

Color Harmonization

Daniel Cohen-Or

Olga Sorkine

Ran Gal

Tommer Leyvand

Tel Aviv University*

Ying-Qing Xu

Microsoft Research Asia†



original image



harmonized image

Figure 1: Harmonization in action. Our algorithm changes the colors of the background image to harmonize them with the foreground.

Abstract

Harmonic colors are sets of colors that are aesthetically pleasing in terms of human visual perception. In this paper, we present a method that enhances the harmony among the colors of a given photograph or of a general image, while remaining faithful, as much as possible, to the original colors. Given a color image, our method finds the best harmonic scheme for the image colors. It then allows a graceful shifting of hue values so as to fit the harmonic scheme while considering spatial coherence among colors of neighboring pixels using an optimization technique. The results demonstrate that our method is capable of automatically enhancing the color “look-and-feel” of an ordinary image. In particular, we show the results of harmonizing the background image to accommodate the colors of a foreground image, or the foreground with respect to the background, in a cut-and-paste setting. Our color harmonization technique proves to be useful in adjusting the colors of an image composed of several parts taken from different sources.

Keywords: image enhancement, color harmonization, cut-and-paste, computational aesthetics

1 Introduction

Much of what we perceive and feel about an image is experienced through colors. Although our perception of colors depends on the context, and is culture-related, it is impossible to look at an image without being affected by the harmony of its colors. Harmonic

colors are sets of colors that hold some special internal relationship that provides a pleasant visual perception. Harmony among colors is not determined by specific colors, but rather by their relative position in color space. Generating harmonic colors has been an open problem among artists and scientists [Holtzschue 2002]. Munsell [1969] and Goethe [1971] have defined color harmony as *balance*, in an effort to transfer the concept of color harmony from a subjective perspective to an objective one. Although currently there is no formulation that defines a harmonic set, there is a consensus among artists that defines when a set is harmonic, and there are some forms, schemes and relations in color space that describe a harmony of colors [Matsuda 1995; Tokumaru et al. 2002].

Professional artists usually rely on experience and intuition to choose their favorite harmonic colors. The artist can choose a harmonic set from prescribed sets provided in handbooks (e.g., [Krause 2002]) or by using an interactive application (e.g., [Meier 1988]). Once the set is defined, the artist needs to color or recolor his/her product with this set, a task that can be tedious when the image is complex and contains many colors.

In this paper we introduce a novel application that provides the user with an automatic recoloring tool, which is useful in different settings. Given an arbitrary image, possibly a photograph, the colors of the image are modified to enhance the relationship among them and to increase their harmony. We term this color-enhancement process *color harmonization*.

Our technique can deal with an arbitrarily complex image or color palette, with a rich variety of colors. Color harmonization frees the designer from choosing a specific harmonic set, since he/she can use any desired color palette, and our automatic method subsequently optimizes the image towards a harmonic setting while remaining as faithful as possible to the original color palette. The application is suitable for both professional designers and amateurs, seeking to enhance their artistic work.

The harmonization technique can be applied in a compositing scenario in which some regions of the input image remain intact. This allows, for instance, harmonizing the colors of a background image with respect to the foreground or adapting the colors of a fore-

*{dcor|sorkine|galran|tommer}@tau.ac.il

†yqxu@microsoft.com

ground object to the background, so that together they form a harmonic color set (see Figure 1). In general, our algorithm is useful for enhancing colors in images that are comprised of a collection of parts originating from different sources and whose colors require harmonization.

2 Background and Related Work

The study of color harmony is historically intertwined with the study of the physical nature of light and color. Early discoveries in the theory of color harmony were made by such masters as Newton, Goethe, Young, and Maxwell. Modern color theory, which was developed at the beginning of the 20th century, deals mainly with representations of colors, but it also discusses color harmony [Munsell 1969; Ostwald and Birren 1969; Itten 1960]. Moon and Spencer [1944] introduced a quantitative representation of color harmony based on the Munsell color system [Munsell 1969]. At the same time, Granville and Jacobson [1944] presented a quantitative representation of color harmony based on the Ostwald color system [Ostwald and Birren 1969]. To a large degree, these works define harmony as order.

Itten [1960] introduced a new kind of color wheel in which he described color harmony, with an emphasis on hue. Itten’s color harmony theory is based on the relative positions of the hues on the color wheel. For example, from the three primary colors of cyan, magenta, and yellow, Itten designed a hue wheel of twelve colors. He referred to complementary colors as a two-color harmony. Itten also recognized the three-color harmony of hues that form an equilateral triangle, the four-color harmony of hues forming a square, the six-color harmony of a hexagon, etc. His schemes have been widely adopted by artists and designers. Based on Itten’s schemes and extensive psychophysical research, Matsuda [1995] introduced a set of 80 color schemes, defined by combining several types of hue and tone distributions. These schemes were used in [Tokumaru et al. 2002] for harmony evaluation and color design. Our color harmonization method is also based on these schemes.

There are various interactive tools that provide designers with harmonic sets (e.g., [Color Schemer 2000; Color Wheel Expert 2000; Nack et al. 2003]). Such applications provide the user with a set of harmonic colors that accommodates the user’s requirements specified by a color seed and possibly a number of other parameters. Meier et al. [1988] presented a system for designing colors based on several color rules, and applied them to a graphical user interface (GUI) building tool. The primary goal of their system was to test whether an automated mechanism would be a viable solution to the problem of choosing effective and tasteful colors. None of the above systems offers a means to harmonize a given arbitrary color image. The method we introduce in this paper automatically harmonizes a given color palette through an optimization process, and provides a means to automatically recolor an arbitrary image.

Our work is also related to general recoloring methods [Reinhard et al. 2001; Welsh et al. 2002; Levin et al. 2004; Gooch et al. 2005; Ironi et al. 2005; Rasche et al. 2005]. Automatic recoloring techniques require the user to provide a reference image. The relationship between the colors of the input and the reference images are learned and transferred to recolor the given image. One of the challenges in these techniques is to recolor the image in a coherent way [Ironi et al. 2005]. In other words, contiguous spatial regions in the input image should remain contiguous after the recoloring. Our color harmonization process uses a graph-cut optimization to enforce contiguous modification of colors in image space.




Figure 2: Harmonic templates on the hue wheel. A collection of colors that fall into the gray areas is considered to be harmonic. The templates may be rotated by an arbitrary angle. The sizes of the sectors are specified in the Appendix.

3 Harmonic Schemes

The notion of color harmony in this work is based on the schemes developed by Matsuda [Matsuda 1995; Tokumaru et al. 2002], which descend from Itten’s notions of harmony [Itten 1960], widely accepted in applicable fields involving colors. Figure 2 illustrates the eight harmonic types defined over the hue channel of the HSV color wheel. Each type is a distribution of hue colors that defines a harmonic template: colors with hues that fall in the gray wedges of the template are defined as harmonic according to this template. We refer to these distributions as templates, since they define the radial relationships on the color wheel, rather than specific colors (meaning that any template may be rotated by an arbitrary angle). The harmonic templates may consist of shades of the same colors (types *i*, *V* and *T*), possibly with complementary colors (see templates *I*, *Y*, *X*) or more complex combinations (template *L* and its mirror image). The sectors of these templates are the domains over which simple membership functions are defined. Color harmony is mainly affected by the hue channel; however, Tokumaru et al. [2002] also addressed tone distribution functions for the values of the *S* and *V* channels, and fuzzy rules for the correlation between the hue templates and the tone distributions. For details, the reader is referred to [Tokumaru et al. 2002].

The type-*N* template corresponds to gray-scale images and thus is not dealt with in this work. Note that each of the remaining seven templates consists of one or two sectors. Each hue *h* on the color wheel is then associated with one of these sectors. The simplest way is to associate *h* with the closest (in terms of arc length) sector. Thus, we define $E_{T_m(\alpha)}(p)$ as the sector border hue of template T_m with orientation α that is closest to the hue of pixel p ($m \in \{i, I, L, T, V, X, Y\}$).

Given an image, we fit a harmonic template T_m to the hue histogram of the image. We define a distance between the histogram and a template, and determine the template that best fits our image by solving an optimization problem. A template T_m together with an associated orientation α defines a *harmonic scheme*, denoted by (m, α) . Given a harmonic scheme (m, α) , we define a function $F(X, (m, \alpha))$ which measures the harmony of an image *X* with respect to the scheme (m, α) :

$$F(X, (m, \alpha)) = \sum_{p \in X} \|H(p) - E_{T_m(\alpha)}(p)\| \cdot S(p), \quad (1)$$

where *H* and *S* denote the hue and the saturation channels, respectively; the hue distance $\|\cdot\|$ refers to the arc-length distance on the hue wheel (measured in radians); hues that reside inside the sectors of T_m are considered to have zero distance from the template.




Figure 3: Overview of the color harmonization process. (a) The original image. (b) The hue histogram of the image before and after harmonization. The top histogram refers to the original image, with best-fitting I -type template superimposed. The bottom histogram shows the hues shifted to match the template sectors. (c) The resulting harmonized image. Note that the harmonization tried to preserve the original colors as much as possible. (d) The user manually rotates the template (top), and the hues are shifted accordingly (bottom). (e) The result of the manual choice of template orientation.

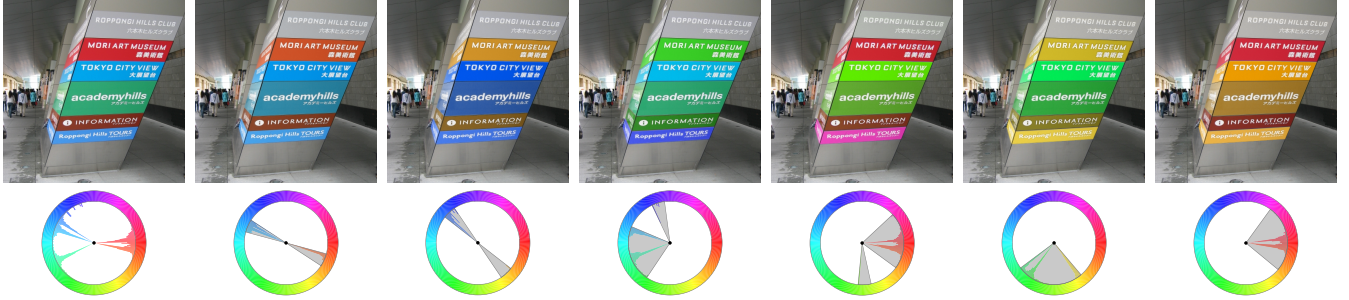


Figure 4: Manual choice of harmonic schemes. The original image and its hue histogram are displayed in the left column. Harmonic templates with various orientations result in different palettes.

Note that the above formula also considers the channel S , since the distances between colors with low saturation are perceptually less noticeable than the distances between those of high saturation. Also note that by summing over all the pixels of the image, we implicitly use a weighted average determined by the histogram of the colors.

Given an image X and a template T_m , the value of angle $\alpha \in [0, 2\pi)$ that minimizes the above expression defines the best harmonic scheme of X under the template T_m :

$$M(X, T_m) = (m, \alpha_0) \text{ s.t. } \alpha_0 = \underset{\alpha}{\operatorname{argmin}} F(X, (m, \alpha)). \quad (2)$$

The best harmonic scheme $B(X)$ of a given image X is determined by minimizing the F function over all possible templates T_m :

$$B(X) = (m_0, \alpha_0) \text{ s.t. } m_0 = \underset{m}{\operatorname{argmin}} F(X, M(X, T_m)).$$

To visualize the process, we use a circular histogram defined over the hue wheel in the HSV color space (see Figure 3(b)). To calculate the best harmonic scheme of an image, we use Brent’s algorithm [Press et al. 1992] to optimize the orientation α (Eq. 2) for each template. However, our application also allows the user to choose a specific harmonic template manually and to optimize its orientation by fitting it to the given image. The template fitting is illustrated in Figure 3.

4 Color Harmonization

The value of $F(X, T_m(\alpha))$ (Eq. 1) reflects the degree of harmony of a given image X under a given harmonic scheme $T_m(\alpha)$ (the smaller the value, the greater the agreement between the image colors and the harmonic scheme $T_m(\alpha)$). Once $T_m(\alpha)$ has been fixed, either automatically or manually, the harmony among the colors can be

optimized with respect to $T_m(\alpha)$ by shifting the colors of X ; this process is called “color harmonization”.

The harmonization process strives to preserve the original colors of the image by shifting them towards the nearest sector of the template. However, this naive definition does not take into account the spatial coherency among the image pixels. This may result in artifacts caused by “splitting” a contiguous region of the image, as demonstrated in Figure 5. This splitting occurs at some singular regions in color space, where arbitrary nearby colors are matched to two different sectors of the harmonic scheme, resulting in a discontinuity of color. The reason for this phenomenon is illustrated in Figure 5(d): hues that are almost equidistant to both sectors of the harmonic template will be shifted to different sectors of the template, which may cause discontinuous recoloring of the image, as shown in Figure 5(b).

To handle this kind of situation, we seek an optimal binary segmentation of the image that defines $E_{T_m(\alpha)}(p)$ (the sector edge of the template T_m to which each pixel p is associated and shifted). We use a segmentation approach similar to [Boykov and Jolly 2001], employing a graph-cut optimization technique. Let us denote by $\Theta_1(p)$ and $\Theta_2(p)$ the clockwise and counterclockwise nearest sector borders of template T_m around the hue of pixel p , $H(p)$. We apply the optimization to each set of pixels Ω whose hues fall into the region of the color wheel enclosed by two consecutive sector boundaries of T_m . Rather than defining $E_{T_m(\alpha)}(p)$ as the closest sector edge border to $H(p)$, the optimization should associate each pixel $p \in \Omega$ with either $\Theta_1(p)$ or $\Theta_2(p)$. This is a classic binary labeling problem, where we need to assign a label $v(p)$ to each pixel p . Note that this is essential even for templates that are comprised of a single sector, since we need to determine to which side of the sector the pixel color should be shifted. The optimal label assign-

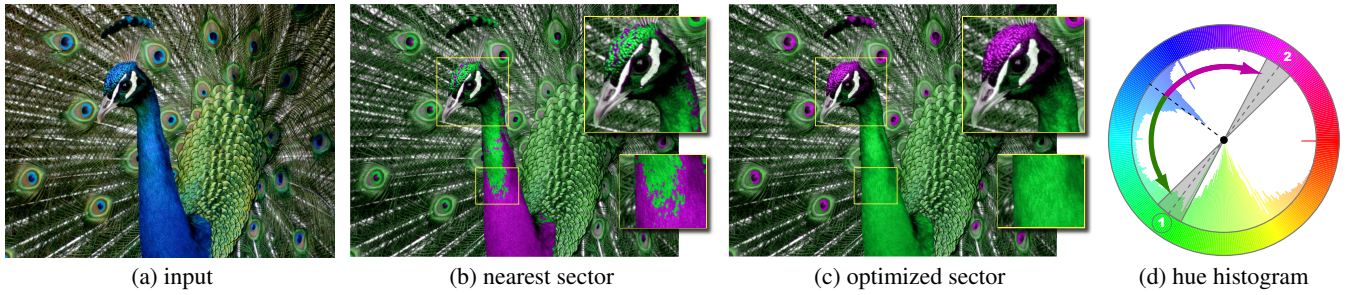


Figure 5: A naive implementation associates colors with their nearest sector, yielding the artifacts in the middle image (b). Using optimized graph-cut labeling alleviates the problem, producing a more coherent result (c). The hue histogram and the harmonic scheme are shown in (d), to visualize the source of the problem: in the naive implementation, hues which are nearly equally close to both sectors of the template may “choose” their sector arbitrarily, and this causes color discontinuities in the resulting image. Note that when the optimization (c) is applied, two pixels with exactly the same color are not necessarily shifted to the same sectors, since we take into account the spatial relation among pixels.

ment $V = \{v(p_1), \dots, v(p_{|\Omega|})\}$ minimizes the energy $E(V)$:

$$E(V) = \lambda E_1(V) + E_2(V),$$

where $E_1(V)$ accounts for the distances between the hues $H(p)$ and $H(v(p))$ (the color of the assigned sector edge), and $E_2(V)$ promotes color coherence between neighboring pixels assigned to the same label. These two terms are weighted by the saturation channel so that saturated pixels contribute more to the expression than pixels with low saturation. Specifically, $E_1(V)$ is defined as follows:

$$E_1(V) = \sum_{i=1}^{|\Omega|} \|H(p_i) - H(v(p_i))\| \cdot S(p_i),$$

where $S(p_i)$ is the saturation of p_i . Similarly,

$$E_2(V) = \sum_{\{p,q\} \in N} \delta(v(p), v(q)) \cdot S_{\max}(p, q) \cdot \|H(p) - H(q)\|^{-1}.$$

The energy $E_2(V)$ penalizes the assignment of different labels to neighboring pixels with similar hues. N is the set of neighboring pixels in Ω (4- or 8-connected); $\delta(v(p), v(q))$ equals 1 if the labels $v(p), v(q)$ are different, and 0 otherwise; $S_{\max}(p, q) = \max\{S(p), S(q)\}$. As before, the distance between the hues is measured in radians. The pixels whose original hues already lie inside a template sector are automatically labeled and serve as hard constraints in the optimization. For details on how to set up the graph-cut optimization, see [Boykov and Jolly 2001].

4.1 The shifting of colors

Once all the pixels $p \in X$ have been assigned to an appropriate sector edge $E_{T_m(\alpha)}(p)$, we can recolor the image by shifting the hues so that they reside inside the harmonic template. Instead of simple linear contraction of the hues, we apply a more elaborate scheme, so that the density of the hues that are already inside the template is not greatly affected. The hue of each pixel p is shifted in the direction of the associated $E_{T_m(\alpha)}(p)$, using the following formula: denote by $C(p)$ the central hue of the sector associated with pixel p ; then,

$$H'(p) = C(p) + \frac{w}{2} (1 - G_\sigma(\|H(p) - C(p)\|)),$$

where w is the arc-width of the template sector and G_σ is the normalized Gaussian function (such that $G_\sigma(x) \in (0, 1]$) with mean 0 and standard deviation σ . Note that similar tasks arise in gamut

mapping [Morovic and Luo 2001] and similar scaling techniques are used for this purpose.

As a result of the mapping above, the hues contract so as to fit inside the template sectors. Note that the hues originally inside the sector, contract around $C(p)$ as well (though to a lesser degree, due to the Gaussian function). This process is necessary to preserve hue monotonicity: two identically labeled pixels whose hues were originally in clockwise order on the hue wheel, will remain so after the harmonization process. The width of the Gaussian σ is a user-defined parameter that may vary between zero and w ; larger values of σ create concentration of hues near the sector centers, whereas smaller values lead to concentration near the sector boundaries. In our implementation we use $\sigma = w/2$, since it was found to provide the best color balance.

By using graph-cut optimization to assist coloring coherence, we inherit the limitation of this optimization approach: it cannot deduce that disconnected areas may belong to the same object, and thus may decide to shift each area to a different sector of the template. This is the case with the girl’s pants in Figure 7(e), since her hand divides them into two disconnected parts. To fix such cases, the user can draw a small scribble on the area that was colored wrong due to erroneous labeling; the graph-cut optimization is performed again, biasing the marked pixels towards the opposite label (see Figure 6). Of course, having a semantic segmentation of the image would enable a fully automatic solution to this problem.

5 Results and Applications

Our color harmonization process is suitable for harmonizing ordinary images, as demonstrated in Figures 1, 3, 11. Images with a large variety of colors appear more pleasing, as the colors migrate to form a harmonic set, avoiding colors that do not fit. While the



Figure 6: Graph-cut optimization may lead to erroneous labeling of disconnected image parts belonging to the same object. User assistance by marking small scribbles in such parts (in blue) and subsequent re-computation of the labeling alleviate the problem.




Figure 7: Harmonization of the background colors with respect to the harmonic scheme of the foreground (a-c) and vice versa (d-e). Note how the color of the background changes to match the different colors of the clothes; in (e) the purple and the complimentary green colors emerge on the clothes.

automatic choice of harmonic schemes strives to maintain the colors as close as possible to the original ones, the user may change the palette of the image altogether by manually rotating the template orientation (see Figures 3(d-e) and 4). Our technique is useful when parts of the image originating from different sources are combined into one image: the algorithm is capable of harmonizing the colors of the entire image, or if desired, to harmonize the colors of one part with respect to the colors of another part. This provides a useful tool for designing posters and logos, presentations, web sites and other kinds of combined imagery.

In particular, color harmonization benefits a cut-and-paste or compositing scenario, where one wants to adjust the colors of the background image to harmonize with the foreground, or vice versa. The strong effect of such harmonization is demonstrated in Figure 12, where the colors of the background image are harmonized according to the harmonic scheme of the pasted flags (which constitute a hard constraint). Figures 1 and 7 show another example, potentially useful for creating fashion catalogs: a photograph of a person is pasted onto some colorful background; the clothes are segmented by the user, and we harmonize the background colors with respect to the harmonic scheme of the foreground image, and vice versa. Note that the harmonized colors are not necessarily found in the source colors for the harmonic template; the colors may instead be complementary. Additional examples of pasting are shown in Figure 8, where posters are designed by pasting some text on top of paintings. Our algorithm assists in choosing colors for the text that harmonize with the paintings (whose colors are not be altered). Finally, the harmonization technique could be applied in interior design: in Figure 13 we demonstrate two suggestions of how to re-color the walls of a living room so that they harmonize with the rest of the room colors. Such an application can be useful for designers since it helps validate or improve a color design concept on the computer, prior to the actual execution.

Color harmonization is most useful when saturated, “man-made” colors are present in the image. Mild tones lead to mild, less striking harmonization results (see Figure 10). Natural images have inherent harmony, and our method does not change such images (at least not by automatic template choice), since their colors usually fall into some harmonic template almost perfectly (see Figure 9). Manually choosing templates with poor matching scores leads to unnatural results, as demonstrated in Figure 9. Moreover, since harmonization sometimes reduces the set of colors, objects with different colors may end up with identical coloring, as evident with the flower stems and young blossoms in Figure 9, which hinders the semantical separation of objects in the image.

6 Discussion and Conclusions

We have presented a color harmonization method that quantifies the color harmony of an image and shifts the colors towards a harmonic setting. Our technique enables the user to harmonize images and to easily tune the automatic results by simple manipulation of the color wheel. Color harmonization is particularly useful when parts originating from different sources are combined into one image, such as in a cut-and-paste setting. Our method extends the notion of cut-and-paste from the spatial domain to the color domain.

It should be noted that our technique cannot change colors with low saturation since it only alters the hue channel. Therefore, dark or



Figure 8: Harmonization of text pasted on top of a poster. The left column shows the original colors and the right displays the harmonization result. The colors of the text were shifted according to the best-fitting harmonic template of the background image. This improves the harmony score of the entire image: for example, for the top left image $F = 2771$ and for the top right $F = 609$.




Figure 9: The colors of natural images usually follow some harmonic template, and their harmonization with respect to that template leads to little or no change. Harmonization with respect to a poorly matching template gives unnatural results.

grey regions of the image remain unchanged. In addition, when changing hues, the apparent color contrast may change since human perception may react differently to varied color combinations. In future work, we would like to incorporate appropriate transformation of the saturation and lightness channels to help preserve the color contrasts.

Our method can be constrained to leave certain regions of the image intact (by masking out these regions). In future work, we would like to be able to constrain certain *colors* to remain unchanged, and optimize the shifting of the other colors with respect to such hard constraints. This will allow the user to constrain colors whose shifting might lead to unnatural results (for example, the color of the sky). In addition, color harmonization may be viewed as a special case of a more general problem of histogram matching. The optimization techniques presented in this paper can be a good starting point for developing such algorithms. In future research, we plan to explore this direction.

Acknowledgements

We would like to thank Hagit Hel-Or, Dani Lischinski and Qi Ren for insightful discussions, and Gila Hameiri, Becky Sundling, Svetlana Matusova, Andrew Nealen, Dongyu Cao and Tsui Hon Hei Hilda for their help with preparing the paper. The background image in Figure 1 is courtesy of Yanxi Liu and James H. Hays. This work was supported in part by grants from the Israel Science Foundation (founded by the Israel Academy of Sciences and Humanities) and the Israeli Ministry of Science.

References

BOYKOV, Y., AND JOLLY, M.-P. 2001. Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images. In *Proceedings of ICCV*, 105–112.



Figure 10: Harmonization of the foreground object (the girl) with respect to the background image with mild, non-saturated colors. The harmonization result is also quite mild in this case.

COLOR SCHEMER, 2000. <http://www.colorschemer.com/>.

COLOR WHEEL EXPERT, 2000. <http://www.abitom.com/>.

GOETHE, J. V. W. 1971. *Goethe's color theory*. Translated by Rupprecht Matthei. New York: Van Nostrand Reinhold Company.

GOOCH, A. A., OLSEN, S. C., TUMBLIN, J., AND GOOCH, B. 2005. Color2Gray: salience-preserving color removal. *ACM Trans. Graph.* 24, 3, 634–639.

GRANVILLE, W. C., AND JACOBSON, E. 1944. Colorimetric specification of the color harmony manual from spectrophotometric measurements. *Journal of the Optical Society of America* 34, 7, 382–395.

HOLTZSCHUE, L. 2002. *Understanding Color*. John Wiley & Sons, Inc., New York.

IRONI, R., COHEN-OR, D., AND LISCHINSKI, D. 2005. Colorization by example. In *Eurographics Symposium on Rendering*, 201–210.

ITTEN, J. 1960. *The Art of Color*. New York: Van Nostrand Reinhold Company.

KRAUSE, J. 2002. *Color Index: Over 1,100 Color Combinations, CMYK and RGB Formulas, for Print and Web Media*. F&W Publications, Inc.

LEVIN, A., LISCHINSKI, D., AND WEISS, Y. 2004. Colorization using optimization. *ACM Trans. Graph.* 23, 3, 689–694.

MATSUDA, Y. 1995. *Color design*. Asakura Shoten. (in Japanese).

MEIER, B. J. 1988. ACE: a color expert system for user interface design. In *Proceedings of the 1st Annual ACM SIGGRAPH Symposium on User Interface Software*, ACM Press, 117–128.

MOON, P., AND SPENCER, D. E. 1944. Geometrical formulation of classical color harmony. *Journal of the Optical Society of America* 34, 1, 46–60.

MOROVIC, J., AND LUO, M. R. 2001. The fundamentals of gamut mapping: A survey. *Journal of Imaging Science and Technology* 45, 3, 283–290.

MUNSELL, A. H. 1969. *A Grammar of Colors*. New York: Van Nostrand Reinhold Company.

NACK, F., MANNIESING, A., AND HARDMAN, L. 2003. Colour picking: the pecking order of form and function. In *Proceedings of the 11th ACM International Conference on Multimedia*, ACM Press, 279–282.



Figure 11: Harmonization of posters. We applied our harmonization process to change the colors of the bottom image so that they match the harmonic scheme of the top logo.



Figure 12: Harmonization of the image with respect to the harmonic scheme of the flags. Note how the colors change to match the colors of the flags. We believe that our method can potentially be extended to handle general color histogram matching.



Figure 13: Harmonization of the color of the walls with respect to rest of the room (only the walls are changed while the rest remains fixed). The original image (a) has harmony score $F = 3809$, best fitting the X template. When we harmonize the walls using this template, the colors are markedly changed (b), and the harmony score of the image in this case is $F = 1881$. Using the mirror- L template results in a more mild change (c), with harmony score $F = 2853$ (which is still an improvement compared to the original image).

OSTWALD, W., AND BIRREN, F. 1969. *The Color Primer*. New York: Van Nostrand Reinhold Company.

PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., AND FLANNERY, B. P. 1992. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA.

RASCHE, K., GEIST, R., AND WESTALL, J. 2005. Re-coloring images for gamuts of lower dimension. *Computer Graphics Forum* 24, 3, 423–432.

REINHARD, E., ASHIKHMIN, M., GOOCH, B., AND SHIRLEY, P. 2001. Color transfer between images. *IEEE Comput. Graph. Appl.* 21, 5, 34–41.

TOKUMARU, M., MURANAKA, N., AND IMANISHI, S. 2002. Color design support system considering color harmony. In *Pro-*

ceedings of the IEEE International Conference on Fuzzy Systems, IEEE Press, 378–383.

WELSH, T., ASHIKHMIN, M., AND MUELLER, K. 2002. Transferring color to greyscale images. *ACM Trans. Graph.* 21, 3, 277–280.

Appendix

The precise sizes of the harmonic template sectors (see Figure 2) are as follows: the large sectors of types V , Y and X are 26% of the disk (93.6°); the small sectors of types i , L , I and Y are 5% of the disk (18°); the large sector of type L is 22% (79.2°); the sector of type T is 50% (180°). The angle between the centers of the two sectors of I , X and Y is 180° , and for L it is 90° .