# Color2Gray: Salience-Preserving Color Removal 



Figure 1: A color image (Left) often reveals important visual details missing from a luminance-only image (Middle). Our Color2Gray algorithm (Right) maps visible color changes to grayscale changes. Image: Impressionist Sunrise by Claude Monet, courtesy of Artcyclopedia.com.


#### Abstract

Visually important image features often disappear when color images are converted to grayscale. The algorithm introduced here reduces such losses by attempting to preserve the salient features of the color image. The Color2Gray algorithm is a 3 -step process: 1 ) convert RGB inputs to a perceptually uniform CIE $L^{*} a^{*} b^{*}$ color space, 2) use chrominance and luminance differences to create grayscale target differences between nearby image pixels, and 3) solve an optimization problem designed to selectively modulate the grayscale representation as a function of the chroma variation of the source image. The Color2Gray results offer viewers salient information missing from previous grayscale image creation methods.


CR Categories: I.4.3 [Image Processing and Computer Vision]: Enhancement-Grayscale manipulations I.4.10 [Image Processing and Computer Vision]: Image Representations-Multidimensional

Keywords: non-photorealistic, image processing, color reduction, perceptually-based rendering
*http://www.color2gray.info


Figure 2: Isoluminant changes are not preserved with traditional color to grayscale conversion. Converting an image of a reddish square whose luminance matches that of the blue background (Left) to grayscale (Middle) results in a featureless gray image. The Color2Gray algorithm incorporates chrominance changes (Right).

## 1 Introduction

If digital images are regarded solely as an optical record, then a grayscale image only needs to record light intensities using a flat spectral response. Current color to grayscale conversions already meet this goal. However, as viewers, we often expect a more ambitious result: we want digital images to preserve a meaningful visual experience, even in grayscale. We are less concerned with the accuracy of light intensities and more concerned with the preservation of visual cues that help us detect the most important, or salient, scene features. Accordingly, a black-and-white line drawing is sometimes far more expressive than a color photograph, and a garish cartoon-shaded rendering can often make important features, such as the shape, position, and reflectance of an object, more apparent.

Color documents printed in grayscale are often indecipherable. Figures and graphs in papers with saturated colors look good printed in color, but when printed in grayscale, a "red line" may appear to be the same shade of gray as a "green line". Grayscale mappings of color images that are constructed solely by approximating spectral uniformity are often woefully inadequate because isoluminant visual cues signaled only by chromatic differences are lost [Livingstone 2002], such as the reflection of the sun in Figure 1.


Figure 3: Comparison of grayscale conversions by previous methods (a-d), Photoshop's grayscale mode with auto contrast (e), and our Color2Gray algorithm (f).


Figure 4: Converting the source image (Left) to grayscale with Photoshop's grayscale mode results in a nearly featureless image (Middle). The Color2Gray algorithm creates circles that are distinct from the background (Right). Parameters: $\theta=180^{\circ} ; \alpha=8 ; \mu=$ entire image.

Vision scientists hypothesize that the human visual system does not perceive absolute values, and instead chrominance and luminance perception are based upon relative assessments, in part due to the center-surround organization of cells in the early stages of visual processing [Lotto and Purves 1999]. We propose that preserving relationships between nearby pixels in an image is much more important than representing absolute pixel values. Our work contributes to the growing trend in computer graphics of using change-based mappings for image representation and manipulation [Fattal et al. 2002; Pérez et al. 2003; Levin et al. 2004]. The technical contribution of this work is a color to grayscale image conversion algorithm based upon the human visual system's sensitivity to change. Additionally, we provide a new signed chrominance distance calculation in the CIE $L^{*} a^{*} b^{*}$ chrominance plane. The Color2Gray algorithm creates images that maintain the salience of color images by mapping chrominance and luminance changes in a source image to changes in a grayscale image.

Finding the most appropriate method to depict color differences in grayscale is also central to this research. For example, there are no luminance changes in the color image in Figure 2; adding them might seem at first to be a distortion of the original grayscale values. However, we claim that all visible color changes should cause visible changes if the grayscale image is to convey a more effective visual experience in comparison to current methods.

## 2 Related Work

Previous methods for converting RGB images to grayscale employed dot products or weighted sums to map a three dimensional color space to a single dimension, as shown in Figure 3. Other methods adopted by programs like Adobe Photoshop [Volk 2000; Brown 2004; Adobe Photoshop 2004], devised custom non-linear projections and required users to set image-dependent parameters by trial and error. These methods applied a fixed function to map a set of 2D manifolds in color space to a set of 1D points (luminance) and were ineffective at preserving chrominance differences between isoluminant pixels (Figure 4). Contrast enhancement techniques increased the dynamic range of the image, map-
ping isoluminant pixels with different chrominance to the same gray value (Figure 3e).

Converting a color image to grayscale is a dimensionality reduction problem. In the course of our research, we have explored both linear and non-linear dimensionality reduction techniques, such as principal component analysis (PCA) and space-filling curves. PCA can be employed to compute an ellipsoid in color space that is a least-squares best fit for the cloud of points formed by all the image color values. Color values in the image can then be projected on a luminance axis defined by the primary axis of this ellipsoid. The effectiveness of PCA depended upon the color space; we found that the larger the ratio between the primary and secondary axes, the more likely the technique created distinctive gray values for different colors. For example, color space plots for a CIE $L^{*} a^{*} b^{*}$ image appeared twisted and stretched in RGB space and yielded different principal component axes and different grayscale mappings. We also explored a non-linear dimensionality reduction technique, similar to Teschioni et al. [1997], which created clusters in the color space and applied space-filling curves to build a 1D parameterization of 3D color space. Both space-filling curves and PCA often created results with a high dynamic range because the length of any 1 D pathway that adequately fills a 2 D plane or 3 D volume quickly approaches infinity.

We also experimented with Poisson solvers [Fattal et al. 2002] and discovered that they work well on images such as the sunrise image (Figure 1), but not on images with large disconnected isoluminant regions (Figure 4) because Poisson solvers compute gradients over nearest neighbors, ignoring difference comparisons over distances greater than one pixel.

The color to gray problem is also similar to color quantization [Heckbert 1982] and compression of gamut dimensionality. The research by Power et al. [1996] created an image with a reduced set of inks or colors. However, both color quantization and gamut compression typically preserved luminance contrast and ignored contributions from chrominance for single ink/color images.

Contemporaneous research by Rasche et al. [2005a; 2005b] on the color to gray problem maintained the proportionality between perceived color difference and perceived luminance difference and ignored spatial arrangement of pixels. They also provided an extension to reduce three dimensional color space to a two dimensional surface to render images as they would appear to color-deficient observers. Our solution incorporates changes from luminance, chrominance, and distance combined as well as provides three simple parameters to enable users to create aesthetic and perceptually salient grayscale images.

## 3 Algorithm

The Color2Gray algorithm has three steps: first, we convert a color image to a perceptually uniform color space,


Figure 5: Color2Gray results for the color image in the center of the figure over several values of the parameter $\theta$, which divides the chrominance plane. Note, for this colorblind image, you should not be able to see the " 45 " made of dots in the source color image unless you are colorblind. Original image copyright Jay Neitz.
then compute target differences in order to combine luminance and chrominance differences, and finally, we use a least squares optimization to selectively modulate the source luminance differences in order to reflect changes in the source image's chrominance.

We express the color differences between pixels in the color image as a set of signed scalar values, and then build a grayscale version of the image with those values. For each pixel $i$ and neighbor pixel $j$, we find a signed distance scalar, $\delta_{i j}$, based upon luminance and chrominance differences between $i$ and $j$. Using these values we then derive values for the grayscale image, $g$. We refer to the grayscale difference between pixel $i$ and a neighboring pixel $j$ as $\left(g_{i}-g_{j}\right)$. A successful color to gray conversion consists of finding $g$ such that all $\left(g_{i}-g_{j}\right)$ values closely match the corresponding $\delta_{i j}$ values.

Specifying $\delta_{i j}$ is fairly involved, and can benefit from minimal amounts of user interaction (Section 3.1). Once we formalize $\delta_{i j}$ (Section 3.2), we find the output image $g$ by an iterative optimization process (Section 3.3).

### 3.1 Parameters

The Color2Gray algorithm encodes differences from a color image into luminance differences in a grayscale image. However, generating a satisfying result sometimes requires making aesthetic decisions, such as the choice to make the orange sun brighter, not darker, than its isoluminant surrounding. Thus some level of user control is desirable. Previous interactive techniques required users to perform region-based adjustments, modify color channels or individually alter pixels. The Color2Gray algorithm allows users to control the mapping of color differences to grayscale differences via three simple parameters:

- $\theta$ : controls whether chromatic differences are mapped to increases or decreases in luminance value (Figure 5).
- $\alpha$ : determines how much chromatic variation is allowed to change the source luminance value (Figure 6).
- $\mu$ : sets the neighborhood size used for chrominance estimation and luminance gradients (Figure 7).


Figure 6: Changing the $\alpha$ parameter: $\alpha=5,10,25$, respectively. Increasing alpha increases the contribution from chrominance differences but may cause dynamic range mapping problems $\left(\theta=45^{\circ}\right)$.


Figure 7: Left: Equiluminant fade from gray to blue on a background with the same luminance. Middle: Color2Gray result with $\mu=9$ neighborhood. Right: Color2Gray result with a full neighborhood ( $\theta=270^{\circ} ; \alpha=8$.). Using a small neighborhood produces a dark band through the middle of the fade and produces artifacts where the edge of the fade meets the background; whereas using the entire set of pixels as the neighborhood (full neighborhood) preserves the horizontal fade with respect to its background.

The angle $\theta$ divides the chrominance plane and determines whether a chromatic difference will darken or lighten the source luminance difference. To understand the intuition behind $\theta$, consider mapping the chrominance changes in Figure 1 to shades of gray. The sun and its reflection would benefit aesthetically from mapping the oranges to brighter gray values (making the sun lighter) and mapping the blues to darker gray values. Such a cool to warm mapping corresponds to $\theta=45^{\circ}$. Though we have implemented tools for automatically selecting $\theta$, we make the parameter available to the user. As shown by applying the Color2Gray algorithm to the color blindness test image of Figure 5, varying $\theta$ may provide insight into the salient cues available in the image.

The second parameter of the algorithm, $\alpha$ controls the amount of chromatic variation applied to the source luminance values. The luminance values, $L^{*} \in[0,100]$, are in a much smaller range than either of the chrominance axes; $a^{*}$ is in the range $[-500,500]$ and $b^{*}$ is in the range $[-200,200]$. The parameter $\alpha$ sets an upper and lower bound for how much of an effect a large chrominance difference can have on the current luminance value for a given pixel, such that the shift due to the chrominance values is within $[-\alpha, \alpha]$. In the next section, we define a function crunch() that allows us to compress large-amplitude values into the valid range, while leaving small yet significant variations around zero almost untouched. By default, $\alpha=10$. However, some images may allow for more chrominance adjustment, such as the colorblindness test image in Figure 9, created with $\alpha=15$. Figure 6 shows the effect of varying $\alpha$.

The third parameter of the algorithm, $\mu$, controls the size of the neighborhood, indicating whether a user is more interested in preserving local or global changes. Smaller neighborhoods may result in non-neighboring isoluminant regions with different chrominance values in the source image to be represented with the same grayscale value in the output image $g$. Additionally, preserving only local changes can create gradients that are not present in the source image, as shown in Figure 7. All of the figures in this paper were generated


Figure 8: Signed chrominance distance between pixels. Given (a) pixel pairs (i.e. $\left(C_{i}, C_{j}\right)=\left(C_{1}, C_{2}\right)$ or $\left.\left(C_{3}, C_{4}\right)\right)$, and (b) their chromatic difference, $\overrightarrow{\Delta C}_{i j}$, we divide the space of color differences into positive and negative halves based upon the parameter $\theta$. We set the sign of the chromatic difference to be the same as the sign of $\left(\overrightarrow{\Delta C}_{i j} \cdot \overrightarrow{v_{\theta}}\right)$, where $\overrightarrow{v_{\theta}}$ is a normalized vector defined by $\theta$ relative to the $\Delta a^{*}$ axis, as illustrated in (b).
using the entire image for neighborhood pixel comparisons, unless otherwise noted.

### 3.2 Computing Target Differences

We begin by converting the color input image to CIE $L^{*} a^{*} b^{*}$ color space, because its Euclidean distances (L2 norm) closely correspond to perceptual dissimilarity [Wyszecki and Stiles 2000]. $L_{i}$ refers to the luminance of the $i^{\text {th }}$ pixel, and $\Delta L_{i j}$ is shorthand for $\left(L_{i}-L_{j}\right)$. Similarly, $\Delta A_{i j}$, and $\Delta B_{i j}$ refer to the $a^{*}$ and $b^{*}$ channel differences between a pair of pixels. $\overrightarrow{\Delta C}_{i j}$ refers to $\left(\Delta A_{i j}, \Delta B_{i j}\right)$, the chromatic difference vector between pixel $i$ and it's neighbor pixel $j$.

In order to determine the target difference, $\delta_{i j}$, we compare the luminance difference, $\Delta L_{i j}$, with the chrominance difference, $\overrightarrow{\Delta C}_{i j}$. Since $\Delta L_{i j}$ is a scalar and $\overrightarrow{\Delta C}_{i j}$ is a 2 D vector, we first map $\overrightarrow{\Delta C}_{i j}$ onto a single dimension using the Euclidean norm, $\left\|\overrightarrow{\Delta C}_{i j}\right\|$. Next we need to choose the appropriate sign for $\left\|\overrightarrow{\Delta C}_{i j}\right\|$ because it is always positive and $\delta_{i j}$ is a signed scalar. Therefore, we introduce an angle $\theta$, which parameterizes the $\Delta a^{*} \Delta b^{*}$ chrominance plane. Let $\overrightarrow{v_{\theta}}$ be a normalized vector defined by $\theta$ relative to the $\Delta a^{*}$ axis. We set the sign of the chromatic difference to be the same as the sign of $\left(\overrightarrow{\Delta C}_{i j} \cdot \overrightarrow{v_{\theta}}\right)$, as illustrated in Figure 8. The parameter $\theta$ controls which chromatic differences are mapped to increases or decreases to the source luminance value (Figure 5), thus specifying the transformation of the $\left\|\overrightarrow{\Delta C}_{i j}\right\|$ s into a signed quantity.

Finally, we define the target differences, $\delta_{i j}$. If the absolute luminance difference is smaller than the chrominance difference, then we set $\delta_{i j}$ to be a measure of the chromatic differences; otherwise, $\delta_{i j}$ is set to the luminance differences. Formally, $\delta_{i j}$ is defined as follows.

$$
\text { Given: } \quad \begin{array}{ll}
\operatorname{crunch}(x) & =\alpha * \tanh (x / \alpha) \\
\overrightarrow{v_{\theta}} & =(\cos \theta, \sin \theta)
\end{array}
$$

then:
$\delta(\alpha, \theta)_{i j}= \begin{cases}\Delta L_{i j} & \text { if }\left|\Delta L_{i j}\right|>\operatorname{crunch}\left(\left\|\overrightarrow{\Delta C}_{i j}\right\|\right) \\ \operatorname{crunch}\left(\left\|\overrightarrow{\Delta C}_{i j}\right\|\right) & \text { if } \overrightarrow{\Delta C}_{i j} \cdot \vec{v}_{\theta} \geq 0 \\ \operatorname{crunch}\left(-\left\|\Delta \vec{C}_{i j}\right\|\right) & \text { otherwise. }\end{cases}$

We tried several methods for combining $\Delta C_{i j}$ and $\Delta L_{i j}$ in order to define $\delta_{i j}$, and of those that we considered, the discontinuous case above worked best. Using a linear or smooth blend between $\Delta C_{i j}$ and $\Delta L_{i j}$ can make the resulting images muddier because when the two terms have opposite signs, the differences cancel out.

We can automatically determine $\theta$ by examining the spread of the chrominance values of the source image and finding an axis through the $a^{*} b^{*}$ plane that divides these chrominances with the largest linear separation [Healey and Enns 1996]. If the images do not have a single best axis, such as images with a large range of chrominance values, we utilize a cool-warm parameterization with $\theta=45^{\circ}$. However, for some images and certain choices of $\theta$, there may be troublesome discontinuities in the calculated signed color differences which will push the optimization in contradicting directions. We are currently investigating adaptive color difference signing and unsigned methods that would avoid these difficulties. However, in practice, these discontinuities have not produced noticeable artifacts.

### 3.3 Solving the Optimization

Given a set of desired signed differences $\delta_{i j}$, we find a gray image $g$ that minimizes the following objective function, $f(g)$, where $\mathcal{K}$ is a set of ordered pixel pairs $(i, j)$ :

$$
\begin{equation*}
f(g)=\sum_{(i, j) \in \mathcal{K}}\left(\left(g_{i}-g_{j}\right)-\delta_{i j}\right)^{2} . \tag{1}
\end{equation*}
$$

We initialize $g$ to be the luminance channel of the source image, and then descend to a minimum using conjugate gradient iterations [Shewchuk 1994]. Equation 1 is convex, but has multiple global minima, as for any grayscale image $g$, an image with an identical optimization score can be created by shifting all pixel values up or down by a constant value $(f(g)=f(g+c))$. In order to choose a single solution from the infinite set of optimal $g$, we shift the $g$ vector returned by our solver such that it becomes as close as possible to the source luminance values (where "closeness" is defined using sum squared differences).

## 4 Performance

The cost of setting up and solving the optimization problem is proportional to the size of $\mathcal{K}$. By default our, algorithm compares every pixel to every other pixel, i.e. $\mathcal{K}$ contains all ordered pixel pairs. Given such a $\mathcal{K}$, the algorithm scales poorly. However, one can also use a reduced set, such as that implied by only considering pairs which cohabitate a $\mu \times \mu$ neighborhood. For a square $S \mathrm{x} S$ image, the cost of generating all pixel pairs will be $O\left(\mu^{2} S^{2}\right)$, or $O\left(S^{4}\right)$ for the full neighborhood case. Using an Althon ${ }^{T M} 643200+$ processor, computing images with full neighborhoods requires $12.7 \mathrm{sec}-$ onds for a $100 \times 100$ image, 65.6 seconds for a $150 \times 150$ image, and 204.0 seconds for a $200 \times 200$ image, just as you would expect for a $O\left(S^{4}\right)$ algorithm.

However, these runtimes can be improved considerably by using programmable graphics hardware to perform the pixel comparisons in parallel. In theory, the ability to process an $S \times S$ block of values simultaneously should increase the performance of the algorithm from $O\left(S^{4}\right)$ to $O\left(S^{2}\right)$. In practice, the performance of our GPU implementation is closer to $O\left(S^{3}\right)$, requiring 2.8 seconds to solve a $100 \times 100$ image, 9.7 seconds for a $150 \times 150$ image, and 25.7 seconds for a $200 \times 200$ image, using an NVIDIA Geforce ${ }^{T M}$ GT6800 graphics card.

Source code for our Color2Gray algorithm implementation, as well as additional example images, can be found at http://www.color2gray.info/.

## 5 Results and Discussion

Figure 9 compares the Color2Gray algorithm to Photoshop's grayscale mode on a wide variety of images. Isoluminant colors for state parks and water in the image in Row 2 make the island disappear in the Photoshop result. The effects of the Color2Gray algorithm are sometimes subtle, but color changes visibly darkened the region behind the head of the butterfly (Row 1) and darkened the green block while lightening the yellow block in the Cornell Box image (Row 4).

Our algorithm generally does not perform better than the traditional grayscale representation when trying to map widely spread chrominance gradations to something independently distinguishable. The first row of Figure 10 illustrates this limitation. The image contains such a wide variety of isoluminant colors that no single span of 255 gray values can express them all separately.

The Color2Gray algorithm does not provide large improvements for scenes with high dynamic range, especially natural scenes, which typically have a wide range of luminance changes. However, the method does improve any image that contains large isoluminant regions with a small number of different chrominance values such as the last row of Figure 9. Here, Color2Gray makes the two sets of hills behind the small car clearly distinguishable.

### 5.1 Augmented Color Images

We found that "Color2Gray result plus chrominance" images aid in making color images more easily comprehensible. We simply add the source image's chrominance channels to the Color2Gray result in CIE $L^{*} a^{*} b^{*}$ color space and then convert the image to RGB. This technique modifies the color images only slightly but reveals Color2Gray results in a grayscale-only printout. For example, we may be able to create a document in which all of the figures will print legibly in both color and grayscale. Adding color to the Color2Gray result also provides useful feedback for visualizing the subtle adjustments made with the Color2Gray algorithm, as illustrated in Figure 11.

## 6 Conclusion and Future Work

We are exploring extensions to the algorithm which would remove the need to specify $\theta$, for example, using an optimization function designed to match both signed and unsigned difference terms. We are also interested in investigating multi-scale methods which could allow for faster solutions on higher resolution images.

We believe better target differences, $\delta_{i j}$, should include image complexity measures as well. For example, although human easily compare widely separated luminance and chrominance values in simple images (e.g. Figure 4, with colored circles on gray), comparisons in complex images over these same distances are far more difficult (e.g. Figure 1). In complex images, reuse of the same grayscale values for different nearby colors becomes more acceptable. We plan to conduct perceptual experiments to understand the threshold for this difference as well as validating the degree to which our algorithm preserves the saliency of the color image.


Figure 9: Comparison of Color Source image, Photoshop Grayscale, and Color2Gray Results. Source images courtesy of Ana Vasileva, Yahoo!/NAVTEQ, Claude Monet, Dani Lischinski, Jay Neitz, Paul Chapman, and Leon Bli, respectively.


Figure 10: Images with a wide range of colors (Top) or with subtle changes (Bottom) result in Color2Gray results that do not provide large improvements over Photoshop grayscale. Top source image courtesy of Ana Vasileva.

The Color2Gray algorithm produces grayscale versions of color images, minimizing isoluminant mappings. While Color2Gray cannot improve all source images, we believe that its results will not be less perceptually salient than the source luminance values alone. We encourage readers to print this paper on a high quality grayscale printer and examine their printed versions of the original color, Color2Gray, and Color2Gray enhanced images.

## 7 Acknowlegments

We thank Rachel Gold, the anonymous reviewers, and many more for helping us iron out the kinks; Holger Winnemoeller for help with ISOMAP and SOR solvers; Ben Watson for pointing us to PCA; Cindy Grimm, Raquel Bujans, and Rennie Bailey, who gave us DKL Color Space pointers and code; and Pat Hanrahan, helped with pointers to CIECAM'97. Thanks to Vidya Setlur and the rest of the Northwestern University Computer Graphics Group for endless support, edits and feedback. Thanks to MidGraph participants for useful discussions and to Feng Lui for the implementation of a contrast attention model. This material is based upon work supported by the National Science Foundation under Grant No. 0415083. Additionally, this research was generously supported by a Helen and Robert J. Piros Fellowship. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

## References

Adobe Photoshop, 2004. Photoshop: Converting color images to black and white. http://www.adobe.com/tips/phs8colorbw/ main.html.

Brown, R., 2004. Photoshop tips: Seeing in black \& white. http: //www.russellbrown.com/tips_tech.html.
Fattal, R., Lischinski, D., and Werman, M. 2002. Gradient domain high dynamic range compression. ACM Transactions on Graphics 21, 3 (July), 249-256.

Healey, C. G., and Enns, J. T. 1996. A perceptual colour segmentation algorithm. Tech. Rep. TR-96-09, UBC CS.


Figure 11: Color2gray results combined with chrominance $\left(a^{*} b^{*}\right)$ from the source image make color difference more visible, and even grayscale printing reveals the improved Color2Gray result. Source images courtesy of Leon Bli, Claude Monet, and Yahoo!/NAVTEQ respectively.

Heckbert, P. S. 1982. Color Image Quantization for Frame Buffer Display. In Proceedings of SIGGRAPH, 297-307.
Levin, A., Lischinski, D., and Weiss, Y. 2004. Colorization using optimization. In ACM Transactions on Graphics.
Livingstone, M., Ed. 2002. Vision and Art: The Biology of Seeing. Harry N. Abrams, Inc., Publishers.
Lotto, R. B., and Purves, D. 1999. The effects of color on brightness. Nature Neurosci. 2, 11, 1010-1014.
Pérez, P., Gangnet, M., and Blake, A. 2003. Poisson image editing. ACM Transactions on Graphics 22, 3 (July), 313-318.
Power, J. L., West, B. S., Stollnitz, E. J., and Salesin, D. H. 1996. Reproducing color images as duotones. In Proceedings of SIGGRAPH, ACM Press, 237-248.
Rasche, K., Geist, R., and Westall, J. 2005. Detail preserving reproduction of color images for monochromats and dichromats. IEEE Comput. Graph. Appl. 25, 3.
Rasche, K., Geist, R., and Westall, J. 2005. Re-coloring images for gamuts of lower dimension. Eurographics/Computer Graphics Forum 24, 3.
Shewchuk, J. R. 1994. An introduction to the conjugate gradient method without the agonizing pain. Tech. rep., Pittsburgh, PA.
Teschioni, A., Regazzoni, C. S., and Stringa, E. 1997. A Markovian approach to color image restoration based on space filling curves. In International Conference on Image Processing Proceedings, vol. 2, $462-465$.

Volk, C., 2000. Adobe Photoshop Tip of the Week Tutorial. http://www.carlvolk.com/photoshop21.htm.
Wyszecki, G., and Stiles, W. S., Eds. 2000. Color Science: Concepts and Methods, Quantitative Data and Formulae. Wiley-Interscience.

