

An Efficient Naturalness-Preserving Image-Recoloring Method for Dichromats

Giovane R. Kuhn, Manuel M. Oliveira, and Leandro A. F. Fernandes

Abstract— We present an efficient and automatic image-recoloring technique for dichromats that highlights important visual details that would otherwise be unnoticed by these individuals. While previous techniques approach this problem by potentially changing all colors of the original image, causing their results to look unnatural to color vision deficient, our approach preserves, as much as possible, the image's original colors. Our approach is about three orders of magnitude faster than previous ones. The results of a paired-comparison evaluation carried out with fourteen color-vision deficient (CVDs) indicated the preference of our technique over the state-of-the-art automatic recoloring technique for dichromats. When considering information visualization examples, the subjects tend to prefer our results over the original images. An extension of our technique that exaggerates color contrast tends to be preferred when CVDs compared pairs of scientific visualization images. These results provide valuable information for guiding the design of visualizations for color-vision deficient.

Index Terms—Color-contrast enhancement, Color-vision deficiency, Recoloring algorithms, Information and Scientific Visualization

1 INTRODUCTION

Individuals with normal color vision, called *trichromats*, present three kinds of cone cells (retinal photoreceptors associated with color sensation). People who are missing one class of cones are called *dichromats*, and can be further classified as *protanopes*, *deuteranopes*, and *tritanopes*, depending on whether the missing cones are more sensitive to the long, medium, or short wavelengths of the visible spectrum, respectively. Natural variations of some proteins may result in changes in the cones' photopigments, making them more sensitive to a different band of the visible spectrum [26]. Individuals with altered photopigments are called *anomalous trichromats* and, likewise, can be classified as *protanomalous*, *deuteranomalous*, or *tritanomalous*. A much rarer situation is characterized by individuals having a single or no kind of cone cells, who are called *monochromats*. Some estimates indicate that for the male population, approximately 7.40% of European descendants, 4.17% of Asians, and 2.61% of Africans have some kind of red-green color vision deficiency (CVD) [26]. For the female population, the numbers are smaller: 0.50%, 0.58%, and 0.54%, respectively. According to these numbers, approximately 200 million people should have some kind of CVD. Among the European descendants, the only ethnic group for which one can find detailed statistics, the numbers are summarized in Table 1.

Color vision deficiency tends to impose several limitations, especially for dichromats and monochromats. Children often feel frustrated by not being able to perform color-related tasks [8], and adults tend to face difficulties to perform some daily activities. Some professional fields, such as biology, chemistry, geology, fashion design, electronics, and others that require interpreting scientific and information visualization data may be especially challenging for color vision deficient [20]. Recently, several techniques have been proposed to recolor images highlighting visual details missed by dichromats [11, 13, 23, 24, 31]. Although these techniques use different strategies, they all approach the problem by potentially changing all colors of the original image. In consequence, their results tend to look unnatural to color vision deficient. Moreover, they tend to present high computational costs, not scaling well with the number of colors and the size of the input images.

- *Giovane R. Kuhn, Manuel M. Oliveira, and Leandro A. F. Fernandes* with Instituto de Informática - UFRGS. E-mails: {grkuhn | oliveira | laffernandes}@inf.ufrgs.br.

Manuscript received 31 March 2008; accepted 1 August 2008; posted online 19 October 2008; mailed on 13 October 2008.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

We present an efficient and automatic technique for enhancing color contrast for dichromats that preserves, as much as possible, the naturalness of the original colors (Section 3). Given the dichromats' reduced color gamut, in the context of this paper *naturalness preservation* should be understood under the *as much as possible* constraint. Our algorithm is based on a mass-spring optimization and presents several other desirable features: (i) it is deterministic; (ii) it satisfies a *global consistency* property (*i.e.*, all pixels with the same color in the original image will be mapped to the same shade in the recolored image); (iii) it preserves the original image luminance and (iv) it can be efficiently implemented on modern GPUs. A paired-comparison evaluation with fourteen color-vision deficient indicated their preference for the results produced by our technique for information visualization examples. For a set of scientific visualization images, the subjects preferred an extension of our technique that exaggerates color contrast.

Figure 1 illustrates a result produced by our technique and compares it with an image obtained using the approach of Rasche et al. [24]. Figure 1 (a) shows a photograph (839×602 pixels) used as reference, while (b) is a simulated view of a protanope for the same image. The simulation was performed using the algorithm described by Brettel et al. [1]. Figure 1 (c) shows a simulated view of a protanope after the image in (a) has been recolored using Rasche et al.'s approach. Figure 1 (d) shows the simulated view of a protanope for the result produced by our technique for the same image. Note that although the colors of the petals have changed to enhance contrast, the colors of the flowers' nuclei and the background foliage have been preserved as much as possible. For this example, the GPU implementation of our algorithm performs about 2,000× faster than Rasche et al.'s approach. Our CPU implementation is still approximately 370× faster than Rasche et al.'s implementation.

2 RELATED WORK

Recoloring techniques for color-vision deficient can be broadly classified as *user-assisted* and *automatic* techniques.

User-Assisted Recoloring Techniques: Daltonize [4] uses three user-provided parameters for recoloring images for protanopes and deuteranopes. Such parameters specify how much of the red-green channel should be stretched, projected into the luminance channel, and projected into the yellow-blue channel, respectively. Working on the HSL color space, Iaccarino et al. [9] modulate the original image colors using six user-provided parameters. The quality of the results obtained with user-assisted recoloring techniques is highly dependent on the user-provided parameters.

Automatic Recoloring Techniques: Ichikawa et al. [10] used an objective function to recolor web pages for anomalous trichromats preserving the ratio between the color differences perceived by trichro-



Fig. 1. Images of natural flowers: (a) Photograph. (b) Same image as perceived by protanopes (*i.e.*, individuals without *red* cones). (c) Simulated view of a protanope for a contrast-enhanced version of the photograph recolored by Rasche et al.'s [24] approach. (d) Simulated view of a protanope for the result produced by our technique. Note how our approach enhanced the overall image contrast by selectively changing only the colors for which there is a significant perceptual difference between trichromats and dichromats. As a result, it preserved the naturalness of the colors (from the perspective of the dichromat) of the flowers' nuclei and of the background foliage (compare images (b) and (d)). For this 839×602 -pixel image, our approach performs approximately $2,000 \times$ faster than Rasche et al.'s technique.

mats and by anomalous trichromats. A genetic algorithm was used to minimize the objective function. Ichikawa et al. [11] extended their previous technique for use on color images, but they did not consider the preservation of color naturalness. Wakita and Shimamura [31] proposed a technique to recolor documents (*e.g.*, web pages, charts, maps) for dichromats using three objective functions aiming, respectively, at: (i) color contrast preservation, (ii) maximum color contrast enforcing, and (iii) color naturalness preservation. However, the colors for which naturalness should be preserved must be specified by the user. The objective functions are then combined by weighting user-specified parameters and optimized using simulated annealing. They report that documents with more than 10 colors could take several seconds to be optimized. Jefferson and Harvey [13] select a set of key colors from the difference between the trichromat's and dichromat's image histogram and use four objective functions to preserve brightness, color contrast, colors in the available gamut, and color naturalness. Again, the user must specify the set of colors whose naturalness should be preserved. They optimize the combined objective functions using a method of preconditioned conjugate gradients. They report times of several minutes to optimize a set of 25 key colors on a P4 2.0 GHz using a Matlab implementation.

Rasche et al. [23] proposed an automatic recoloring technique for dichromats as an optimization that tries to preserve the perceptual color differences between all pairs of colors using an affine transformation. Such transformation, however, does not capture color variations along many directions and does not ensure that the mapped colors are all within the available gamut. Rasche et al. [24] addressed these limitations using a constrained multivariate optimization procedure applied to a reduced set of quantized color, which are in turn used to optimize the entire set of colors. The authors did not consider the problem of naturalness preservation and the technique can arbitrarily change the colors of the original images (Figure 1 c). Moreover, the algorithm does not scale well with the number of quantized colors and the size of the input images. Our technique can optimize hundreds of colors in real time, and can be used to create images that have a much more natural look (Figure 1 d). Contrary to all previous automatic techniques, our approach is deterministic, always producing the same result for a given input image.

Recoloring images for dichromats is a dimensionality reduction problem and, as such, is somehow related to recent techniques used to map color to grayscale techniques [6, 7, 14].

2.1 Mass-Spring Systems

A mass-spring system consists of a set of particles (nodes) connected by springs that deform in the presence of some external forces. When compressed or stretched, the springs apply internal reaction forces to maintain their rest length [5]. Mass-spring systems are simulated by assigning some position and mass to each particle, and some rest length to each spring. The system must obey Newton's second law:

$$F_i = m_i a_i \quad (1)$$

Classification	Incidence (%)	
	Male	Female
Anomalous trichromacy	5.71	0.39
Protanomaly	1.08	0.03
Deuteranomaly	4.63	0.36
Tritanomaly	0.0001	0.0001
Dichromacy	2.28	0.03
Protanopia	1.01	0.02
Deuteranopia	1.27	0.01
Tritanopia	0.002	0.001
Monochromacy	0.003	0.00001

Table 1. Incidence of CVD among caucasians [26, 25].

where m_i is the mass of node P_i , a_i is the acceleration caused by force F_i , which is the composition of internal and external forces. Therefore, F_i can be obtained from Hooke's law by summing the tensions of all the springs that connect P_i to its neighbors P_j :

$$F_i = \sum_{j \in N} k_{ij} \left(1 - \frac{l_{ij}}{l'_{ij}}\right) (p_j - p_i) \quad (2)$$

where N is the set of neighbors linked to P_i , l_{ij} and l'_{ij} are, respectively, the rest length and current length of the spring between P_i and P_j , k_{ij} is the stiffness of that spring, and p_i and p_j are the current positions of P_i and P_j , respectively. Verlet integration is often used to express the dynamics of each node. This type of integration is frequently used in simulations of small unoriented mass-points, being especially interesting when it is necessary to place constraints on the distances between the points. With a time step Δt , the new position of a node P_i at time $t + \Delta t$ can be computed as:

$$p_i(t + \Delta t) = \frac{F_i(t)}{m_i} (\Delta t)^2 + 2p_i(t) - p_i(t - \Delta t) \quad (3)$$

Recently, some researchers have demonstrated efficient implementations of mass-spring systems on GPUs [3, 5, 28]. In each integration step, the forces acting on each mass point P_i are accumulated in a fragment shader, requiring information about the system topology, which is usually stored in two textures [3, 28]. Our recoloring approach is modeled as a mass-spring system with every mass point P_i connected to every other mass point P_j by a spring S_{ij} . This fixed and implicitly defined topology lends itself to an efficient GPU implementation, with no topology setup needed.

3 THE RECOLORING ALGORITHM FOR DICHROMATS

Our algorithm uses a mass-spring system to optimize the colors in the input image to enhance color contrast for dichromats. The color gamut of each class of dichromacy can be represented by two half-planes in

the LMS color space [1], which can be satisfactorily approximated by a single plane passing through the luminance axis [30]. Thus, for each class of dichromacy, we mapped its color gamut to the approximately perceptually-uniform CIE L*a*b* color space and used least-squares to obtain a plane Π that contains the luminance axis and best represents the corresponding gamut. This is similar to what has been described by Rasche et al. [24], who suggested the use of a single plane for both protanopes and deuteranopes as a further simplification. The angles between the recovered planes and the L*b* plane are $\theta_p = -11.48^\circ$, $\theta_d = -8.11^\circ$, and $\theta_t = 46.37^\circ$, for protanopes, deuteranopes, and tritanopes, respectively (Figure 2). These angles are used to align their corresponding planes to the L*b* plane, reducing the optimization to 1D along the b* axis (the luminance values are preserved). After the optimization, the new colors are obtained by rotating the corresponding plane back to its original orientation.




Fig. 2. Approximating each dichromatic color gamut by a plane in the CIE L*a*b* color space. $\theta_p = -11.48^\circ$, $\theta_d = -8.11^\circ$, and $\theta_t = 46.37^\circ$.

Our algorithm has three main steps: (1) *image quantization*, (2) *mass-spring optimization* of the quantized colors, and (3) *reconstruction* of the final colors from the optimized ones. The first step consists in obtaining a set Q of quantized colors from the set of all colors C in the input image I . This can be performed using any color quantization technique, such as uniform quantization, k-means, minimum variance, median cut, or color palettes.

3.1 Modeling the Problem as a Mass-Spring System

Working in the CIE L*a*b* color space, each quantized color $\vec{q}_i \in Q$ is associated to a particle P_i with mass m_i . The position \vec{p}_i of the particle P_i is initialized with the coordinates of the perceived color by the dichromat after the rotation of plane Π :

$$\vec{p}_i = M_\theta D \vec{q}_i \quad (4)$$

where D is the matrix for orthographic projection onto the dichromat's plane Π (see Appendix A), and M_θ is a rotation matrix that aligns Π with the L*b* plane. We connect each pair of particles P_i and P_j with a spring S_{ij} with elasticity coefficient $k_{ij} = 1$ (Equation 2), and rest length $l_{ij} = \|\vec{q}_i - \vec{q}_j\|$, the (Euclidean) distance between colors \vec{q}_i and \vec{q}_j in the L*a*b* space. At each optimization step, we update the positions \vec{p}_i and \vec{p}_j and compute S_{ij} 's current length as $l'_{ij} = \|\vec{p}_i - \vec{p}_j\|$. Given the restoring forces of the springs, the system will try to converge to a configuration for which $l'_{ij} = l_{ij}$ for all S_{ij} . Thus, after stabilization (or a maximum number of iterations has been reached), the perceptual distances between all pairs of new colors/positions (\vec{p}_i, \vec{p}_j) will have approximately the same perceptual distances as their corresponding pairs of quantized colors (\vec{q}_i, \vec{q}_j) from Q . The set T of optimized colors \vec{t}_i is obtained by applying, to each resulting color \vec{p}_i , the inverse rotation used in Equation 4:

$$\vec{t}_i = M_\theta^{-1} \vec{p}_i \quad (5)$$

In order to enforce color naturalness preservation, we define the mass m_i of each particle P_i as the reciprocal of the perceptual distance (in the L*a*b* space), between \vec{q}_i and $D \vec{q}_i$:

$$m_i = \frac{1}{\|\vec{q}_i - D \vec{q}_i\|} \quad (6)$$

This equation enforces that *any color perceived similarly by both trichromats and dichromats will have large masses, causing their*

corresponding particles to move less. If trichromats and dichromats perceive \vec{q}_i exactly the same way (e.g., achromatic colors), the particle would have infinite mass. In this case, we simply set the forces acting on the particle to zero (i.e., $F_i = 0$ in Equation 1).

3.2 Dealing with Local Minima

Like other optimization techniques, mass-spring systems are prone to local minima. Figure 3 (left) depicts the problem with a configuration obtained right after the quantized colors \vec{q}_i have been rotated: $\vec{q}_i^r = M_\theta D \vec{q}_i$. Particles P_1 and P_2 have large masses (m_1 is infinite) since they are perceived as having, respectively, the same and very similar colors by trichromats and dichromats. The springs ($S_{13}, S_{14}, S_{23}, S_{24}$) connecting P_1 and P_2 to P_3 and P_4 apply forces that constrain P_3 and P_4 from moving to the other half of the b* axis. Figure 4 illustrates this situation, where the resulting optimized image (c) does not represent any significant improvement over the original image perceived by the dichromat (b).

Once the plane that approximates the dichromat's gamut has been aligned to the L*b* plane, pairs of ambiguous colors with considerable changes in chrominance will have their a* color coordinates with opposite signs (e.g., the red and green colors in Figure 3 left). We use this observation and the topology of our mass spring system to deal with local minima using the following heuristic: *we switch the sign of the b* color coordinate of all rotated quantized colors whose a* coordinates are positive and whose perceptual distance between the color itself and how it is perceived by the class of dichromacy is bigger than some threshold τ* (Equation 7).

$$p_i^{b*} = \begin{cases} -p_i^{b*} & \text{if } (qr_i^{a*} > 0) \text{ and } (\|\vec{q}_i - D\vec{q}_i\| > \tau) \\ p_i^{b*} & \text{otherwise} \end{cases} \quad (7)$$

The threshold τ enforces that colors that are perceptually similar to both dichromats and trichromats should not have the signs of their b* coordinates switched.

Although at first this might look too naive because one would just be replacing some ambiguities with another ones, it has some rationale: (i) it avoids the ambiguities among some colors found in the original image and (ii) although there is the possibility of introducing new ambiguities, as we switch the sign of the b* coordinate for some colors, we are also compressing and stretching their associated springs, adding to the system a lot of potential energy that will contribute to drive the solution. Although such a heuristic cannot guarantee that the system will not run into a local minimum, it works well in practice. Figure 3 (right) illustrates the configuration obtained by applying our heuristics to the example shown to its left. Figure 4 (d) shows the corresponding result of applying this heuristic to the example shown in Figure 4 (a). According to our experience, $\tau = 15$ tends to produce good results in general and was used for all images shown in this paper and in the accompanying video.




Fig. 3. Dealing with local minima. (left) A configuration, obtained right after plane rotation, which leads to a local minimum: since P_1 cannot move at all (it is an achromatic color) and P_2 can only move a little bit, they will constrain the optimization to the yellow portion of b* axis. (right) By switching the sign of the b* coordinate of \vec{q}_r4 , the system escapes the local minimum.




Fig. 4. Color naturalness preservation: (a) Reference image. (b) Simulation of a deuteranope’s view of the image in (a). (c) Simulation of a deuteranope’s view after recoloring the reference image using the mass-spring optimization without the heuristic described by Equation 7. (d) Result produced by our technique with the use of Equation 7.

3.3 Computing the Final Color Values

The last step of the algorithm consists in obtaining the color values for all pixels of the resulting image from the set of optimized colors t_k , $k \in \{1, \dots, \|T\|\}$, where $\|T\|$ is the cardinality of T . For this task, we have developed two distinct interpolation solutions: *projection of the perceptual difference* and an *optimization with respect to T* .

Projection of the perceptual difference: Let $C_k \subset C$ be a cluster formed by all colors in C that are represented by the optimized color \vec{t}_k . The final value \vec{t}_m^k associated to color $\vec{c}_m^k \in C_k$ is then obtained as

$$\vec{t}_m^k = \vec{t}_k + (\vec{d}.L^*, r_k \vec{d}.a^*, r_k \vec{d}.b^*) \quad (8)$$

where $\vec{d}.L^*$, $\vec{d}.a^*$, and $\vec{d}.b^*$ are respectively the L^* , a^* , and b^* coordinates of the difference vector $\vec{d} = (\vec{c}_m^k - \vec{q}_k)$. $\vec{q}_k \in Q$ is the quantized color associated to the optimized color \vec{t}_k . Equation 8 guarantees that the final color \vec{t}_m^k has the same luminance as \vec{c}_m^k . r_k is an interpolation of ratios that indicates how close the transformed value \vec{t}_k is to the optimal solution. This interpolation is guided by Shepard’s method [27], a standard technique for distance-weighted interpolation:

$$r_k = \frac{\sum_{i=1}^{\|T\|} w_{ki} \frac{\|\vec{t}_k - \vec{t}_i\|}{\|\vec{q}_k - \vec{q}_i\| + \varepsilon}}{\sum_{i=1}^{\|T\|} w_{ki}} \quad (9)$$

where ε is a small value just to avoid division by zero, and $w_{ki} = 1/(\|\vec{q}_k - \vec{q}_i\|^2 + \varepsilon)$ is the distance-weighted term suggested by Shepard [27]. Equation 8 can be efficiently implemented both on a CPU and on a GPU. On the CPU, the cost to compute all cluster ratios using Equation 9 is $O(\|Q\|^2)$ for the set Q of quantized colors, and the cost to interpolate each pixel using Equation 8 given an image with $N \times N$ pixels is $O(N^2)$. Thus, the total cost of this strategy is $O(\|Q\|^2 + N^2)$.

Since the final colors are computed by optimizing the set of quantized colors Q , the quality of the results depends directly on the quality of the quantization algorithm used. According to our experience, the transformation expressed by Equation 8 produces excellent results in combination with k-means. Unfortunately, k-means is not very fast. Table 2 compares uniform quantization and k-means execution times for various image resolutions and number of clusters. In the case of uniform quantization, we discretize the RGB space using a uniform $10 \times 10 \times 10$ grid. The quantized color is given by the grid color closest to the input one. Uniform quantization is fast, but tends to lead to poorer results when used in combination with the transformation defined by Equation 8.

Optimization with respect to T : One can benefit from the speed of uniform quantization by performing a more expensive reconstruction of the final set of colors, achieving results similar to the ones obtained when using k-means. In this case, the final shading of each pixel is obtained by optimizing it against the set of already optimized colors T . This is modeled by setting up a mass-spring system, and creating springs between the current pixel (treated as a particle initialized with Equation 7) and all optimized colors t_k , $k \in [1, \dots, \|T\|]$. For this refining optimization stage, we force the particles associated to the

Quantization technique	Image resolution (pixels)		
	640x480	800x600	1024x768
Uniform	0.055	0.087	0.150
K-means 64	2.082	3.338	5.517
K-means 128	3.932	6.029	10.432
K-means 256	7.545	11.972	20.049

Table 2. Time (in seconds) to quantize images with various resolutions and numbers of clusters on a 2.2 GHz PC.

shades in T to remain stationary by setting the forces acting on them to zero (F_i in Equation 1). For each color $\vec{c}_m^k \in C$, we define its mass as $m_{c_m} = 1/\|\vec{c}_m^k - D \vec{c}_m^k\|$. This way, we allow a color to change by an amount directly proportional to the difference of how it is perceived by trichromats and dichromats. *This mechanism guarantees the color naturalness in the resulting image.*

The cost of this optimization procedure is $O(\|Q\|^2 + \|Q\|N^2)$ for an $N \times N$ image and can be significantly higher than the mapping defined by Equation 8. However, since the color of each output pixel can be computed independently of each other, the computation can be efficiently implemented in a fragment program. Table 3 compares the times for recoloring images with various resolutions using different algorithms. *MS CPU* and *MS GPU* are respectively the CPU and GPU versions of our mass-spring algorithm using Equation 8 to obtain the final colors. *MS GPU full* optimizes the final colors with respect to T as described in the previous paragraph. Table 3 shows that in all of its variations, our approach is a few orders of magnitude faster than Rasche et al.’s approach. All images (and execution times) reported for the technique of Rasche et al. [24] were obtained using the software available at [22]. Since Rasche et al.’s technique is not deterministic, for each example shown in the paper, we ran their software three times and selected the best image result.

Recoloring technique	Image resolution (pixels)		
	640x480	800x600	1024x768
Rasche	225.16	349.31	580.49
MS CPU	0.41	0.46	0.54
MS GPU	0.20	0.22	0.26
MS GPU (full)	0.22	0.23	0.27

Table 3. Performance of various algorithms on images of different resolutions. Times measured in seconds on a 2.2 GHz PC with 2 GB of memory and on a GeForce 8800 GTX GPU. Quantization times not included. All techniques used 128 quantized colors. Mass-spring (MS) optimized the set of quantized colors using 500 iterations. The *full* version obtained the final colors using 100 iterations per pixel.

3.4 Exaggerated Color Contrast

For applications involving non-natural images (*e.g.*, scientific and information visualization) contrast enhancement is probably more important than preserving color naturalness. This comes from the fact that colors used in such applications tend to be arbitrary, usually having little connection to the viewer’s previous experiences in the real world. In scientific visualization, the details presented by the datasets are interactively explored via transfer functions [21]. Until now, transfer function design has largely ignored the limitations of color-vision deficient. Popular color scales usually range from red to green, colors that are hard to be distinguished by most CVDs.

By supporting real-time recoloring of transfer functions and images, our approach can assist color-vision deficient to exploit the benefits of scientific visualization. Such assistance can be even improved with the use of some color contrast exaggeration, achieved with the following changes in our image-recoloring algorithm:

1. Modifying the springs’ rest lengths to exaggerate the contrast between colors during the optimization process: $l_{ij} = \alpha \|\vec{q}_i - \vec{q}_j\|$, where α is a scalar used to exaggerate the perceptual difference between any pair of colors \vec{q}_i and \vec{q}_j ;

2. Defining the mass of particle P_i as: $m_i = 1/\|(a_i^*, b_i^*)\|$, where $\|(a_i^*, b_i^*)\|$ is the distance from color \bar{q}_i to the luminance axis L^* . Thus, less saturated colors present bigger masses and tend to move less. This preserves achromatic colors;
3. Initializing the mass-spring system with $\tau = 0$ (Equation 7), since we do not need to preserve the naturalness of colors.

4 RESULTS

We have implemented the described algorithms using C++ and GLSL, and used them to recolor a large number of images. The reported times were measured using a 2.2 GHz PC with 2 GB of memory and on a GeForce 8800 GTX with 768 MB of memory. Figures 1, 5, and 6 compare the results of our technique against Rasche et al.’s approach, which is the only truly automatic competing technique. Table 4 summarizes the performances of both algorithms. For these experiments, we used two different quantization techniques for the purpose of illustration: *k-means* and *uniform quantization*. For k-means we allowed up to 128 clusters. Uniform quantization was performed on a $10 \times 10 \times 10$ grid, allowing up to 1,000 clusters. The actual numbers of quantized colors (clusters) were obtained automatically by these algorithms observing the specified limits. k-means and uniform quantization were used in combination with the following reconstruction strategies: projection of the perceptual difference, and optimization with respect to T, respectively. Figure 1 (*Flowers*) used uniform quantization with 227 colors; Figure 5 (*Bell Peppers*) used k-means with 127 colors; Figure 6 (*Chinese Garden*) used k-means with 128 colors; Figure 7 (*Flame*) used uniform quantization with 84 colors; and Figure 8 (*Carp*) used uniform quantization with 131 colors.

Image (size)	Rasche Time	CPU Time	GPU Time
Flowers (839 × 602)	315.84	0.85	0.16
Bell Peppers (321 × 481)	114.63	0.27	0.09
Chinese Garden (239 × 280)	44.06	0.23	0.08
Flame (288 × 184)	52.15	0.19	0.08
Carp (629 × 241)	121.25	0.26	0.09

Table 4. Performance comparison between our technique and Rasche et al.’s for images of various sizes and different quantization strategies. Time (in seconds) shows that our technique scales well with image size.

Figure 1 (*Flowers*) has 839×602 pixels and shows an example of an image recolored for protanopes. Our GPU implementation recolored this image in 0.158 seconds. This is $2,000\times$ faster than Rasche et al.’s approach. Our CPU implementation was still $372\times$ faster than Rasche et al.’s for this example.

Figure 5 (*Bell Peppers*) illustrates the case of image recoloring for deuteranopes. While Rasche et al.’s approach (c) enhanced the contrast among the colors of the peppers, our technique also preserved the colors of the crates, yellow peppers, and other vegetables in the background as much as possible (d).




Fig. 5. Color peppers: (a) Original image. (b) Simulation of a deuteranope’s view of image (a). (c) Simulation of a deuteranope’s view for the results produced by Rasche et al.’s technique. (d) Simulation of a deuteranope’s view for the results produced by our approach.




Fig. 6. Photograph of a Chinese garden: (a) Color image. Simulated views of tritanopes for: (b) the original image, (c) the recolored image by Rasche et al.’s approach, and (d) the recolored image using our technique. Note the blue sky and the enhanced contrast for the flowers.

Chinese Garden (Figure 6) provides an example of image recoloring for tritanopes. Note how our technique preserved the blue sky, while enhancing the contrast of the flowers. Rasche et al.’s approach, on the other hand, recolored the sky as pink and did not sufficiently enhance the contrast of the flowers.

Figures 7 and 8 illustrate the use of our exaggerated color-contrast approach. Figure 7 shows our results on a simulated flame. Note how difficult it is for protanopes to distinguish the red and green regions (b). Figures 7 (c) and (d) show the simulated views of a protanope for the results produced by our image-recoloring and exaggerated color-contrast approaches, respectively. Note how these two images help color vision deficient to see the dataset details.

Figure 8 shows the visualization of the carp dataset (a) and how it is perceived by deuteranopes (b). Figures 8 (c) and (d) show simulated views of a deuteranope for the results produced by our image-recoloring technique for dichromats and by our exaggerated color-contrast approach, respectively.

4.1 Evaluation by Color Vision Deficients

In order to assess the quality of our results, we performed an evaluation with fourteen individuals selected from a group of fifteen male volunteers who declared themselves as color vision deficient. The subjects had no previous experience with scientific visualization. First, the volunteers were submitted to an Ishihara test [12]. The results of the test indicated that the deficiency of one of the volunteers was very mild and, for this reason, he did not participate of the evaluation of the proposed technique. For the other volunteers, the Ishihara test suggested the following classification: 4 protanopes (ages 20 to 29), 4 protanomalous (ages 24 to 58), 2 deuteranopes (ages 17 to 32), and 4 deuteranomalous (ages 22 to 27). One should note, however, that a precise diagnosis of protanopia and deuteranopia cannot be achieved using only the Ishihara test and would require the use of an anomaloscope. The volunteers were then arranged into two groups: *protans* (protanopes and protanomalous) and *deutans* (deuteranopes and deuteranomalous). Each group evaluated a set of recolored images for its related class of dichromacy. Even though our recoloring strategy was designed for dichromats, we decided to also include anomalous trichromats (*i.e.*, protanomalous and deuteranomalous) in our study. The reason is twofold: (i) we would like to evaluate the effect of the recoloring for these individuals, whose cones’ spectral sensitivities tend to be similar to dichromats’, and (ii) to have a larger number of subjects.

For each class of dichromacy, we selected eighteen images divided into three groups of six images each. The first group (*Natural*) was intended to contain natural elements that the volunteers could relate to their previous experiences, and consisted of photographs of natural scenes. The second group (*InfoVis*) consisted of information visualization images, while the third group (*SciVis*) consisted of scientific visualization images. For the second and third groups, the colors have no particular association to the volunteers’ previous experiences. For each selected image, we took the original one and recolored it using both Rasche et al.’s and our technique. The images in the *InfoVis* and *SciVis* groups were also recolored using the exaggerated-contrast technique described in Section 3.4. Figure 9 shows thumbnail versions of




Fig. 7. Simulation of a flame: (a) Color image. Simulated views of protanopes for: (b) original image, (c) result produced by our image-recoloring technique for dichromats, and (d) result produced by our exaggerated color-contrast approach using $\alpha = 2$.




Fig. 8. Visualization of a carp dataset using a multi-dimensional transfer function: (a) Color image. Simulated view of deuteranopes for: (b) original image, (c) result produced by our image-recoloring technique for dichromats, and (d) result produced by our exaggerated color-contrast approach using $\alpha = 2$.

the images used in the experiment.

For the evaluation of the techniques, we used the method of paired comparisons [29]. For each image in the *Natural* group, the volunteers were presented with the following pairs of side-by-side images: (*Or*, *R*), (*Or*, *MS*), and (*R*, *MS*) and were asked to indicate their binary preferences according to the following subjective questions: (i) *Which image looks more natural?*, (ii) *Which image has the most pleasant contrast?*, and (iii) *What is your overall preference*. *Or*, *R*, and *MS* stand for original, recolored with Rasche et al. technique, and recolored with the proposed technique, respectively. For the *InfoVis* and *SciVis* groups, the pairs of images included all combinations of *Or*, *R*, *MS*, and *EMS* (images with exaggerated contrast, recolored as described in Section 3.4 using $\alpha = 2$). For such groups, the volunteers were asked to indicate their binary preferences only with respect to questions (ii) and (iii). Each volunteer compared a total of 18 (*i.e.*, 6×3), 36 (*i.e.*, 6×6), and 36 pairs of images for the *Natural*, *InfoVis*, and *SciVis* groups, respectively. The order in which the pairs of images were presented as well as the order of the images in each pair were defined randomly. Each question was asked for all pairs of images before changing to the next question. The order of the questions varied between subjects. The images were displayed on an LG L1953HS LCD monitor (19 inches, 1280x1024 pixels, NVIDIA GeForce 9600 GT) at approximated 50 cm from the subject. The average time for completing the experiment by the volunteers in the *protans* and *deutans* groups

were 24 and 29 minutes, respectively.

We analyzed the paired-comparison data using Turnstone’s Law [29], Case V. Turnstone’s Law allows one to measure individuals’ preference orderings for some stimuli, from a set of discrete binary choices. It is the classic tool for ranking items based on subjective choices. For the *protans*, the ranking of the preferred technique was based on 144 (8×18) binary choices for the natural images, and on 288 (8×36) binary choices for both the *InfoVis* and *SciVis* groups. For the *deutans*, the number of binary choices were 108 (6×18) and 216 (6×36), respectively. We verified that Case V fits the data well by also analyzing the results of our experiments using dual scaling [19] and Mosteller’s χ^2 test of goodness-of-fit [17], working with a significance test at the 5% level.

Figure 10 shows the average preference scores for the groups of *protans* (left column) and *deutans* (right column). For the natural images, the χ^2 has 1 degree of freedom (dof), causing the goodness-of-fit for Thurstone’s Case V to be confirmed for $\chi^2 < 3.84$. For the *InfoVis* and *SciVis* images, the χ^2 has 3 dofs, satisfying Case V for $\chi^2 < 7.81$. The numbers shown in parenthesis next to each evaluated criteria (at the bottom of the graphs) are the corresponding computed χ^2 values, which satisfy the stated conditions. For the average preference scores obtained using Case V, we also computed the corresponding 95% confidence intervals (the vertical error bars in Figure 10) using [16].

Figure 10 (a) shows the scores for the *protans* group considering only the set of natural images. As expected, the original images were ranked best regarding naturalness. Considering the preferred contrast, given the significant overlap of the confidence intervals, one can assume a tie between the original image and the one recolored with our technique. Overall preference of these subjects was for the original image. Figure 10 (c) presents the scores of the *protans* group for the set of *InfoVis* images. In this case, the images recolored with the mass-spring algorithm without contrast exaggeration was preferred both in terms of contrast and overall preference. It is interesting to note, however, that when comparing the *SciVis* images (Figure 10 e), *protans* showed a clear preference for the ones with exaggerated contrast.

Figure 10 (right) shows the corresponding scores for the *deutans* group, which are similar to the ones observed for the *protans*. The only difference is that for the set of *InfoVis* images (d), the subjects considered the contrast of the original images as good as the recolored ones. This might be explained by the small number of subjects (six) in this group, with only two deuteranopes.

We have also performed a paired-comparison evaluation of the same set of images using a group of 23 male trichromats (ages 18 to 35) with no previous experience with scientific visualization. One should recall that the recolored images matched the reduced color gamut of dichromats. Therefore, the original natural images were preferred by trichromats. When considering the set of *InfoVis* images, trichromats found the contrast of the original images slightly better than the ones of the images recolored using our approach (*MS*). However, when asked about their overall preference, there is a tie between the original and the recolored images (Figure 11 left). For the *SciVis* group, the original and recolored images with exaggerated contrast (*EMS*) were equally preferred by trichromats both in terms of contrast and overall preference (Figure 11 right). These rankings of trichromats were computed using Thurstone Law, Case V, from a total of 828 binary choices.

5 CONCLUSIONS

We presented an efficient technique for enhancing color contrast for dichromats based on mass-spring optimization. Contrary to previous automatic techniques, our approach preserves, as much as possible, the naturalness of the images as perceived by color vision deficient. Our technique can be efficiently implemented on modern GPUs and both its CPU and GPU versions are significantly faster than previous approaches. We also described an extension to the proposed image-recoloring algorithm for exaggerating color contrast.

We have evaluated the results produced by our technique using the method of paired-comparison [29] with 8 *protans*, 6 *deutans*, and 23




Fig. 9. Thumbnail version of the images used in the paired-comparison test. The images were grouped in three classes of six images each: *Natural*, information visualization (*InfoVis*), and scientific visualization (*SciVis*). Each image in the natural group contains three versions: (i) original, (ii) recolored using Rasche et al.’s approach, and (iii) recolored using our approach (MS). The *InfoVis* and *SciVis* groups also contain a recolored image with exaggerated contrast (EMS). This figure only shows the recolored images for deuteranopes. The ones for protanopes look similar.

trichromats. According to our user study, protans and deutans preferred the contrast-enhanced images recolored with the proposed regular and exaggerated-contrast algorithms when visualizing information and scientific data, respectively. It is worth noting that even though the recoloring algorithms were originally designed for dichromats, they seem to be also applicable to anomalous trichromats. This could be explained by the fact that for protanomalous and deuteranomalous the spectral responses of the cones specialized in long and medium wavelengths tend to be very similar [2], causing these two kinds of cones to behave almost as a single one.

The quality of the results produced by our techniques depends on the quality of the quantization performed in its first stage. Although we ensure a continuous mapping among the colors in any given cluster, our approach does not guarantee a continuous mapping between different clusters, which could result in false edges or spurious bands. Obtaining the final colors of all pixels by optimizing them against the set of already optimized quantized colors T (Section 3.3) minimizes the occurrence of this problem. In practice, we have not noticed any objectionable artifacts due to these limitations. The proposed techniques do not incorporate the Helmholtz-Kohlrausch color appearance effect [18], which causes a colorful stimulus to appear brighter than another equiluminant but less colorful stimulus. As a result, there is no guarantee that for two isoluminant colors, the brighter color will always be perceived as brighter after the recoloring. This effect can be observed in Figure 1, where the red petals in (a) appear brighter than the blue petals in (d), even though the two images are pixelwise isoluminant. By preserving the luminance values of the original colors, we avoid the undesirable occurrence of luminance polarity reversal between pairs of colors. A consequence of this design decision is that our approach is not effective for enhancing contrast in images that already span essentially all the chrominance plane (Figure 12).

We believe our techniques can have a positive impact on the way dichromats interact with digital media, as it finally provides a practical way of disambiguating colors without contradicting, as much as possible, their memories about how the world looks like. Our results also provide important information for guiding the design of visualizations for color-vision deficient. We hope they will inspire the design of new applications and interfaces for CVDs. We are currently extending these ideas to video sequences.

APPENDIX A: THE PROJECTION ONTO PLANE Π

The orthographic projection matrix D (Equation 4) for each class of dichromacy is obtained by the following matrix composition: $D = M_\theta^{-1} M_b M_\theta$, where M_b is the matrix for orthographic projection onto the L^*b^* plane. Thus, Equation 4 can be rewritten as:

$$\vec{p}_i = M_b M_\theta \vec{q}_i = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} q_{i_{L^*}} \\ q_{i_{a^*}} \\ q_{i_{b^*}} \end{pmatrix}$$

For each class of dichromacy, θ is the negative of the angle shown in Figure 2, as we need to align Π with the L^*b^* plane.

ACKNOWLEDGEMENTS

We are deeply thankful to our volunteers. We also thank Karl Rasche for making his code available, and the anonymous reviewers for their insightful comments. This work was sponsored by CAPES and CNPq (grants 305613/2007-3 and 142627/2007-0). Nvidia generously donated the GeForce 8800 GTX used in this research. Microsoft Brazil provided additional support. The natural images (except Figure 4, which is from [24], and the green and red apples) are from [15]. The InfoVis images are from <http://commons.wikimedia.org>. The SciVis test images were provided by Francisco Pinto, (except Flame). We are grateful to Carlos Dietrich for the discussions about mass-spring systems and for the base code to solvers on the CPU and GPU. Bárbara Bellaver and Eduardo Pons produced the supplemental video.

REFERENCES

- [1] H. Brettel, F. Viénot, and J. D. Mollon. Computerized simulation of color appearance for dichromats. *J. Opt. Soc. Am.*, 14(10):2647–2655, 1997.
- [2] P. DeMarco, J. Pokorny, and V. C. Smith. Full-spectrum cone sensitivity functions for x-chromosome-linked anomalous trichromats. *Journal Opt. Soc. of America A*, 9:1465–1476, 1992.
- [3] C. A. Dietrich, J. L. D. Comba, and L. P. Nedel. *Storing and accessing topology on the GPU: a case study on mass-spring systems*, chapter ShaderX 5, pages 565–578. 2006.
- [4] R. Dougherty and A. Wade. Daltonize. (<http://www.vischeck.com/daltonize/>), 2002. Accessed on Oct/06.
- [5] J. Georgii and R. Westermann. Mass-spring systems on the gpu. *Simul. Modelling Practice and Theory*, 13:693–702, 2005.
- [6] A. A. Gooch, S. C. Olsen, J. Tumblin, and B. Gooch. Color2gray: saliency-preserving color removal. *ACM Trans. Graph.*, 24(3):634–639, 2005.

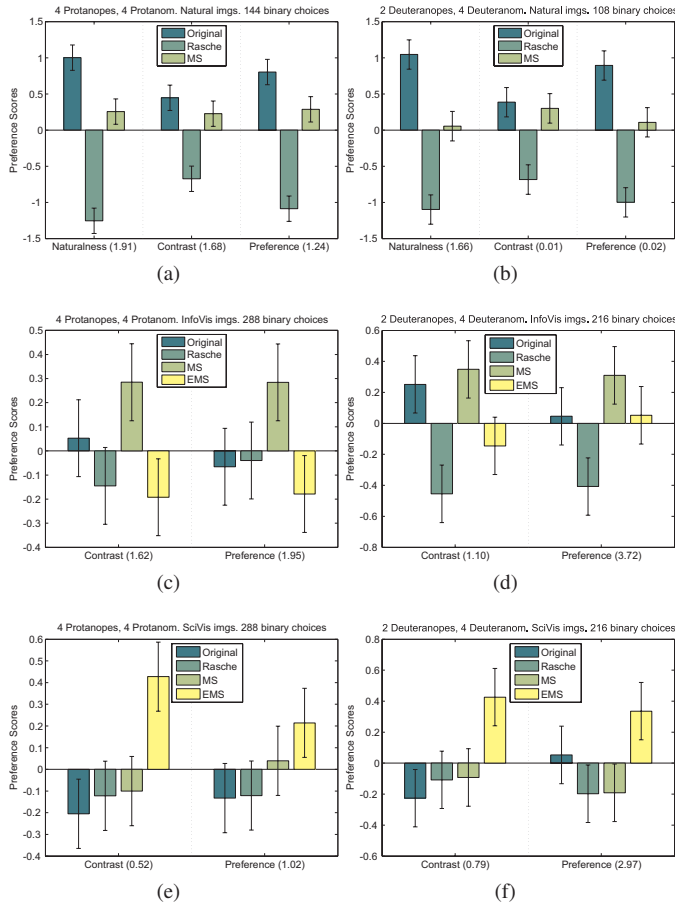


Fig. 10. Average preference scores in a paired-comparison test, shown with a 95% confidence intervals (vertical bars). The lengths of the intervals are due to the small number of subjects (6 or 8 per test). Left column shows the results for the protans. Right column shows the preferences considering only the natural images. First row: preferences considering only the natural images. Middle row: preferences for the set of *InfoVis* images. Bottom row: preferences for the set of *SciVis* images. MS: proposed technique. EMS: proposed technique with exaggerated contrast.

[7] M. Grundland and N. A. Dodgson. Decolorize: Fast, contrast enhancing, color to grayscale conversion. *Pattern Recogn.*, 40(11):2891–2896, 2007.

[8] G. Heath. The handicap of color blindness. *Journal of the American Optometric Association*, 45(1):62–69, 1974.

[9] G. Iaccarino, D. Malandrino, M. D. Percio, and V. Scarano. Efficient edge-services for colorblind users. In *Proc. of WWW '06*, pages 919–920, 2006.

[10] M. Ichikawa, K. Tanaka, S. Kondo, K. Hiroshima, K. Ichikawa, S. Tanabe, and K. Fukami. Web-page color modification for barrier-free color vision with genetic algorithm. *Lecture Notes in Computer Science*, 2724:2134–2146, 2003.

[11] M. Ichikawa, K. Tanaka, S. Kondo, K. Hiroshima, K. Ichikawa, S. Tanabe, and K. Fukami. Preliminary study on color modification for still images to realize barrier-free color vision. In *IEEE SMC '06*, volume 1, pages 36–41, 2004.

[12] S. Ishihara. *Tests for colour-blindness*. Tokio: Kanehara Shuppan Co., 1979.

[13] L. Jefferson and R. Harvey. Accommodating color blind computer users. In *Proc. of ASSETS '06*, pages 40–47, 2006.

[14] G. R. Kuhn, M. M. Oliveira, and L. A. F. Fernandes. An improved contrast enhancing approach for color-to-grayscale mappings. *The Visual Computer*, 24(7-9):505–514, July 2008.

[15] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.




Fig. 11. Average preference scores for the same paired-comparison test applied to 23 trichromats. Results for the set of: *InfoVis* images (left) and *SciVis* images (right).




Fig. 12. Simulation of a Raleigh-Taylor instability problem (courtesy of the Lawrence Livermore National Labs). The colors span most of the red-green and yellow-blue axes of the $L^*a^*b^*$ chrominance plane (a). In such a case, our technique cannot enhance the dichromat's perceived contrast. Simulated views of a protanope for the original and recolored images using our regular and exaggerated techniques (b)-(d).

[16] E. Montag. Louis Leon Thurstone in Monte Carlo: creating error bars for the method of paired comparison. In *Proc. SPIE*, volume 5294, pages 222–230, 2003.

[17] F. Mosteller. Remarks on the method of paired comparisons: A test of significance for paired comparisons when equal standard deviations and equal correlations are assumed. *Psychometrika*, 16(2):207–218, 1951.

[18] Y. Nayatani. Simple estimation methods for the Helmholtz-Kohlrausch effect. *Color Res. Appl.*, 22:385–401, 1997.

[19] S. Nishisato. *Elements of dual scaling: an introduction to practical data analysis*. Lawrence Erlbaum Associates, New Jersey, 1994.

[20] M. M. Olsen and K. R. Harris. *Color vision deficiency and color blindness: an introduction to the problem*. Fern Ridge Press, 1988.

[21] H. Pfister, B. Lorenzen, C. Bajaj, G. Kindlmann, W. Schroeder, L. S. Avila, K. Martin, R. Machiraju, and J. Lee. The transfer function bake-off. *IEEE Comput. Graph. Appl.*, 21(3):16–22, 2001.

[22] K. Rasche. Detail preserving color transformation. (<http://www.fx.clemson.edu/~rkarl/c2g.html>), 2005.

[23] K. Rasche, R. Geist, and J. Westall. Detail preserving reproduction of color images for monochromats and dichromats. *IEEE Comput. Graph. Appl.*, 25(3):22–30, 2005.

[24] K. Rasche, R. Geist, and J. Westall. Re-coloring images for gamuts of lower dimension. *Comput. Graph. Forum*, 24(3):423–432, 2005.

[25] C. Rigden. The eye of the beholder - designing for colour-blind users. *British Telecommunications Engineering*, 17, 1999.

[26] L. T. Sharpe, A. Stockman, H. Jägle, and J. Nathans. *Color Vision: From Genes to Perception*, chapter Opsin genes, cone photopigments, color vision, and color blindness, pages 3–51. Cambridge University Press, 1999.

[27] D. Shepard. A two-dimensional interpolation function for irregularly-spaced data. In *Proc. of the 23rd ACM Nat. Conf.*, pages 517–524, 1968.

[28] E. Tejada and T. Ertl. Large steps in GPU-based deformable bodies simulation. *Simulation Practice and Theory*, Special Issue on Programmable Graphics Hardware:703–715, 2005.

[29] L. L. Thurstone. A law of comparative judgment. *Psych. Rev.*, 34:273–286, 1927.

[30] F. Viénot, H. Brettel, and J. D. Mollon. Digital video colourmaps for checking the legibility of displays by dichromats. *Color Research and Application*, 24:243–252, 1999.

[31] K. Wakita and K. Shimamura. Smartcolor: disambiguation framework for the colorblind. In *Proc. of ASSETS '05*, pages 158–165, 2005.