Using Color In Computer Applications: A Psychophysical Perspective

Bernice E. Rogowitz
IBM T.J. Watson Research Center, Yorktown Heights, New York

Deciding how to use color in computer applications is a very difficult problem. Some of the difficulty resides in the characterizing the output medium. Color production systems vary in their underlying color models, their intrinsic color gamuts, and even the angular distribution of the light stimulus reaching the human observer. Some of the difficulty resides in understanding the human perceptual system which is receiving this input. Perhaps most elusive of the difficulties, however, is how the interaction between the physical stimulus and the observer depends on the task the user is trying to perform.

The Physical Stimulus

The first stage of the human visual system is designed to detect and process photons of light. At the retina, light absorbed by the receptor outer segments engages a neural response which becomes part of a complex network of neural signals throughout the brain. At this initial stage, however, the first response is determined by the spectral composition of the photon which has been absorbed, independent of whether it has been reflected from a piece of film or emitted from a CRT display.

Surprisingly, very few studies begin with a complete analysis of the light distribution reaching the eye, or to be more complete, with an analysis of the light distribution actually reaching the retina (McCann, 1995). Some studies even begin with the assumption that the stimulus is an r,g,b value or a dye's spectral reflectance distribution, ignoring the fact that the r,g,b value is often just a set of bit values for the monitor's d/a converters and the dye's spectral reflectance doesn't take paper properties into account. In both cases, the "color" property is expressed in terms which device manufacturers specify, not in terms which are relevant to perception.

Colorimetric measurements have provided a convenient way of characterizing the physical stimulus in perceptual terms. In colorimetry, the perceptual effect of the physical stimulus is expressed in terms of the degree to which three color filters are stimulated, where a linear transformation of these filter functions has been derived from human color-matching experiments.

Colorimetry provides a tremendous shorthand for expressing color. Based on the trichromatic properties of human vision, it lets us express the stimulus in terms of three values instead of requiring us to specify the complete spectral excitation function. In the family of colorimetric domains, standardized by the CIE in 1931 and in various other forms since then, two patches of stimulus which have the same colorimetric value, when viewed under the same illuminant, will appear to have the same color. Of course, what that color "looks like" depends on the size of the patches, the color of the areas surrounding them, what the observer was looking at just before making the observation, etc.

The Visual System

This list of caveats immediately shows us that the perception of color is not based on a point-by-point analysis of the physical stimulus. A spot of light made up of a fixed spectral distribution and having a fixed chromaticity, for example, may look "white" when it subtends a few minutes of arc, but "blue" when it subtends several degrees. And that same spot of light may look "blue/green" if surrounded by a red area, and "black" if surrounded by a bright surround.

These effects suggest that the appearance of a colored stimulus depends not only on the chromaticity at a point, but on the spatial environment in which it lies. The role of spatial context has a long history (see, for example, Hunt (1957), Judd (1960), and Land (1959)), and there are still many unresolved questions.

These perceptual effects are joined by other higher-level effects. For example, the "blue" stimulus will appear more saturated to the observer if it has a shape which is associated with the color blue, a flower or a bird, than if it has a neutral shape or a shape associated with another color, say a tree. These are examples of cases where the chromaticity measured at a point fail to predict color appearance at that point.

The inverse can also be true. Two patches of stimulus can have very different chromaticities yet appear to be identical. When color constancy holds, for example, a stimulus viewed under a bathing illumination will maintain its apparent color, independent of the chromaticity of the illuminant. A blue patch, under a broad-field pink light, thus, would look "blue" not "purple," even though the light reflected to the eye provides significantly more long wavelength excitation.

So, perceived color depends not only on the excitation of the cone filters, but also on the spatial and temporal distribution of that excitation. But, even knowing the spatial, temporal and chromatic characteristics of the visual scene and the spatial, temporal, and chromatic response of early visual mechanisms does not provide a complete model of perception.

This is just the first step in the network. These retinal and cortical responses feed to areas of the brain where higher-level cognitive functions, such as memory, attention and language are mediated, and to mechanisms involved in emotion and aesthetics. Furthermore, this is not simply a bottom-up network, where all the information flows in one direction. Higher-level evaluations, such as the semantic meaning of an object, can affect lower level processes, as in the case we discussed where the meaning of an object affected the perception of its color.

Color Vision and Tasks in Computer Applications

When humans look at a visual display, all the stages of the perceptual system are engaged. Which perceptual mechanisms predominate, however, depends on the task the observer is trying to achieve. This section examines several computer applications with respect to the perceptual demands of different tasks.

Image Reproduction

The major work in applying principles of color vision and colorimetry to computer applications has been in the area of image reproduction. This work has laid a foundation for the applied problem currently called "device independent color," the problem of creating a visual representation which appears the same whether printed on a laser printed, viewed on a CRT display, or developed photographically.

In some cases, we have the luxury of being able to create representations which are spectrally identical on a point-by-point basis. Sometimes we have the luxury of being about to create representations which are colorimetrically identical on a point-by-point basis. In most cases, however, the physical process limits our ability to reproduce the original or create representations which are identical. The color gamut of the output devices, for example, may not match each other, or match the color gamut of the original image. For example, there may be colors in the original image for which there is no equivalent color in the output display's gamut. In order to make sure that each pixel in the output image has a value, certain compromises must be made.

How we solve this problem depends on our model of what the human observer is doing to evaluate the differences between the original and the representation, or between the representations. If we think that the observer is evaluating the two images on a pixel-by-pixel basis, we may preserve all values where the gamuts overlap and select the closest possible value in the output gamut for every value in the original image which cannot be directly represented. Perhaps a saturated purple in the original image can be replaced by a saturated dark blue. If we think that the important characteristic to preserve is the relationship of all the colors to the white point, we can pin the white points of the two gamuts and remap every point in the original gamut to a new point in the output gamut. If we think a simple transformation of all the values, simulating, for example, the effects of color constancy, will cause the two to look equivalent, we might try to find a transformation which simulates the effect of a chromatic illuminant. If we think that preserving the color names of all the major regions will cause the two to look equivalent, we might segment the image into regions and map the colors according to a color-naming scheme.

The success of the different algorithms, however, depends inevitably on the task the user is trying to accomplish with the image. If the images are photographs of human faces, for example, an algorithm which sacrifices hue in order to preserve saturation and luminance may not be acceptable. If the images are color-coded pseudocolor maps and the color purple has a special semantic meaning to the user, an algorithm which maps all purple regions to blue will not be acceptable, no matter how good the ΔE score is.

Use of Color in Sampled and Compressed Images

In digital technology, images are often represented on devices which have lower spatial, temporal or chromatic resolution than desirable. To facilitate storage and communication, images are often compressed, often to the point where the uncompressed image does not preserve all the visual information in the original image. In both these cases, the image the user sees is visually different from the ideal image which it represents. One strategy for improving visual quality has been to use information about the spatial, temporal and chromatic properties of human perception to develop sampling and compression algorithms. A sampling algorithm, for example, may "hide" alias frequencies in spatial/temporal/chromatic regions with low visual sensitivity; a compression algorithm might take advantage of differential visual bandwidths for spatial variations in color and luminance information.

Color in the Design of Computer Applications

As access to the technology for creating, processing and using color images becomes more prevalent, new color applications arise. Desktop publishing applications, for example, allow users to build color presentations, newsletters, and artwork. These systems give the users interactive control over colors, fonts and formatting options, but provide no guidance in the visual aspects of color, except, perhaps, in the form of template examples. Generally, a user is free to select illegible, garish, or distracting color combinations. One line of research we've been pursuing is the development of intelligent tools to help users select colors and fonts which are legible and color combinations which are aesthetic. We do this by constraining the user's choice of color, font and graphical objects. These constraints are built on rules based on perception, cognition and color theory (Rogowitz and Rabenhorst, 1993). For example, if the user has selected a dark background color, the set of selections offered for text will be limited to those colors which will provide sufficient luminance contrast. The amount of luminance contrast required, however, also depends on the spatial resolution of the font, and so the system also incorporates information about the spatial frequency in delimiting the color choices. By building a framework which allows us to constrain the choices a user has based on perceptual rules, we can systematically explore new types of questions. For example, if we use our knowledge of spatial/color interactions in human perception to constrain the colors and fonts the user can select, does this result in a better, faster, or more effective application? Can we extend this work beyond the body of established psychophysical relationships? For example, are there simple rules which can describe the aesthetic relationship between sets of colors? Can these rules be operationally defined in such a way that they actually assist users in selecting aesthetically-pleasing color combinations? Can these rules be parameterized to reflect different tastes, different cultural sensibilities, different moods? Is it possible to develop rules which help select colors which communicate different emotional characteristics? Addressing these questions have taken us into the realms of visual psychophysics, graphic design, font design, color aesthetics, and emotion.

They also lead us deeper into the technology of color representation. Some unanswered questions include: how well does a display monitor need to be calibrated in order to be able to enable a rule-based system which helps users select legible color combinations? - aesthetically-pleasing color combinations? - color combinations which will be discriminable to color deficient users? How much does it matter how the colors are represented to the user? How much practical difference is there between the various perceptual color spaces? Do different color spaces provide better media than others depending on user's task? By embodying our hypotheses into a testbed, we can test how good and how robust they are, and how competing or complementary relationships interact.

Color in Visualization

Another emerging application for color is in the area of visualization, where data from satellites, medical sensors, demographic studies, etc, are represented visually. In visualization systems, physical or abstract dimensions are represented as variations in perceptual dimensions. In modern visualization systems, the user dynamically crafts the visual representation of the data, making decisions about how many variables to represent and how to represent them. These representations are often very complex, with colored planes intersecting colored three-dimensional objects.

Sophisticated tools assist the user in creating visual representations of data. Despite their sophistication, however, these systems provide only minimal color support. For example, the most common color tool is a "default" color map which maps scalar data onto a scalar variable representing a linear interpolation between blue (0,0,255) and red (255,0,0). Mathematically, this represents the mapping of scalar data values onto a scalar visualization variable. Perceptually, however, this scale is not at all continuous. It looks like bands of colors from blue to red in rainbow sequence, varying in hue, brightness and saturation. When mapped onto scalar data, this colormap can give the user the erroneous impression that the data are organized into discrete regions, each represented by one of the rainbow colors. This can lead the user to infer structure which is not present in the data.

The question of visual representation of data has a long history (see reviews by Bertin, 1967; Tufte, 1987), but enormous opportunities still exist for understanding how to match the visualizaton to the structure of the data and to the perceptual processing of the user.

Furthermore, visualization is used to support a wide range of tasks which have very different requirements and place very different demands on human perception and cognition. In some cases, the user is browsing, manipulating and exploring the data, looking for particular features or relationships in the data. In some analysis tasks, it is important for the visual representation to be as faithful to the underlying structure as possible, as, for example, in the analysis of a medical MRI image.

In some analysis tasks, the analyst relies on the visual representation to highlight certain ranges of the data, or to segment the regions into semantically-namable areas. For analysis, however, it is important that mapping of the data values to the perceived value, height, size and color of the representation be well-controlled.

In other cases, the data has already been analyzed, relationships and features identified, and now the task is to communicate that message. In this case, the best representation may highlight or draw attention to a particular feature or relationship, or exaggerate some aspect of the data in order to make the best presentation. While analysis tasks require perceptually faithful representations of the data, presentation tasks may require distortions in order to be effective.

In our work, we have begun to develop a taxonomy of visualization tasks, a taxonomy of visual representations, and a framework linking tasks to the representations (Rogowitz, Ling and Kellogg, 1992; Rogowitz and Trienish, 1993a, 1993b, 1994). In particular, we've been developing an intelligent system which will guide the user in selecting an appropriate visual representation. One area we've been exploring is the selection of appropriate color maps, since, as we discussed earlier, the "default" colormap is so fraught with problems.

For tasks such as medical imaging, where the analysis of of the image depends on the faithful representation of the data, we are developing color maps where changes in the magnitude of the variable being represented are perceived as equivalent changes in the representation. For tasks requiring the representation of semantically-distinct regions, we are developing color scales whose regions look distinct based on hue, and may also connote magnitude through manipulations of luminance. For tasks which require that certain characteristics or regions be highlighted, we are developing color scales which attract attention. The goal of this research is to understand the perceptual requirements of different tasks and to develop visual representations which support them.

The development of these color maps has drawn heavily on psychophysical and cognitive literature relating to color. For example, although monotonic variations in luminance (grayscale) and saturation both create monotonic variations in perceived magnitude (Stevens, 1951), the high spatial-frequency sensitivity of the luminance channel in human vision suggests that grayscale would provide the better representation of high-spatial frequency data. We find, that, in fact, a continuous color map which includes a strong luminance component provides a more faithful perceptual representation of high spatial-frequency data.

This work has also raised a number of interesting questions. For example, how do different perceptual scales interact in guiding perception? Does a color map which varies both in luminance and in saturation provide a greater perceptual dynamic range than a scale of a single dimension? Is is possible to construct a color map which effectively represents complex images containing both low and high spatial frequency information? How is the number of semantically-definable regions limited by color discrimination, color memory, and spatial resolution? How can multiple meanings be effectively layered onto the same data? How do complex spatial objects differ from simple targets in the way they "pop" into attention (Treisman and Gelade, 1983).

Conclusions

Computer systems are evolving. Data and color-manipulation functions only available on high-end systems a few years ago are now available on desktop workstations. Color representations are being used in an increasingly wider range of applications, fueled by advances in electronic photography, digital image compression, and portable storage media. Most of the color applications available today, however, have been developed without a fundamental knowledge of human color perception, which means that many applications involving color are frustrating to users. They create presentation graphics which are illegible, they inadvertently introduce visual artifacts into scientific visualizations, they create color images which sometimes appear with mysterious colors when printed. These problems, however, define the next generation of problems in color research. They ask us to broaden our perspectives beyond the physical measurement of color, beyond the colorimetric representation of color, beyond the spatial, temporal and chromatic filters of early vision. These new applications ask us to broaden our study of color, placing greater emphasis on higher level mechanisms of color perception, such as image and color semantics, color memory, aesthetics, attention and emotion. These areas will fuel the development of future systems.

References

1. Bertin, J., *Semiology of Graphics*, 1967, reprinted by University of Wisconsin Press, Madison, WI, 1983.
2. Hunt, R.W.G., "*The Reproduction of Colour in Photography, Printing and Television*," London: Bath Press, 1987, first edition, 1957.
3. Judd, D. B., "Appraisal of Land's Work on Two-Primary Color Projections" *Journal of the Optical Society of America*, **50**, 254, 1960.
4. Land, E. J., "Color in the Natural Images, Parts I and II," *Proceedings of the National Academy of Sciences,* **45**, 115-129; 636-634, 1959.
5. McCann, J. J., (1995) "Subtracting Stimulus Physics from Psychophysical Conclusions," *Proceedings of the SPIE/ IS&T Conference on Human Vision, Visual Processing and Digital Display*, San Jose, February, 1995, in press.
6. Rogowitz, B. E., Ling, D. T. and Kellogg, W. A., Task Dependence, Veridicality, and Pre-Attentive Vision: Taking Advantage of Perceptually-Rich Computer Environments. *Proceedings of the SPIE,* Volume **1666***, Human Vision, Visual Processing and Digital Display*, pp. 504-513, 1992.
7. Rogowitz, B.E. and Rabenhorst, D. A., "CRAFT: A Tool for Customizing Color and Ront Selections Guided by Perceptual Rules," *IS&T/SID Color Imaging Conference*, Scottsdale, AZ, 1993.
8. Rogowitz, B. E., and Treinish, L.A., "Data Structures and Perceptual Structures," *Proceedings of the SPIE Conference on Human Vision, Visual Processing and Digital Display*, Volume **1913**, 600-612, 1993.
9. Rogowitz, B. E., and Treinish, L.A., "An Architecture for Rule-Based Visualization," *Proceedings of the IEEE Conference on Visualization '93*, San Jose, California, October 25-29, 1993, 236-243.
10. Rogowitz, B. E., and Treinish, L.A., "Using Perceptual Rules in Interactive Visualization," *Proceedings of the SPIE, Human Vision Visual Processing the Digital Display,* Volume **2179**, 1994.
11. Senay, H. and Ignatius, E., "Compositional Analysis and Synthesis of Scientific Data Visualization Techniques", TR-90-16, *Center for Excellence in Space Data and Information Sciences*, 1990 and SIGGRAPH '90 Course State-of-Art in Scientific Visualization, August 1990, Dallas, TX.
12. Stevens, S. S., "Matching Functions Between Loudness and Ten Other Continua," *Perception and Psychophysics*, **1**, pp. 5-8, 1966.
13. Treisman, A. and Gelade, G., (1980) "A Feature Integration Theory of Attention," *Cognitive Psychology*, **18**, 643-662.
14. Tufte, E., The Visual Display of Quantitative Information, Connecticut: *Graphics Press*, 1983.