with simulation of a color deficient viewer. The simulator is used to determine which pairs of colors in the current palette would have perceived distances below a critical threshold in a color-deficient gamut. Such colors are marked for re-consideration in the palette design. The editor can also select a default palette for which all colors meet threshold specifications.

Daltonization [DW] is a procedure for re-coloring an image for viewing by a color-deficient viewer. Here, a user specifies parameters for stretching contrast between red and green hues, as well as parameters for modulating blue and yellow contrast and modulating luminance. Simulation of a deutanopic viewer is provided for evaluating the output image. Parameter selection is not automated, although three “default” parameter settings are offered. The results are highly dependent upon the choices for the parameters.

Ichikawa *et al.* [ITK*03, ITK*04] propose an automated method for re-coloring images that is based on a genetic algorithm search. A small subset of the colors in the image is selected and used to construct an objective function that both maintains color distances and constrains the extent of color remapping. The result is then interpolated across the entire image. In this process, there is no consideration of constraints on luminance consistency or constraints to keep colors within a target gamut. The method also has numerous parameters to adjust in order to control the quality of the output image.

Rasche *et al.* [RGW05] discuss an automatic method for both grayscale conversion and re-coloring images. They use a subset of the image colors to find a linear transform that




Figure 3: Luminance Consistency and Linear Projection Mapping (a) to (b) violates luminance consistency; shadows and highlights appear reversed. Mapping (c) to (d) is based on a linear projection; widely differing colors are mapped to similar gray values.

minimizes their objective function, which is designed to enforce proportional color differences across the remapping. For the grayscale case, the transformation involves projecting the colors onto a single vector. For the re-coloring case, it involves a linear transformation in homogeneous CIELAB coordinates. Because their method is based on linear transformations, it has advantages in speed and simplicity, but it also carries attendant limitations. In particular, it has difficulties with images that exhibit smooth variation along multiple color dimensions, and it does not offer any constraints on luminance consistency. Examples of each are shown in Figure 3. In the re-coloring case, their method does not constrain the transformed colors to the available gamut.

2. Constrained Multidimensional Scaling

We seek a mapping of color to grayscale that preserves contrasts and maintains luminance consistency. The two goals can be expressed in the form of a constrained, multidimensional scaling problem. We first define a quadratic objective function that incorporates contrast preservation. We briefly review the technique, known as “majorization,” that is commonly used to optimize functions of this form. We then add constraints that enforce luminance consistency and observe that an approximation will allow us to express the desired solution in terms of a sequence of linear programming problems.

2.1. Majorization

As noted in Section 1, we would like our grayscale mapping to preserve global contrast. We express this in terms of an objective function by requiring that differences between mapped gray values should remain proportional to differences between the original color values. Specifically, we seek a minimum of

$$s(g_0, g_1, \dots, g_{n-1}) = \sum_{i=0}^{n-2} \sum_{j=i+1}^{n-1} (|g_i - g_j| / \delta_{ij} - K)^2 \quad (1)$$

where g_i is the mapped gray value corresponding to color c_i , K is the target proportionality constant, and $\delta_{ij} = |c_i - c_j|$, the distance between c_i and c_j in CIELAB space.

Equation 1 is referred to in the MDS literature as STRESS [CC94]. STRESS is non-linear. It is generally solved by a method known as “majorization”, in which the objective function is minimized by solving an iterative sequence of unconstrained quadratic problems. The sequence of minima will converge to a local minimum of (1).

To establish the bounds, we first expand the summands as

$$(|g_i - g_j| / \delta_{ij})^2 - 2K|g_i - g_j| / \delta_{ij} + K^2 \quad (2)$$

and observe that the sum of the quadratic terms can be written succinctly

$$\sum_{i=0}^{n-2} \sum_{j=i+1}^{n-1} (|g_i - g_j| / \delta_{ij})^2 = g^T Q g \quad (3)$$

where

$$Q_{ij} = \begin{cases} -\delta_{ij}^{-2} & i \neq j \\ \sum_{k \neq i} \delta_{ik}^{-2} & i = j \end{cases} \quad (4)$$

To bound the linear terms, we observe that for any vector y of length n ,

$$|g_i - g_j| |y_i - y_j| \geq (g_i - g_j)(y_i - y_j) \quad (5)$$

and so if $y_i \neq y_j$,

$$\sum_{i=0}^{n-2} \sum_{j=i+1}^{n-1} (|g_i - g_j| / \delta_{ij}) \geq \sum_{i=0}^{n-2} \sum_{j=i+1}^{n-1} \frac{(g_i - g_j)(y_i - y_j)}{\delta_{ij} |y_i - y_j|} \quad (6)$$

We conclude that

$$\sum_{i=0}^{n-2} \sum_{j=i+1}^{n-1} (|g_i - g_j| / \delta_{ij}) \geq g^T L y \quad (7)$$

where

$$L_{ij} = \begin{cases} -1/(\delta_{ij}|y_i - y_j|) & i \neq j, \quad y_i \neq y_j \\ 0 & i \neq j, \quad y_i = y_j \\ -\sum_{k \neq i} L_{ik} & i = j \end{cases} \quad (8)$$

We now have an upper bound for (1):

$$s(g_0, g_1, \dots, g_{n-1}) \leq g^T Qg - 2Kg^T Ly + ((n-1)n/2)K^2 \quad (9)$$

which is minimized by the solution of the linear system

$$Qg = KLy \quad (10)$$

where Q is symmetric and positive semidefinite. Thus, to minimize (1), we can choose an initial random value for y , and iterate: solve (10) for g , set $y = g$, and repeat until convergence. Note that a solution of (10) with $y = g$ is a minimum of the original function (1).

2.2. Constrained Majorization

Maintaining luminance consistency within narrow chrominance bands will require that we add constraints of the form $g_i \geq g_j$ to our optimization of (1). In terms of majorization, this implies that (10) will no longer (necessarily) have a solution, and we must approximate. We do this by introducing n additional variables, v_i , used to represent per-component error in (10). We then express our goal in terms of the linear programming problem

$$\begin{aligned} & \text{minimize} && \sum_i v_i && (11) \\ & \text{subject to} && g_j - g_k \geq 0, && (j, k) \in D \\ & && -g_j \geq -100 && \text{all } j \\ & && v_i + Q_i g \geq K L_i y, \\ & && v_i - Q_i g \geq -K L_i y \end{aligned}$$

where Q_i and L_i are the i th rows of Q and L , and D is the set of indices of those colors with order-constrained luminance values. Note that the second constraint arises from the maximum value, 100, for any gray level in CIELAB space, and that a constraint, $v_i \geq 0$, is implicit in the final pair of constraints. All gray levels are non-negative.

3. Implementation

Several implementation details must yet be addressed. These include selection of the target proportionality constant, K , luminance constraint selection, and handling images with large numbers of colors.

Proportionality constant: There are at least two reasonable choices for the target proportionality constant, K . The most obvious is $K = 100/c_{max}$, where c_{max} denotes the maximum CIELAB distance between any pair of colors in the original image, and, as noted earlier, 100 is the maximum distance between gray values in CIELAB space. A second choice, which tends to heighten contrast, is to weight gray value differences equally with color differences, i.e., $K = 1$. We provide comparisons in Section 5. In an ideal solution, the set of ratios, $|g_i - g_j|/|c_i - c_j|$, would have mean K and variance 0.

Constraints: To preserve luminance consistency, we must choose constraints so that colors with similar chrominance are mapped to a grayscale order that matches their luminance order. For this, we need a threshold, ϵ , on the difference

$$\Delta_{ij} = \sqrt{(a_i^* - a_j^*)^2 + (b_i^* - b_j^*)^2} \quad (12)$$

in CIELAB color space. If $\Delta_{ij} \leq \epsilon$, the pair of colors are deemed similar and a constraint is added. If the luminance of c_i is less than that of c_j , the constraint, $g_i \leq g_j$, is added. The threshold, ϵ , should be the lowest value that still avoids introduction of the luminance-reversal artifacts such as lightened shadows or darkened highlights. In practice, it appears that values of ϵ between 10 and 20 generally produce visually pleasing images. This step can introduce a large number of redundant constraints, and we eliminate those that are implied by transitivity.

Handling large numbers of colors: Images of interest often have hundreds of thousands of colors. The execution-time performance of MDS does not scale well to such large data sets. To attempt to work around this deficiency, several proposals for “sparse MDS” have been put forth. One method is to choose a subset of the total points to be “landmark points”, which carry increased weight in determining the mapping. In one approach to the use of such landmark points, the system is solved by considering only those pairs of points for which at least one member of the pair is a landmark point. This approach is used by Gansner *et al.* [GKN04] for drawing large graphs. In another, two-pass approach, a dense MDS is carried out only on the landmark points, and then the remaining points are mapped based on the results of the dense MDS. This technique is used by Kruskal and Hart [KH66] and de Silva and Tenenbaum [dST03].

We use color quantization to select 256 colors as landmark points and then employ a two-pass approach. After mapping the landmark points, we then wish to map the remaining points to minimize

$$t(g_0, g_1, \dots, g_{n-1}) = \sum_{i=0}^{n-1} \sum_{j=0}^{255} (|g_i - \tilde{g}_j|/\delta_{ij} - K)^2 \quad (13)$$

where g_i is a gray value to be determined, \tilde{g}_j is the gray value already determined for the j^{th} landmark color, and δ_{ij} is the CIELAB distance from c_i to the j^{th} landmark color. This minimization has a form similar to (1), except that the \tilde{g}_j values are known. Expanding (13) as before, we have a bound:

$$\sum_{i=0}^{n-1} \sum_{j=0}^{255} (|g_i - \tilde{g}_j|/\delta_{ij} - K)^2 \leq g^T Hg - 2Kg^T J + R \quad (14)$$

where

$$H_{ij} = \begin{cases} \sum_{k=0}^{255} \delta_{i,k}^{-2} & i = j \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

	Colors	Solving	Interpolating
Berries	1,415,424	22.4 sec	1260.4 sec
Candy	367,136	34.8 sec	365.0 sec
Flower	338,010	22.5 sec	307.5 sec
Grid	43	0.041 sec	-

Table 1: Timing: Results for the grayscale images shown in Figure 1. In each case except the “Grid” image, the solution was two-phase, with an initial MDS over 256 landmark colors followed by an interpolation step. The “Grid” image was solved for all colors without interpolation.

$$J_i = \sum_{j=0}^{255} \left(\frac{\tilde{g}_j}{K\delta_{ij}^2} + \frac{(y_i - \tilde{g}_j)}{\delta_{ij}|y_i - \tilde{g}_j|} \right) \quad (16)$$

R is a constant, and vector y is to be used, as before, in iteration. Since H is diagonal, the right side of inequality 14 has a minimum value

$$g_i = \frac{KJ_i}{H_{ii}} \quad (17)$$

To maintain luminance consistency, we must use constraints when interpolating the results. We find upper and lower bounds on each interpolated value by thresholding the difference in chrominance between c_i and all the landmark colors. If $\Delta_{ij} \leq \epsilon$, \tilde{g}_j is a candidate for either an upper or lower bound on g_i , depending on the luminance order between c_i and the j^{th} landmark color. The interpolated solution is then

$$g_i = \min \left(\max \left(\frac{KJ_i}{H_{ii}}, g_i^l \right), g_i^u \right) \quad (18)$$

where g_i^l and g_i^u are the lower and upper bounds of g_i .

4. Results

Our method is easily implemented using any of a large collection of linear programming libraries. We used the CLP solver that is part of the COIN-OR collection of optimization software (<http://www.coin-or.org>). Timing results for the images shown in Figure 1 are given in Table 1. For the first three images we used the two-phase approach described in the previous section and quantized the original image to 256 colors. The MDS solution for the quantized set was then interpolated over the original image data using (18). The fourth image (blocks of constant color), contained relatively few colors, and so the mapping was computed directly. In each case the proportionality constant, K, was $100/c_{max}$

5. Evaluation

To evaluate the effectiveness of our grayscale transformation, we performed a small experiment to measure viewer preference. The color images from the first three rows of Figure 1, as well as those from Figure 4, were used. All images were converted to grayscale by three methods, the standard mapping of luminance to gray, our method with $K = 1$,



Figure 4: Additional Images: Images used, along with the first three images in Figure 1, for experimental evaluation.

and our method with $K = 100/c_{max}$. A sequence of monitor pages, each containing a color image and a pair of corresponding grayscale images, was presented to each of 17 study participants. All participants were graduate students with some background in either computer graphics, image processing, or design. An exhaustive paired comparison was used for each of the six images. Each participant was instructed to choose the grayscale image that best represented the visual detail of the color original. We used Thurstone’s Law of Comparative Judgment [Tor67] to arrange the paired preferences for each image in a linear scale. A 95% confidence interval was computed, as discussed by Braun *et al.* [BFA96]. There were no unanimous decisions, and so no special handling of such cases was necessary for the scaling.

To determine the overall scales for the techniques, we would ordinarily average the scale values across all images, but there is a logical partition in our data. The grayscale images produced from the color images of Figure 4 were all visibly similar, but those produced from the color images of Figure 1 were appreciably different from one another. To avoid skewed results, we averaged separately across the scale values of the images in each group. The results for the “similar” grayscale images can be seen in Figure 5 (a). Here, there was no significant difference among any of the three methods we tested. This is not surprising, given the visual similarity of the grayscale images. Figure 5 (b) shows the results for the images from Figure 1. Here, the standard mapping of luminance to grayscale and our mapping with $K = 1$ did not show any significant difference, but our mapping with $K = 100/c_{max}$ was significantly preferred. Figure 5 (c) shows the results of handling all images together.

6. Re-coloring for Color-Deficient Observers

We can use also Equation 1 to create a false color image that will preserve image information that might be lost if the orig-




Figure 5: Experimental results: *Scalings resulting from paired-choice experiment. (a) Scales from images in Figure 4, (b) scales from the first three images in Figure 1, and (c) scales from all data. The vertical scale is dimensional-less with arbitrary origin.*

inal color image were viewed by a color-deficient observer. The objective here is *not* to create an image that contains the same level of detail for all viewers, that is, for observers with normal vision and those with various color deficiencies. Rather, it is to create multiple versions of the image, each tailored to the visual characteristics of the viewer.

We wish to use MDS to map three-dimensional image colors to the two-dimensional surface of colors distinguishable by each class of color-deficient observer. In CIELAB color space, the surface of distinguishable colors, as simulated by Brettel *et al.* [BVM97], is not flat. This suggests a need to parameterize the surface and map the computed configuration with this parameterization. Nevertheless, this surface has low curvature, and we can approximate using a plane containing the luminance axis. We have found normals of the approximating planes to be $(0.00, 0.99, 0.14)$ for the protan and deutan cases and $(0.0, -0.58, 0.81)$ for the tritan case.

We simultaneously eliminate the need for constraints on luminance and reduce the problem to one-dimensional by setting output luminance to input luminance. Nevertheless, we will need to constrain the g_i values to the available gamut, and the chromaticity range is a function of luminance. We compute the bounds of the gamut, along the approximating plane, at the luminance value of each input color. These boundaries then form the upper and lower bound constraints, g_i^u and g_i^l . This results in the linear program

$$\begin{aligned} & \text{minimize} && \sum_i v_i && (19) \\ & \text{subject to} && v_i + Q_i g && \geq K L_i y, \\ & && v_i - Q_i g && \geq -K L_i y \\ & && g_i^l && \leq g_i \leq g_i^u \end{aligned}$$

Figures 6 and 7 show the results of our re-coloring method. In each case the proportionality constant, K , was 1. The similarity of each of the re-colored images to the simulated, color-deficient view of that image demonstrates the validity of our planar approximation of the color-deficient gamut. Timing results recorded on a 2.4 Ghz PC are shown in Table 2.

	Colors	Solving	Interpolating
Berries	1,415,424	24.5 sec	723.1 sec
Flower	338,010	25.7 sec	164.7 sec
Balls	416,880	20.0 sec	253.4 sec

Table 2: Timing: *Results for the re-colored images shown in Figures 6 and 7. The method was applied to a quantized set of 256 colors and then interpolated over the image colors.*

7. Conclusion

We have suggested a technique for mapping color images to grayscale images in a way that preserves image information and preserves important luminance gradients. We have also shown that a simple extension of this technique can be used to re-color images for viewing by color-deficient observers. In each case, the solutions are approximate minima of a STRESS function that are found by iteratively solving a sequence of linear programming problems.

Controlled assessment of our grayscale mapping results by a group of 17 human subjects has revealed that, for some images, our technique performs no worse than the traditional method, and, for other images, our technique performs much better than the traditional method, provided we use a proportionality constant of $K = 100/c_{max}$.

The ability to easily interpolate a solution for a small set of colors over an arbitrary set of colors is important for sequences of images as well as sub-images. To avoid flickering artifacts, some sense of temporal coherence must be maintained, and the interpolation mechanism allows for a variety of strategies. We are currently evaluating a number of techniques related to this problem.

8. Acknowledgments

This work was supported in part by the ERC Program of the U.S. National Science Foundation under award EEC-9731680 and the ITR Program of the National Science Foundation under award ACI-0113139. Thanks to Andrew Duchowski for his help in experiment design.




Figure 6: Re-coloring two images from Figure 1: (a) *Simulated protanope (first row) and deuteranope (second row) views, (b) our re-coloring, and (c) simulated protanope (first row) and deuteranope (second row) views of (b).*



Figure 7: Tritanopic deficiencies: (left to right) original color image, image as seen by a simulated tritanopic viewer, image recolored for a tritanopic viewer, and the recolored image as seen by a simulated tritanopic viewer.

References

[Ada81] ADAMS A.: *The Negative*. Little, Brown, and Company, 1981.

[BFA96] BRAUN K., FAIRCHILD M., ALESSI P.: Viewing techniques for cross-media image comparisons. *Color Research and Application* 21, 1 (1996).

[BVM97] BRETTEL H., VIÉNOT F., MOLLON J.: Computerized simulation of color appearance for dichromats. *Journal of the Optical Society of America A* 14, 10 (October 1997).

[CC94] COX J. F., COX M. A. A.: *Multidimensional Scaling*. Chapman & Hall, 1994.

[dST03] DE SILVA V., TENENBAUM J.: Global versus local methods for nonlinear dimensionality reduction. *Advances in Neural Information Processing Systems* 15 (2003), 705–712.

[DW] DOUGHERTY R., WADE A.: Daltonize. <http://www.vischeck.com/daltonize/>.

[FS75] FLOYD R. W., STEINBERG L.: An adaptive algorithm for spatial gray scale. *SID 75 Digest* (1975), 36–37.

[GKN04] GANSNER E., KOREN Y., NORTH S.: Graph drawing by stress majorization. In *12th International Symposium on Graph Drawing* (2004).

[GOTG05] GOOCH A., OLSEN S., TUMBLIN J., GOOCH

- B.: Color2gray: Saliency-preserving color removal. *ACM Transactions On Graphics* 24, 3 (2005).
- [HAM02] HUBERT L. J., ARABIE P., MEULMAN J. J.: Linear unidimensional scaling in the l_2 -norm: Basic optimization methods using MATLAB. *Journal of Classification* 19 (2002), 303–328.
- [HB96] HARIKUMAR G., BRESLER Y.: Feature extraction for exploratory visualization of vector valued imagery. *IEEE Transactions on Image Processing* 5 (September 1996), 1324–1334.
- [Hec82] HECKBERT P.: Color image quantization for frame buffer display. In *Computer Graphics (Proceedings of ACM SIGGRAPH 82)* (1982), pp. 297–307.
- [ITK*03] ICHIKAWA M., TANAKA K., KONDO S., HIROSHIMA K., ICHIKAWA K., TANABE S., FUKAMI K.: Web-page color modification for barrier-free color vision with genetic algorithm. *Lecture Notes in Computer Science* 2724 (2003), 2134–2146.
- [ITK*04] ICHIKAWA M., TANAKA K., KONDO S., HIROSHIMA K., ICHIKAWA K., TANABE S., FUKAMI K.: Preliminary study on color modification for still images to realize barrier-free color vision. In *IEEE International Conference on Systems, Man and Cybernetics* (2004).
- [Jol02] JOLLIFFE I. T.: *Principal Component Analysis*, second ed. Springer-Verlag, 2002.
- [KH66] KRUSKAL J. B., HART R. E.: A geometric interpretation of diagnostic data from a digital machine: Based on a study of the Morris, Illinois electronic central office. *Bell System Technical Journal* 45, 6 (1966).
- [Kon90] KONDO S.: A computer simulation of anomalous color vision. In *Color Vision Deficiencies*. Kugler & Ghedini, 1990, pp. 145–159.
- [Man96] MANDUCA A.: Multispectral image visualization with nonlinear projections. *IEEE Transactions on Image Processing* 5 (October 1996), 1486–1490.
- [MG88] MEYER G., GREENBERG D.: Color-defective vision and computer graphics displays. *IEEE Computer Graphics and Applications* 8, 5 (September 1988), 28–40.
- [MLM95] MITRA S. K., LI H., MANJUNATH B. S.: Multisensor image fusion using the wavelet transform. *Computer Vision, Graphics, and Image Processing* 57, 3 (1995), 627–640.
- [Nat04] NATIONAL SCIENCE FOUNDATION: Grant proposal guide, 2004. NSF 04-23.
- [RAGS01] REINHARD E., ASHIKHMIN M., GOOCH B., SHIRLEY P.: Color transfer between images. *IEEE Computer Graphics and Applications* 21, 5 (2001), 34–41.
- [RGW05] RASCHE K., GEIST R., WESTALL J.: Detail preserving reproduction of color images for monochromats and dichromats. *IEEE Computer Graphics & Applications* 25, 3 (May 2005).
- [RS00] ROWELS S. T., SAUL L. K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* 290, 5500 (December 2000), 2323–2326.
- [SCB88] STONE M. C., COWAN W. B., BEATTY J. C.: Color gamut mapping and the printing of digital color images. *ACM Transactions on Graphics* 7, 4 (October 1988), 249–292.
- [Soc00] SOCOLINSKY D. A.: *A Variational Approach to Image Fusion*. PhD thesis, The Johns Hopkins University, April 2000.
- [SOM*87] SCHMIEDL U., ORTHENDAHL D. A., MARK A. S., BERRY I., KAUFMAN L.: The utility of principal component analysis for the image display of brain lesions: A preliminary, comparative study. *Magnetic Resonance in Medicine* 4 (1987), 471–486.
- [SOS98] STOLLNITZ E. J., OSTROMOUKHOV V., SALESIN D.: Reproducing color images using custom inks. In *Proceedings of ACM SIGGRAPH 1998* (1998), pp. 267–274.
- [SW02] SOCOLINSKY D. A., WOLFF L. B.: Multispectral image visualization through first-order fusion. *IEEE Transactions on Image Processing* 11, 8 (August 2002).
- [TdSL00] TENENBAUM J. B., DE SILVA V., LANGFORD J. C.: A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 5500 (December 2000), 2319–2323.
- [Tor67] TORGERSON W.: *Theory and Methods of Scaling*. Wiley, 1967.
- [TR93] TUMBLIN J., RUSHMEIER H.: Tone reproduction for realistic images. *IEEE Computer Graphics and Applications* 13, 6 (November 1993), 42–48.
- [WA97] WALRAVEN J., ALFERDINCK J. W.: Color displays for the color blind. In *IS&T and SID 5th Color Imaging Conference* (1997), pp. 17–22.
- [Wan95] WANDELL B.: *Foundations of Vision*. Sinauer Associates, Inc., 1995.