

Natural Language Color Editing

Geoff Woolfe

Xerox Innovation Group

Webster, New York

Introduction

There are many ways to specify color and color difference. Color imaging scientists and engineers describe color using precise, numeric color specifications. Such color specifications are often based on the color matching behavior of a standard human observer and are referred to as device-independent colorimetric specifications. Commonly used colorimetric specifications include CIE XYZ, and the more perceptually uniform, CIE Lab system. In the world of digital printing, color is also frequently described in terms of the device-dependent control values needed to generate a color on a specific device. Colors can also be specified using color order systems such as the Munsell Book of Color, the Swedish Natural Color System, or the Pantone Color Formula Guide. These types of color specifications also provide a precise color specification, but are more commonly used by professionals in the color graphics and design industries rather than color imaging.

Another ubiquitous form of specifying color is to use color names in natural language. Although this is a far less precise method of color specification than those discussed above, it is nonetheless the most widespread and best understood method of color specification used by all consumers of color. This method of color specification uses common color names, such as red, green, blue, etc. It also uses combinations of common color names to refine the specification. Examples of such combinations include reddish-brown, greenish-blue, yellowish-green etc. In addition, natural language provides many modifying adjectives to provide further subtle discrimination in color specification. Examples of such modifying adjectives include light, dark, bright, saturated, vivid, muddy, moderate, dull, pale, washed-out etc.

In addition to specifying colors, natural language is also commonly used for specifying color differences or color changes. Natural language examples of color difference words and phrases include “slightly less yellow”, “much darker”, “more saturated”, “greener”, “significantly punchier” and “a smidge lighter”. Now, while these expressions are certainly imprecise, many people commonly use them to describe the changes required during proofing of their print jobs.

Color control and adjustment systems generally require the user to develop an understanding of the behavior of the various controls provided in the user interface. Users must determine the appropriate adjustment tool and the correct settings for that tool in order to achieve their desired adjustment. Many people, lacking specialized training, therefore find such systems difficult to use. Accordingly, most consumers of color images and documents find it difficult to adjust the colors in images or documents. It then becomes the job of the graphics professional or printer to translate the spoken or written color requirements of the customer into appropriate settings in the control tools of image editing or device control applications. Color graphics professionals require extensive training and experience to successfully and efficiently manipulate controls in such applications to achieve an aesthetic effect that can be stated simply and concisely in

verbal terms. It is advantageous therefore to provide a natural language interface for color adjustment and image processing applications to address this usability issue.

While, verbal descriptions of color and color difference are less precise than the numerical specification of color spaces, it can nonetheless be argued that a less precise, but better understood communication system is still preferable to a highly precise, but difficult-to-use, interface.

Developing a mapping between natural language color specifications and the precise numerical color encodings used in color image processing and device control applications is complicated by a number of issues. First, there is no uniquely defined natural color language. The words and grammar used to describe color can vary based on culture, geographical location, professional affiliation, and individual preference. Second, the boundaries between named colors are not precisely defined – indeed; they are somewhat fuzzy and can vary, to some extent, between individuals. Third, the lower level of precision afforded by a natural language interface may be perfectly adequate in some applications but be unsuitable for others. This paper will illustrate one approach to developing a natural language interface to a color editing application and will demonstrate its use.

Discussion

Verbally specified color adjustment instructions can generally be parsed into two pieces of information - the target range of colors that a user wants to adjust and the color modification to be implemented. The parsing of several examples of natural language color adjustment instructions is shown in table 1.

Table 1

Natural Language Color Specification Phrase	Target Color Range	Color Modification to be Implemented
Make the image moderately lighter	All colors	Moderately lighter
Make the blues slightly less purple	Blue colors	Slightly less purple
Make the red and purple regions much less contrasty	Red and purple colors	Much less contrast
Make the yellows but not greenish-yellow a little more orange	Yellow, but not greenish-yellow colors	A little more orange

The target color range information can be used to construct an image mask that defines the regions of the image in which the color modification is to be applied. Such a mask is illustrated in figure 1. The result of applying a natural language color adjustment to the image of figure 1 is shown in Figure 2.

In order to construct the image mask one needs to determine, for each pixel in an image, whether it falls in the target color range. Thus each color name that might be used in the target color range specification must be associated with a region or sub-volume in the color space of the image. There are several ways that this partitioning problem can be solved.

In the prototype natural language applications we have developed to date, this has been accomplished by defining a standardized dictionary of color names and associating, with each of those names, a prototypical location in color space. The color space used for these prototypical locations is the CIE Lab color space. This space has the advantage of

being widely understood, device-independent and approximately perceptually uniform. Any other color space could be used, but perceptually uniform, device independent spaces are generally preferred. Device independence is important because it allows the algorithm to be easily implemented without regard to the specific printing or display hardware on which the colors will eventually be rendered. Perceptual uniformity is valuable because it greatly simplifies the task of partitioning the color space into named regions that match the human perception of color.




Figure 1: Original image (above) and mask generated by the color name 'orange' (below).

The standard dictionary of color names used in this application is based on the NBS-ISCC standard color name dictionary¹⁻³. We have supplemented this dictionary with several additional names for neutral gray colors. The NBS-ISCC dictionary consists of 267 color names. The names are structured and composed of between 1 and 3 words. At the most fundamental level there are simple one-word color names such as red, green, yellow, blue, black etc. At the next level there are hue modifiers attached to these fundamental color names. This gives colors such as yellowish-green, reddish-blue, purplish-red etc. At the next level are modifiers that distinguish levels of lightness and colorfulness. These modifiers include words like dark, light, pale, vivid, deep etc.



Figure 2: Image after application of the transform 'very-strongly more saturated'.

Having defined a color name dictionary with its associated prototypical locations in a color space, a means to partition the color space into regions assigned to the various names of the color name dictionary is required. A variety of methods may be used to define these regions. The only requirement is that every location in the color space be assigned to at least one region. The regions may be disjoint or overlapping. There is no restriction on the topological genus of a region and a region may consist of multiple, non-connected sub-regions.

A computationally and experimentally efficient means to partition the color space is to compute the Voronoi partition^{4,5} of the set of prototypical locations in the color space. This approach computes a convex Voronoi cell around each prototypical color space

location such that all colors inside the cell are closer to the prototypical location than to any other prototypical location. An illustration of a Voronoi partition of a set of points in two dimensions is shown in figure 3 below.

Another method, mathematically distinct but functionally equivalent to Voronoi partitioning, is to assign the prototypical color locations to the nodes of a space-partitioning data structure such as a kd-tree^{6,7}. Computationally efficient nearest-neighbor searches can then be performed on the kd-tree to determine the closest named color for any arbitrary input color.

A further means of assigning regions of color space to the terms in a color dictionary include using a statistical analysis of the results of psychophysical color naming studies. This method would work to directly map color terms with regions of color space or could be used to define the regions around prototypical locations in color space. This method has the drawback of being experimentally difficult and very time consuming.




Figure 3. Example of a Voronoi partition of a set of points in two dimensions.

One problem associated with using a standard color name dictionary is that it would require users to be familiar with, and use, this dictionary. This is an unreasonable expectation and hence we have developed methods by which much larger dictionaries, consisting of commonly used color names, can be mapped onto the standard dictionary. Currently the common color name dictionary consists of almost 2000 color names in common (and not so common) usage. If users make entries that are not included in this larger common name dictionary, then the nearest matching entry is chosen using a minimum string edit distance algorithm.⁸

Another way of selecting the colors the user wishes to change is to select the region of the image that contains these colors. There are tools available that let the user select a point in the image, and set the range of colors in the vicinity of the selected color that is to be modified, but most users find it difficult to set the range.

The second part of a natural language color editing specification is the color modification instruction. The color modification instruction can be used to create a color transform that is applied to the desired colors. Three pieces of information must be determined from the color modification instruction. These are the color property to be modified, the direction of the modification and the magnitude or extent of the modification. When, however, the color change is large, and the selected range of colors is small, unpleasant color reversal artifacts can be produced. We have therefore developed methods of increasing the range of the colors that will be modified to maintain smoothness and monotonicity, thereby avoiding this problem.

The creation of a color transform from the verbal color change specification involves mapping the change direction and magnitude onto numerical scales and then selecting

from a family of pre-defined transform functional forms based on the color property being modified. The numerical values corresponding to the color change magnitude and direction are used as parameters in the transform function. It is necessary however to condition these parameter values based on the actual colors in the image in order to avoid the color reversal artifacts mentioned above.

Having created both an image mask, based on the target color range, and a color modification transform, based on the verbal color change specification, it is relatively simple to apply the transform to the image, modulated by the mask. The verbal presentation of this paper will discuss the details of how mappings between verbal descriptions and numerical color spaces are achieved and demonstrate the results that can be obtained using such a system.

References

1. K. L. Kelly und D. B. Judd, The ISCC-NBS Method of Designating Colors and a Dictionary of Color Names, National Bureau of Standards (USA), 553 (1955)
2. K. L. Kelly und D. B. Judd, Color Universal Color Language and Dictionary of Names, National Bureau of Standards, Special. Publication. 440 (1976)
3. F. T. Simon, "Color Order", in F. Grum und C. J. Bartleson, Optical Radiation Measurements, Volume 2, Color Measurement, New York 1980, pp. 165-235.
4. Atsuyuki Okabe, Barry Boots, Kokichi Sugihara & Sung Nok Chiu (2000). Spatial Tessellations - Concepts and Applications of Voronoi Diagrams. 2nd edition. John Wiley, 2000, 671 pages ISBN 0-471-98635-6
5. Franz Aurenhammer (1991). Voronoi Diagrams - A Survey of a Fundamental Geometric Data Structure. ACM Computing Surveys, 23(3):345-405, 1991
6. Bentley, J. L. 1975. Multidimensional binary search trees used for associative searching. Commun. ACM 18, 9 (Sep. 1975), 509-517.
7. H. Samet, The Design and Analysis of Spatial Data Structures, Addison-Wesley, Reading, MA, 1990.
8. H. Bunke, and J. Csirik, "Parametric String Edit Distance and Its Application to Pattern Recognition," IEEE Trans. Systems, Man, and Cybernetics, vol. 25, pp. 202-206, 1995.

Author Biography

Dr. Geoff Woolfe is currently a Principal Research Scientist in the Xerox Innovation Group. He joined Xerox in 2005 after spending a 21-year career in the Kodak Research Laboratories.

Dr. Woolfe received his B.Sc. (Honors) and Ph.D. degrees in Physical Chemistry from the University of Melbourne (Australia) and the M.S. degree in Imaging Science from the Rochester Institute of Technology. He is a member of the Honor Society of Phi Kappa Phi and a member of the Society for Imaging Science and Technology.

Dr. Woolfe is the author of 23 scientific papers and 38 US and international patents and pending patent applications. His current research interests include development of color imaging algorithms, color management, color application usability and color imaging system modeling.