

Transferring Color to Greyscale Images

Tomihisa Welsh Michael Ashikhmin Klaus Mueller
Center for Visual Computing, Computer Science Department, SUNY at Stony Brook*



Figure 1: Colors are transferred to the second image without user intervention. Source image courtesy © Ian Britton - FreeFoto.com

Abstract

We introduce a general technique for “colorizing” greyscale images by transferring color between a source, color image and a destination, greyscale image. Although the general problem of adding chromatic values to a greyscale image has no exact, objective solution, the current approach attempts to provide a method to help minimize the amount of human labor required for this task. Rather than choosing RGB colors from a palette to color individual components, we transfer the entire color “mood” of the source to the target image by matching luminance and texture information between the images. We choose to transfer only chromatic information and retain the original luminance values of the target image. Further, the procedure is enhanced by allowing the user to match areas of the two images with rectangular swatches. We show that this simple technique can be successfully applied to a variety of images and video, provided that texture and luminance are sufficiently distinct. The images generated demonstrate the potential and utility of our technique in a diverse set of application domains.

Keywords: Image Processing, Color, Texture Synthesis, Video

1 Introduction

Color can be added to greyscale images in order to increase the visual appeal of images such as old black and white photos, classic movies or scientific illustrations. In addition, the information content of some scientific images can be perceptually enhanced with color by exploiting variations in chromaticity as well as luminance.

The task of “colorizing” a greyscale image involves assigning three-dimensional (RGB) pixel values to an image which varies along only one dimension (luminance or intensity). Since different colors may have the same luminance value but vary in hue or saturation, the problem of colorizing greyscale images has no inherently “correct” solution. Due to these ambiguities, human interaction usually plays a large role in the colorization process.

*{tfwelsh,ash,mueller}@cs.sunysb.edu,
<http://www.cs.sunysb.edu/~tfwelsh/colorize>

Copyright © 2002 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212-869-0481 or e-mail permissions@acm.org.
© 2002 ACM 1-58113-521-1/02/0007 \$5.00

Even in the case of pseudocoloring, [Gonzalez and Wintz 1987] where the mapping of luminance values to color values is automatic, the choice of the colormap is commonly determined by human decision.

Since most colorization software used in the movie industry is proprietary, detailed technical documents describing the process are generally not publicly available. However, a few web articles describe software in which humans must meticulously hand-color each of the individual image regions. For example, one software package is described in which the image is first polygonalized so the user can color individual components much like a coloring book. Then the system tracks polygons between frames and transfers colors in order to reduce the number of frames that the user must color manually [Silberg 1998]. Alternatively, photographs can be colorized using photo-editing software to manually or automatically select components of a scene. The area is then painted over with a color selected from a palette using a low opacity level.

There also exist a number of applications for the use of color in information visualization. For example, Gonzalez and Wintz [1987] describe a simple approach for pseudocoloring greyscale images of luggage acquired by X-ray equipment at an airport. The method uses separate transformations for each color channel which results in coloring objects with the density of explosives in bright orange and other objects with a blue tone. Further, color can be added to a range of scientific images for illustrative and educational purposes. In medicine, image modalities which only acquire greyscale images such Magnetic Resonance Imaging (MRI), X-ray and Computerized Tomography (CT) images can be enhanced with color for presentations and demonstrations.

Pseudocoloring is a common technique for adding color to greyscale images such as X-ray, MRI, scanning electron microscopy (SEM) and other imaging modalities in which color information does not exist. Pratt [1991] describes this method as an “image enhancement” technique because it can be used to “enhance the detectability of detail within the image.” In its most basic form, pseudocoloring is a transformation T [Pitas 1993], such that, $c(x, y) = T(f(x, y))$ where $f(x, y)$ is the original greyscale image and $c(x, y)$ is the resulting color vector for the three RGB color channels. A simplified example of this method is the application of an arbitrary color map to the data where a single, global color vector is assigned to each greyscale value. The strength of this approach is that it does not alter the information content of the original data since no extra information is introduced. For example, in an pseudo-colored MRI image there will be a one-to-one correspondence between each density value and color value, even though the color choice is arbitrary. However, by using a colormap

which does not increase monotonically in luminance, pseudocolored images may introduce perceptual distortions. Studies have found a strong correlation of the perceived “naturalness” of face images and the degree to which the luminance values increase monotonically in the colormap [Rogowitz and Kalvin 2001].

Our concept of transferring color from one image to another is inspired by work by Reinhard et al. [2001] in which color is transferred between two color images. In their work, colors from a source image are transferred to a second colored image using a simple but surprisingly successful procedure. The basic method matches the three-dimensional distribution of color values between the images and then transforms the color distribution of the target image to match the distribution of the source image. Further, swatches can be employed to match similar areas between the two images.

In this study, the greyscale image is represented by a one-dimensional distribution, hence only the luminance channels can be matched between the two images. Because a single luminance value could represent entirely different parts of an image, the statistics within the pixel’s neighborhood are used to guide the matching process. Once a pixel is matched, the color information is transferred but the original luminance value is retained. In difficult cases, a few swatches can be used to aid the matching process between the source and the target image. After color is transferred between the source and the target swatches, the final colors are assigned to each pixel in the greyscale image by matching each greyscale image pixel to a pixel in the target swatches using the L_2 distance metric. Thus, each pixel match is determined by matching it only to other pixels within the same image. We have found that this simple procedure works well for a wide range of image types. Further, these methods are easily extended to video. Here the colorization procedure is first applied to a single frame in the video sequence. The other frames in the scene are then assigned a color using the original frame’s colorized swatches. We have found that for scene frames in which the objects do not change dramatically, the colorization procedure works surprisingly well.

2 Color Transfer Algorithm

In this section, we describe the general algorithm for transferring color; the basic idea is then extended to use swatches. The general procedure for color transfer requires a few simple steps. First each image is converted into the $l\alpha\beta$ color space. We use jittered sampling to select a small subset of pixels in the color image as samples. Next, we go through each pixel in the greyscale image in scan-line order and select the best matching sample in the color image using neighborhood statistics. The best match is determined by using a weighted average of pixel luminance and the neighborhood statistics. The chromaticity values (α, β channels) of the best matching pixel are then transferred to the greyscale image to form the final image. Color transfer using swatches involves the same global image matching procedure but only between the source and target *swatches*. The colorized pixels in the target swatch regions are then used as the source pixels for the color transfer to the remaining non-colorized pixels using a texture synthesis approach. More specific details and justification are provided below.

2.1 Global Image Matching

Both color (source) and greyscale (target) RGB images are converted to the decorrelated $l\alpha\beta$ space [Ruderman et al. 1998] for subsequent analysis. $l\alpha\beta$ space was developed to minimize correlation between the three coordinate axes of the color space. The color space provides three decorrelated, principal channels corresponding to an achromatic luminance channel (l) and two chromatic channels α and β , which roughly correspond to yellow-blue and red-green opponent channels. Thus, changes made in one color

channel should minimally affect values in the other channels. The reason the $l\alpha\beta$ color space is selected in the current procedure is because it provides a decorrelated achromatic channel for color images. This allows us to selectively transfer the chromatic α and β channels from the color image to the greyscale image without cross-channel artifacts. The transformation procedure follows directly from Reinhard et al. [2001].

In order to transfer chromaticity values from the source to the target, each pixel in the greyscale image must be matched to a pixel in the color image. The comparison is based on the luminance value and neighborhood statistics of that pixel. The luminance value is determined by the l channel in $l\alpha\beta$ space. In order to account for global differences in luminance between the two images we perform luminance remapping [Hertzmann et al. 2001] to linearly shift and scale the luminance histogram of the source image to fit the histogram of the target image. This helps create a better correspondence in the luminance range between the two images but does not alter the luminance values of the target image.

The neighborhood statistics are precomputed over the image and consist of the standard deviation of the luminance values of the pixel neighborhood. We have found that a neighborhood size of 5x5 pixels works well for most images. For some problematic images we use a larger neighborhood size.

Since most of the visually significant variation between pixel values is attributed to luminance differences, we can limit the number of samples we use as source color pixels and still obtain a significant range of color variation in the image. This allows us to reduce the number of comparisons made for each pixel in the greyscale image and decrease computation time. We have found that approximately 200 samples taken on a randomly jittered grid is sufficient. Then for each pixel in the greyscale image in scan-line order the best matching color sample is selected based on the weighted average of luminance (50%) and standard deviation (50%). We have also included the neighborhood mean and varied the ratio of these weights but have not found significant differences in the results. Once the best matching pixel is found, the α and β chromaticity values are transferred to the target pixel while the original luminance value is retained.

This automatic, global matching procedure works reasonably well on images when corresponding color regions between the two images also correspond in luminance values. However, regions in the target image which do not have a close luminance value to an appropriate structure in the source image will not appear correct. Figure 2c shows the results of transferring color using the global image matching procedure. Here, the sky and trees match reasonably well between the images, but the road in the target does not match to the road in the source.




Figure 2: The two variations of the algorithm. (a) Source color image. (b) Result of basic, global algorithm applied (no swatches). (c) Greyscale image with swatch colors transferred from Figure 2a. (d) Result using swatches.

2.2 Swatches

In order to allow more user interaction in the color transfer procedure and to improve results, swatches are used between corresponding regions in the two images. Figures 2a,b,d demonstrate the basic idea. The first step is to use the general procedure described above to transfer color, but now *only* between the corresponding swatches. This allows the user to selectively transfer colors between the source and target swatches. We also expect the results to be good for individual swatches because there should be less overlap of luminance levels between different color regions within the same swatch. We perform luminance remapping as in the global procedure but only between corresponding swatches. Again, we use random jittered sampling with approximately 50 samples per swatch. The second step is similar to texture synthesis algorithms [Efros and Leung 1999; Efros and Freeman 2001] in which the L_2 distance is used to find texture matches. We define the error distance E using the L_2 metric between neighborhood N_g in the greyscale image and neighborhood N_s in the colorized swatch as:

$$E(N_g, N_s) = \sum_{p \in N} [I(p) - S(p)]^2$$

where I is the greyscale image, S is the luminance channel of the colorized swatch and p are the pixels in these neighborhoods.

Note, at this stage we no longer search the color image for texture matches but only search for matches within the colorized swatches in the target image. The advantage of the approach is that in the first stage we transfer colors to the swatches selectively which prevents pixels with similar neighborhood statistics but from the wrong part of the image from corrupting the target swatch colors. It also allows the user to transfer colors from any part of image to a select region even if the two corresponding regions vary largely from one another in texture and luminance levels. Secondly, since we expect there to be more texture coherence *within* an image than *between* two different images, we expect pixels which are similar in texture to the colorized target swatches to be colorized similarly.

2.3 Video

Colorization of video can be automated using the colorization procedure described above. To colorize all of the frames in a scene, we first transfer color from a source color image to a single target frame. Then every frame in the video sequence can be colorized using the same colorized target swatches used in the single frame. If a single frame is successfully colorized using these procedures, then frames which consist of the same objects as that single frame will be colorized similarly. Three sample clips are provided on the video proceedings.

3. Results

Figure 3 showcases the final results of the algorithm applied to a variety of image domains. Figures 3a-c shows the results of colorizing foliage, face and landscape photographs. The technique works well on scenes where the image is divided into distinct luminance clusters or where each of the regions has distinct textures. In general, however, the current technique does not work very well with faces. Although colors are transferred well into the swatches, the L_2 distance is not always a sufficient measure in classifying the difference between skin and lips and sometimes clothes and hair.

Figures 3d-f demonstrate the use of the algorithm with different types of scientific data. Although this technique is not intended for clinical diagnosis with medical image data, it might be used in anatomical illustrations or to enhance scientific presentations.

Although we show that the algorithm works well in a number of image domains, we do not claim that the technique will work on

most images. It should be clear that when one considers only a small neighborhood size around a pixel it is often impossible to determine whether that neighborhood belongs to one texture or another. However, by using high resolution images and larger neighborhoods we can obtain improved results. Further, we believe that more images can be colorized using the basic method provided but with better texture classification methods at the expense of simplicity and computation time.

The running time of the algorithm for one image can range from 15 seconds to 4 minutes on a Pentium III 900 Mhz CPU using optimized MATLAB code. Running time will vary depending on the number of samples used for comparison, the number of swatches, neighborhood size and the size of the images. Most images can be colorized reasonably well in under a minute.

4. Conclusions

In this paper we have formulated a new, general, fast, and user-friendly approach to the problem of colorizing greyscale images. While standard methods accomplish this task by assigning pixel colors via a global color palette, our technique empowers the user to first select a suitable color image and then transfer the color mood of this image to the greylevel image at hand. We have intentionally kept the basic technique simple and general by not requiring registration between the images or incorporating spatial information. Our technique can be made applicable to a larger class of images by adding a small amount of user guidance. In this mode, the user first transfers the desired color moods from a set of specified swatch regions in the color image to a set of corresponding swatch regions in the greyscale image. Then, in the second and final stage of the colorization process, the colorized swatches are employed, using a texture synthesis-like method, to colorize the remaining pixels in the greyscale image. Currently, the L_2 distance is used to measure texture similarity within the image. In the future we believe the technique can be substantially improved by using a more sophisticated measure of texture similarity.

Our technique of employing an example color image to colorize a greylevel image is particularly attractive in light of the growing sophistication of internet image search engines and the emergence of centralized and indexable image collections which can be used to easily locate suitable color images. Finally, one could also utilize a database of basis texture swatches for the initial color transfer in the user-guided stage of the colorization process.

Acknowledgements

This work was supported by NSF Career Grant ACI-0093157.

References

- EFROS, A. A. AND FREEMAN, W. T., 2001. Image Quilting for Texture Synthesis and Transfer. In *Proceedings of ACM SIGGRAPH 2001*, 341-346.
- EFROS, A. A. AND LEUNG, T. K., 1999. Texture Synthesis by Non-parametric sampling. *7th IEEE International Conference on Computer Vision*, 1033-1038.
- GONZALEZ, R. C. AND WINTZ, P., 1987. *Digital Image Processing*. Addison-Wesley Publishing, Reading MA.
- HERTZMANN, A., JACOBS, C., OLIVER, N., CURLESS, B., SALESIN, D., 2001. Image Analogies. In *Proceedings of ACM SIGGRAPH 2002*, 341-346.
- SILBERG, J., 1998. Cinesite Press Article, http://www.cinesite.com/core/press/articles/1998/10_00_98-team.html
- IOANNIS, P., 1991. *Digital Image Processing Algorithms*, Prentis Hall, 58.
- PRATT, W. K. 1991. *Digital Image Processing*, John Wiley & Sons, 311.
- REINHARD, E. ASHIKHMIN, M., GOOCH B. AND SHIRLEY, P., 2001. Color Transfer between Images, *IEEE Computer Graphics and Applications*, September/October 2001, 34-40.
- ROGOWITZ, B. E. AND KALVIN, A. D., 2001. The "Which Blair Project": A Quick Visual Method for Evaluating Perceptual Color Maps, *Proceedings of IEEE Visualization 2001*.
- RUDERMAN, D. L., CRONIN, T. W. AND CHIAO, C. C., 1998. Statistics of Cone Responses to Natural Images: Implications for Visual Coding, *J. Optical Soc. of America*, vol 15, no. 8, 2036-2045.


(a) Foliage image mapped using a swatch for the bark and the leaves. Source image courtesy of Adam Superchi. Target image courtesy of <http://philip.greenspun.com>.




(b) The results of colorizing a photograph of a face. Four swatches were used for the hair, skin, shirt and background with a 11x11 neighborhood size for the L_2 metric. Images courtesy of Lela.




(c) The results of colorizing an Ansel Adam's photograph. A total of 3 swatches were used. Source image courtesy of Paul Kienitz.



(d) A Landsat 7 satellite image (converted to greyscale) was colorized with another Landsat satellite image using the global matching procedure.



(e) One slice of a colorized MRI volume. We used a color cryosection from the Visible Human Project dataset and two swatches. By using swatches, we could avoid transferring the color of the blue gel.



(f) A scanning electron microscopy (SEM) image colorized with a photograph of an ant using the global matching procedure.



280 Figure 3: The results of applying the algorithm to photographs and scientific datasets. The first column contains the original color image, the second column contains the target greyscale image and the final column is the result of the colorization procedure. Note, the images have been scaled to fit the page. Swatches have been included as colored rectangles on the source and target images.