Toward a Framework for Intuitive Perceptual Colormap Specification and Generation in Exploratory Data Visualization

Mark Wang

June, 1999

Stanford University

Contents

Segmenting

Isosurface Detection

Semiautomatic Generation (Current and future user Interfaces for choosing color)

Single-color specification mechanisms

PRAVDAColor

Dynamic Real-Time Interfaces: Calico

Our approach: A modular, semi-interactive approach: combining endpoint/boundary specification with path generation within the colorspace

Explorer module: randomly walk the tree of colorspaces

Methodology for Automatic Generation

Shape-based and spline-based methods in color space

Implementation: the RIVET and Mineset projects

Future Work

Feasiblity of automatic colormap generation?

The limits of color and colormaps in visualization

Abstract

In almost every contemporary computer visualization example, color is often used to add one or more dimensions of information expressiveness. However, unlike other graphical parameters, such as space or time, the interpretation by the human visual system of color is not objectively defined in a completely known manner, but has many perceptual pitfalls that can make visualizations less effective, and even mislead at worst. Therefore, the problem of choosing an optimal color scheme is an important one.

Furthermore, a given color scheme is not optimal for all visualization tasks, even for a fixed data set, as often times, we want to explore and highlight certain areas, and color provides a compelling tool to do this interactively.

Current color selection and specification methods in many existing visual tools usually fail to address both of the above issues accurately: color selection is manual and in almost all cases, left partially or completely up to the user, without guidance as to good perceptual principles, nor a easy way to intuitively *change* color mapping to facilitate exploring data in a visual manner.

In this report, we begin with a review of human color perception, and identify some well-known "pitfalls" with color perception. We then survey the issues surrounding color selection for information visualization, including issues of working in a color space, and various functional requirements for various visualization tasks and metatypes of data. Finally, we review existing color selection/specification interfaces and suggest a framework for exploring the space of color maps in information visualization, as an interactive experience, with assistance to maintain perceptual optimality in light of these varied guidelines.

Acknowledgements

The initial motivation for focusing on the color map generation question was the RIVET (Rapid Information Visualization EnvironmenT) project. Pat Hanrahan, Robert Bosch, Tamara Munzner, and Chris Stolte provided some early suggestions and showed some visualization examples to get me started. A significant portion of the research was conducted while working in Silicon Graphics in the Data Mining and Visualization Group. Thanks to Barry Becker for overseeing my work, and professor Roger Crawfis for providing suggestions on various visualization tasks that users might want to do. David Peck and David Pinkney also provided hints that encouraged me to look in new directions

Introduction

Over the years, color computer graphics have become a pervasive and integral part of the field of information visualization.

The field of visualization itself can be characterized by its intended function: *expository* visualization aims to communicate known information in a visual, intuitive manner, while *exploratory* visualization aims to assist the user of the visualization in actually discovering information as well as higher-level trends (such as correlations) of the underlying data. The use of visualization to portray known information has, in a sense, been around for centuries, since the first systematic and quantitatively-accurate graphs were produced in the late Middle Ages. However, exploratory visualization has largely been a recent field brought on by the advent of computer graphics systems capable enough to be able to process and manipulate visual representations of data in an acceptably interactive fashion. As data sets have gotten bigger and more unwieldy, the use of visualization not just as an ancillary presentation tool, but as a integral part of the data analysis and discovery process, has grown rapidly.

Many of the guidelines for generating good visualizations such as those put down by Tufte[X], as noted by Uselton[X], have centered around visualization of the expository kind. However, when the nature of the data itself is not known even to the producer of the visualization, it is important that the data being mapped into a visualization is perceptually accurate as possible, without any artifacts that might confuse or even mislead.

Color has been long one of the fundamental graphical quantities to which variables are mapped to in visual representations of data, along with other objective parameters such as spatial or temporal location. Yet, our perception and cognition of color is also quite complicated in comparison to these other elements. Firstly, because color is inherently a cognitive phenomenon, accurately *specifying* color in a quantitative manner has proven to be challenging. While much research has been done to address this issue by coming up with psycho-mathematical color spaces that are attuned to human perceptual processes, the *interaction* of colors with one another remains a thorny issue that must be considered in order to avoid misleading visual illusions. Furthermore, *user interfaces* for color specification, especially those that facilitate interactive exploration of the space of color mappings, are also crucial if color is to have its full potential to communicate information visually.

In this report, we seek to lay a foundation to begin to address these issues. We review the basics of human color perception and color spaces, and then examine the requirements of data visualization to provide a motivation for our color selection guidelines. We then propose a framework for not only choosing a set of suitable colormaps initially by the system, but also one in which the system will assist in *interactive* color selection for various visualization tasks that may arise as the user is exploring the data, such as segmentation or highlighting specific ranges. By combining a dynamic user interface with system-assisted selection of colormaps, our system extends previous work in dynamic user-adjustable colormaps by reducing guesswork and the need to manually select a suitable initial set of mappings to work with, freeing the user to concentrate on the exploration of the colormap and the data.

The architecture we suggest provides a flexible test bed for future work and experimentation with heuristics and other rules for color selection as well as user interfaces, all areas where much future work remains to be done.

Fundamental attributes of human color perception

For the rays [of light], to speak properly, are not colored. In them there is nothing else than a certain power and disposition to stir up a sensation of this or that color. For as sound, in a bell or musical string or other sounding body, is nothing but a trembling motion, and in the air nothing but that motion propagated from the object, … so colors in the object are nothing but a disposition to reflect this or that sort of ray more copiously than the rest.

-- Sir Isaac Newton, 1671

To understand the role of color in augmenting information presentation, and to motivate the guidelines that we outline for choosing effective colors for visualization, it is necessary to have a basic understanding of how the human visual system processes color, and the attributes that psychologically distinguish one color from another.

What *is* color, anyway? A high level scientific definition of color is that it is the perceptual result of perceiving various *spectral power distributions* (SPDs) of wavelengths of light that lie in the visible range. However, in the actual cognitive process of color recognition, the relationship between a given SPD and the resulting perceived color is not straightforward at all. We do not consciously think about anything related to the actual characteristics of the distributions themselves. Indeed, for a given perceived color, an infinite number of different SPDs (known as *metamers*) may give rise to it. Hence, when characterizing color, we do not think in terms of the actual distributions, but rather, several higher-level qualities, known as lightness, hue, and saturation, which perceptually distinguish various SPDs from one another as different colors.

Luminance and Lightness

Firstly, we perceive differing levels of light energy as being distinct which give rise to the notions of *luminance* and *lightness*. Luminance is a psychophysical measure of the absolute amount of light energy emitted by a light source per unit solid angle (known as *radiance*) multiplied by a response curve over the visible spectrum (see the section on CIELAB for further discussion) to compensate for the human eye's varying sensitivity for different wavelengths of light energy.

The human eye is most sensitive to light in the 555 nm range (roughly yellowish-green), which corresponds naturally to maximum perceived lightness for any given radiance level. Mathematically, the *spectral luminous relative efficiency curve*, which describes the human eye's relative sensitivity as a function of wavelength has its maximum at this point, and tapers off in a rough Gaussian-like fashion as one nears the extremes of the visible spectrum.

Human vision has a nonlinear perceptual response to luminance relative to a given reference value: a source having a luminance only 18% of a reference luminance appears about half as bright. The perceptual response to luminance is called *lightness* (sometimes referred to as *brightness*), and has "white" and "black" as the extremes of the scale. Lightness is invariant to the absolute power of a computer display, but luminance is directly dependent on it. In addition, for the eye to be able to discriminate a lightness difference between two adjacent, but not contiguous patches, a luminance difference of about 7% is required. Contiguous patches are perceived as different lightnesses when the difference is only 2%.

Empirical studies have shown that an observer can detect an intensity difference between two patches roughly proportional to the absolute brightness of the patches – i.e., the perception of lightness, given empirical data, is roughly logarithmic. This is known as the *Weber-Fechner law*.

In terms of colormap generation, a map that varies lightness should be perceptually continuous – that is, there should be no perceived line between two samples as the eye moves about the scale, and the rate of change in perceived lightness should be constant throughout the scale. In order to achieve this, we need an additional linearization step, described later in the section on color spaces, to generate a true perceptually based scale of lightness.

Hue and Ordering

Hue is defined as that quantity that distinguishes one color from another, as opposed to a different lightness or saturation level of the same color.

On a high level, the human cognitive system groups hues, both pure, and otherwise into various large regions in color space – colors on those regions are those that tend to be associated with a certain name, such as "blue," "green," or "red." Experiments have shown and verified the intuitive notion that color discrimination and search within a hue region takes significantly more time than searching for a color that lies outside of one. [Kawai95]

The space of hues perceivable by humans can be represented in several different ways. For our purposes of data visualization, all these schemes convey a sense of *ordering* to the user, and hue is usually used to communicate an ordering in the data being visualized. We briefly describe some of the orderings that have been prominent in the history of art and color science.

The concept of a color wheel, where all possible hues are arranged in a circular fashion, can be attributed to Newton. The circular pattern implies a "wrapping around" of hue that does not have a physical basis, since hue is just a mapping of wavelengths. (Note that the hue of purple – the mixture of red and blue -- is the only hue not found in the rainbow spectrum – it was artificially constructed by Newton so that his hue ordering would have a circular property.)

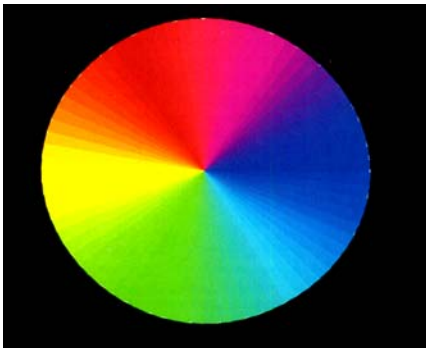

Figure 1: The Newtonian Color Wheel

An alternative geometric formulation of the hue dimension is found in Goethe's color triangle[X]. Instead of the continuous circle of hue found in the Newtonian color wheel, Goethe identified three "primary" hues as the fundamental building blocks of color. The geometric arrangement is that of a equilateral triangle. Hence, any point inside that triangle can be expressed as a linear combination of the vertices, with the interpretation that any realizable hue can be formed from one of the three basic hues.

A third scheme is due to Kandinsky [X], who as both a practitioner and teacher of painting, was searching for a psychological ordering of color that would provide a logical basis for color choices. Prominent in his scheme were three important colors: yellow and blue at the "poles," and red as the central and dominating hue. This triad still is prominent in the teaching of color theory in American art schools, even though the designation of these as "primary colors" is rooted in Kandinsky and Goethe's notions of color psychology rather than in the physical basis of color.

Kandinsky referred to picture the relationship between colors on a linear scale, rather than a circle of color. He visualized yellow as developing out of white, and having a strong relationship to it; likewise blue stemmed from black. Indeed, all the hues could be arranged in a linear fashion from white to yellow through red to blue, and finally black – an arrangement of hues that corresponds to decreasing order in the eye's sensitivity as a function of wavelength. Another progression of Kandinsky goes in the "other direction" – passing through green rather than red at its midpoint.

Saturation

Saturation can be thought of as how *chromatic*, or "washed-out" the color is. In terms of the spectral power distribution, it can be thought of qualitatively as the number of different wavelengths contributing to a sensation of color, or the "width" of the curve. Another way to think about saturation is how pronounced a "dominant" hue is. A fully saturated color is a pure hue that will have a single Dirac-delta spike as its power distribution curve, but except for a few situations (such as laser light), such *monochromatic* colors almost never occur in nature. On the other hand, a desaturated color would have a perfectly uniform SPD, thus being a white or a shade of gray. In other words, there is no distinguishing hue that stands out from any other.

While this definition of saturation holds for all hues, the *level of saturation perception varies with hue*. In particular, yellow appears to be less saturated than any other pure hue – it is closest to being achromatic, and desaturates quickly as other wavelengths are added to it.

Effects of the Attributes on One Another

This brings us to one of the main difficulties with constructing accurate progressions of color is that these three perceptual color variables, while established as the fundamental attributes that our visual perception process recognizes, are not *orthogonal* to each other – varying one will have perceptual effects on the others. For instance, Purdy [X] performed color matching experiments where subjects had to adjust the hue of a displayed patch to match a sample patch which was twice the luminance, but actually the same wavelength. The changes in hue as luminance changed, as determined by the amount of adjusting that the subjects did to make the two patches perceptually match up, was as much as 30 nm for a deep red. The apparent dependence of hue with luminance is known as the Bezold-Brucke effect, and similar perceptual dependencies have been described for lightness-saturation, and hue-saturation [REF]. For perceptually accurate color map work, this is a complicating factor.

For these reasons, while hue, lightness, and saturation is reasonable for humans to specify single colors, for doing color computations, particularly path generation in color space, we desire to work in a different manner than just simply interpolating these three dimensions, as we will see later.

Perceptual Illusions of Color

The above-mentioned dependence of hue perception on luminance is but example of a general class of problem that we need to tackle to construct perceptually-accurate color visualizations. Our perception of color can be, in many ways, affected by the presence or absence of other, surrounding colors with various attributes. Thus, having a mere knowledge of how humans perceive individual colors is not enough, for we need to consider these psychological effects that colors can have on each other's perceptions.

A general principle is to avoid using colors that could potentially arise due to such "illusions" – we want all our perceived colors to be caused by the visualization of the data itself, not by perceptual effects. To this end, we examine several well-studied examples where color perception can be affected, and try to find colors which might be perceptually ambiguous.

Artifacts in Interactions of Color

Perhaps the definitive work on interactions of color with others is Josef Albers' 1963 volume, *Interaction of Color [Albers63]*. In the plates of the book, it shows that the perception of *every* qualitative attribute of color (hue, saturation, value) is dependent on the adjacent surrounding color environment.

As a simple example, changes in value perception can be demonstrated by three shades of gray: A patch of a medium level of gray will look darker when it is surrounded by a light gray background, and likewise lighter when it is surrounded by a darker background.

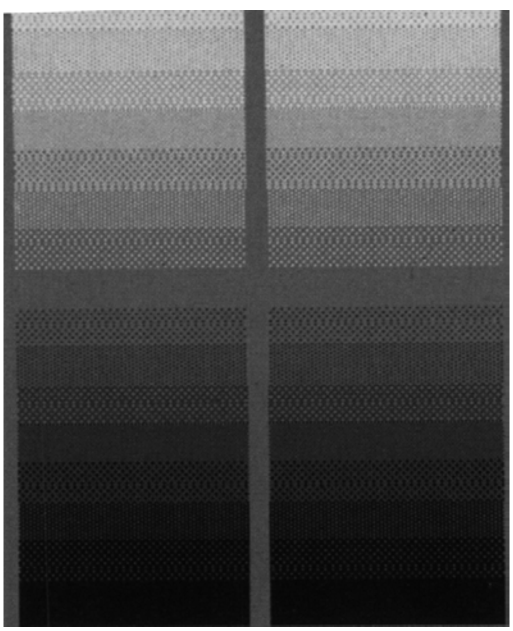

Figure 2: Effect of simultaneous contrast. Note how the top portion of the bar in the middle is perceived to be darker than the bottom part, despite the bar being the same throughout.

Changes in *perceived* hue and saturation are also easily shown, as illustrated in the following figures. Hue can be slightly shifted around the hue spectrum, making a patch of the same hue appear to be two different ones if superimposed over two varying backgrounds. Conversely, two slightly different hues can take on the appearance of being the same color.

Figure 3: The left and right portions of the multiple purple blocks erroneously appear to have a different hue.

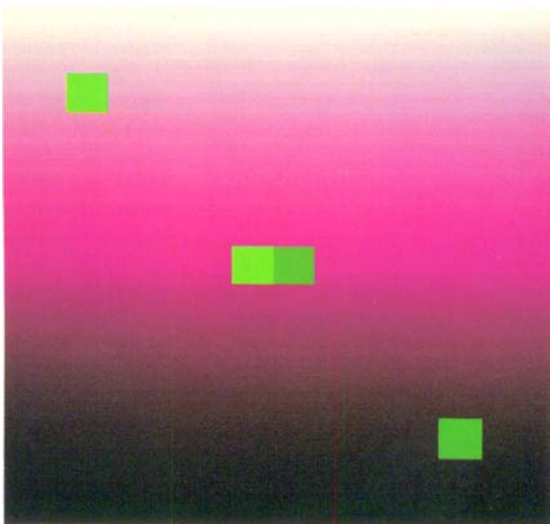

Figure 4: The upper-left and lower-right green blocks will appear to have the same hue, if the middle blocks are covered up.

In light of these potential misleading effects, we should:

- Select hues that appear on backgrounds of other hues using large and unambiguous hue transitions (for instance, deep blue and green, rather than two closely spaced hues of green).
- The corollary is that we should not have multiple hues for visualization of continuous quantitative variables.

Contrasts in Color

Intuitively, hues that are perceived as “warm” and “cool” correspond to those that appear on the opposite sides of the Newtonian wheel. When colors are diametrically opposite each other, they are termed *complementary* to each other.

Interestingly, if the colors are truly complementary to each other, a well-mixed combination of the two (for instance, a checkerboard dither pattern) will appear to be gray when viewed from a distance. Paradoxically, when they are viewed in closer proximity, complementary pairs of colors actually accentuate one another, making each other display their greatest brilliance, and even have “vibration” or “noise” at the borders where the two colors meet.

Figure 5: Complementary colors giving rise to "ringing" where they meet

Itten describes this phenomenon: “Two such colors make a strange pair. They are opposite, they require each other. They incite each other to maximum vividness when adjacent; and they annihilate each other, to gray-black when mixed – like fire and water.” [X]

Indeed, it has been psychologically suggested that the eye seems to require a balance of complementary colors. The use of complementary colors is a often-repeated method to ensure the harmony of colors in graphic and architectural design. Indeed, if the complement is not there, the mind will generate that color spontaneously – this is the phenomenon of *simultaneous contrast*. Among other effects, simultaneous contrast manifests itself as *afterimages*. The effect is easily experienced by staring intently at one image composed of saturated colors for about thirty seconds, and then quickly looking at a area of white. An afterimage, consisting of the original image with its colors complemented should be apparent.

Figure 6: Afterimage effect

Physiologically, what is taking place is that when the eye focuses on the white of the paper after having been exposed to a color, say, red for an extended period of time, the cones of the eyes responsible for color vision were not able to absorb the red component from the uniform spectrum of the white light being reflected from the paper. Instead, it absorbs the spectral complement of red, which turns out to be perceived as cyan.

An understanding of simultaneous contrast can be used to one's advantage – to *draw attention* to boundaries where complementary colors interface with one another, for instance, but it can mislead in this sense as well, if the mind perceives a hue that is non-existent in the visualization due to afterimage effects. Thus:

- Using complementary hue pairs may be effective for delineating borders
- multiple pairs of complementary colors should not be used in a given visualization, lest multiple (and thus ambiguous) gray areas occur due to complementary interference.
- hue scales with complementary colors should avoid including gray.

Linear Separation and Color Search

A recent observation relates the ability of the human visual system to find a specific color amidst other different colors in the background to the *linear separability* of the colors in color space. While this phenomenon does not strictly belong in the same category as those described above, it is still of concern in visualization, for the ability to *search* effectively for a given target color is often (or more) important than accurately perceiving it.

This linear separation effect, originally described by D' Zmura in 1991 [D'Zmura91], manifests itself when one is searching for a color target amidst points in the background with two or more different non-target colors. If the target color *T* was linearly separable (that is, separable by a plane in color space – here CIELUV is used, but the effect is present in any number of other parameterizations) from non-target colors

A and *C*, the time required for an observer to find the target was independent of the number of points. On the other hand, if *T* was collinear (such as with *A* and *B*), then the search time became linearly proportional to the number of points. [Figure 7]

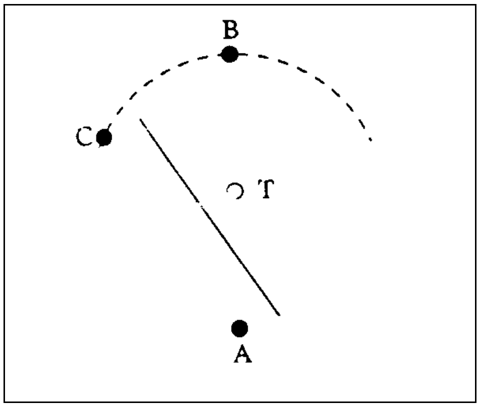

Figure 7: Color T is linearly separable with A and C, but not A and B.

Color spaces

In order to provide a computational basis for color processing and colormap generation, we need to define some way to quantify the space of colors that are realizable to our eyes, as well as the display hardware which must generate them.

RGB Color Space

The standard RGB color space, corresponding to the three primary colors of light used in computer displays, is used internally by computer hardware driving display devices. However, it is deficient for accurate color specification and selection purposes as RGB values are generally device-dependent. Because of the nature of the phosphors used in monitors, one set of RGB values may produce an entirely different perceived color on another monitor, or even another unit of the same model. Furthermore, the RGB space is not based on concepts of human perception, and specifying a color is unintuitive for the person without training. Usually, for the average user, specifying a color using an interface involving direct manipulation of the RGB values is a haphazard process of trial and error. Thus, we choose to do our color selection and processing in another space, and then only transform to RGB for the final display step. We now examine issues in the consideration of color spaces to motivate our specific choice.

Perceptually-Motivated Color Spaces and Their Desired Characteristics

Any system to model the space of perceivable colors must, first and foremost, recognize and incorporate the three fundamental psychological qualities of color described earlier: hue, saturation, and lightness.

It must be *colorimetrically precise and accurate* – any given color that is generated should have a set of quantities attached to it such that one could, given some standard definitions about the output device (such as colorimetric information about the "primaries"), reproduce that color with good faith that it will be perceptually equal.

Ideally, if perceived differences in color are to accurately reflect differences in numerical data, the metric of the color space should be dependent on the color or luminance-discriminating ability of the human visual system. Moreover, we desire our color space to be a *Euclidean space*: *the perceptual difference between any two equidistant points in the color space should be the same*. Such a color space is termed *uniform*.

In such a color space, a natural choice for our metric unit is the quantity known as the Just Noticeable Difference (JND), a general term in cognitive psychology which as the name would suggest, represents the smallest possible difference for any variable stimulus to be perceived as different from another. In color space, the JND is interpreted as the smallest distance between any two perceptually-different colors.

The parameters of hue, lightness, and value should be *geometrically discernable* and able to be controlled separately. In terms of the 3-D color space, there should exist 2-D manifolds (surfaces) of this space where moving on them should cause uniform progression on two elements, while leaving the other one fixed; these manifolds should be intuitive to construct, visualize and work with.

Finally, in order for its use to be practical in a semi-automatic environment in a computer, it must be *computationally and analytically tractable*.

The CIELUV and CIELAB Spaces

We focus the bulk of our attention on the (L*, u*, v*) coordinate system, known as CIELUV and used for additive light conditions, such as that of a CRT screen. (The sister (L*, a*, b*) system, known as CIELAB coordinate system, is used for reflective color, such as printing). Both these spaces have a metric defined in terms of the JND, and they attempt to have a uniform perceptual basis.

The CIELUV space is based on physical colorimetric principles: Central to the CIE space are the *tristimulus values X*, *Y*, and *Z*, defined as:

$$X = \int_V E(\lambda)\bar{x}(\lambda)d\lambda \qquad [1]$$

$$Y = \int_V E(\lambda)\bar{y}(\lambda)d\lambda \qquad [2]$$

$$Z = \int_V E(\lambda)\bar{z}(\lambda)d\lambda \qquad [3]$$

where $E(\lambda)$ is the spectral energy distribution, $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$ are the standard *CIE color matching functions* depicted in the figure, and *v* is the range of visible wavelengths, usually taken to be from 380 to 780 nanometers.

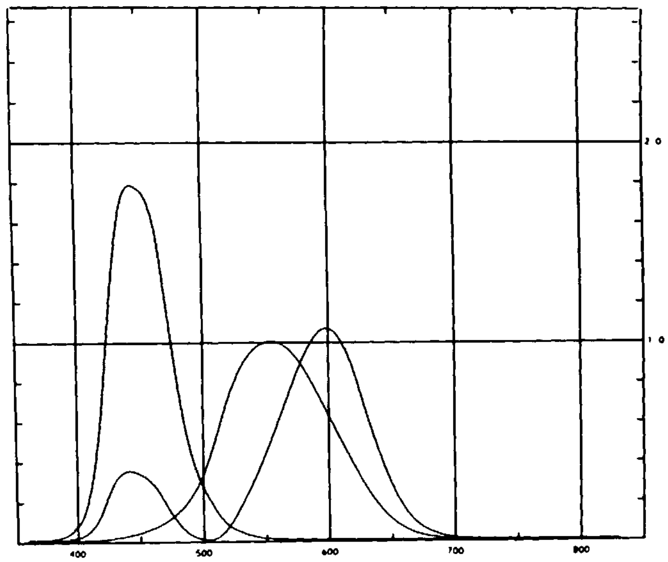

Figure 8: The X, Y, Z curves

In particular, the Y tristimulus value is the *luminance* of the color – two spectral distributions with identical Y values will result in the same perceived brightness. When two differing distributions have identical X, Y, and Z values, they will result in the same color being perceived. (Since for a given set of X, Y, and Z values, there are an infinite number of possible distributions with those values, it follows that it is entirely possible for various spectral distributions to result in the same perceived color – such colors are termed *metamers*.)

In CIE L*u*v* space, L* represents perceptual *brightness*, and is derived from the CIE Y coordinate by the following perceptual transform:

$$L^* = 116\left(\frac{Y}{Y_n}\right)^{1/3} - 16 \qquad \frac{Y}{Y_n} > .008856 \qquad [4]$$

$$L^* = 903.3\left(\frac{Y}{Y_n}\right) \qquad \frac{Y}{Y_n} \le .008856 \qquad [5]$$

where Y represents the CIE Y value, and Y_n is a normalization constant which is defined to be the Y value of the "white" color.

The above step is in addition to *gamma correction*, the step performed (often in hardware) to compensate for non-linear voltage response in the phosphors of the display. Gamma correction linearizes the relation between DAC values stored in the frame buffer and the *physical* luminance of the output on the display. Even with gamma correction, we still must *perceptually* linearize.

The u* and v* parameters can be taken together to express both hue and saturation, even though the latter two pseudo-perceptual notions are never explicitly mentioned in the actual color space specification. Specifically, we can think of a cylindrical parameterization of hue-lightness-saturation space, with L* (lightness) forming the vertical axis, and u* and v* forming a horizontal plane. We can think about saturation being the distance from the center axis, and hue being the angle that the u* and v* coordinates make with the L* axis, ie:

H = arctan(u*/v*)

L = L*

S = $\sqrt{(u*^2 + v*^2)}$

The Munsell Color Space

An alternative color space that is widely used in the design field is the Munsell Color System. It is motivated less from a physical point of view, but a design outlook. While it is not used by our system, we briefly describe it as an alternative example of color space organization.

The Munsell Book of Color is an empirical organization of colors based on human perception, and designed to be uniform using visual judgement, rather starting from an analytical formulation. [Levkowitz p. 41]. Munsell picked five primary colors, red, yellow, green, blue, and purple, as well as five colors to go in-between the primaries, and spaced them out evenly. (This means that truly complementary colors are no longer diametrically opposite each other as they are in the Newton hue wheel.) The distance between adjacent Munsell color samples is adjacent.

The Munsell scale is open-ended. As more intense and vivid colorants are found, they are added radially to the appropriate hue axis, so that the Munsell color solid actually grows outward over time. The overall shape of the Munsell color solid is that of a distorted sphere.

In terms of our desired characteristics for color spaces Munsell Color System meets all of them, except for the analytical representation, since it is first and foremost a organized collection of experimentally determined color samples.

Desired characteristics for data display and colormaps

We turn our attention now to examining how these characteristics of color can be used to our advantage in different types of visualization. Cognitive studies have shown that humans tend to associate various notions with the various perceptual qualities of color: for instance, hue, as mentioned earlier, can be used to convey order, while lightness has been traditionally used to communicate quantitative value. Thus, there should be systematic contentions for mapping aspects of color to various metatypes of data.

The Four Metatypes of Data

Statisticians have identified four metatypes of data which characterize information which be presented in visual form.[X] The former two are qualitative, while the latter two are quantitative.

Nominal data has no notion of distance or ordering between two values. One cannot say, for instance, that "John" is greater than "David." The only attribute that can be communicated is that of equality or distinctness between two values.

Ordinal data represents one step up from nominal data, having both name and ordering. "First place," "second place," and "third place" would be an example of an ordinal characterization. However, there is no notion of a metric, or distance between differences in value.

Interval variables have the property that differences in the numbers represent real differences in the variable. Another way to say this is that equal differences in the numbers on the scale represent equal differences in the underlying variables being measured. For example, consider the difference between 36 degrees and 37 degrees compared to the difference between 40 degrees and 41 degrees on either the Fahrenheit or Celsius temperature scale. Because the differences in the numbers are the same, when you have an interval variable you know temperature intervals are the same.

Ratio data is characterized by having a true "absolute zero" point with the semantic meaning of nothing, which allows one to define the notion of a ratio (division) between two data points meaningfully. Some examples of ratio variables are mass, time, length, energy, and temperature on the Kelvin scale. With ratio variables, the only difference from interval variables one can actually say, for instance, that this quantity is twice as great as this other quantity. This, when visualizing ratio data, it is important to have a point somewhere within the color map which can be clearly identified as the zero point.

Levels of Perception in Bivariate and Univariate Data

Besides the metatypes of data themselves, there are multiple levels of relationships that can be understood in a visualization of a given set of data. An understanding of these levels of perception, and how they manifest themselves in both univariate and bivariate visualizations is important because it provides further guidance for us on the features that our colormap is supposed to convey. An effective mapping of color should be suitable for all these levels.

In particular, Wainer and Francolini [X] establish an empirical approach to evaluating schemes for the representation of statistical variables in color. Following earlier work by Bertin[X], they define three distinct levels at which a statistical map, whether it be color or some other quantity, can be comprehended. At the first elementary level, direct translation from a perceived variable to a quantitative component is made (ie, determining the value of a single data point). At a second, intermediate level, trends between two perceived variables are related (for a visualization of bivariate data); in effect, the local interrelationships between two variables are understood, with the univariate analog being the appreciation of local distribution or other salient features such as maxima or minima. Finally, at the third or *superior* level, the entire structure or distribution of one variable is compared with that of the other, again with the univariate analog being appreciation of global distribution.

The second and third levels are perceptions of overall *form*, largely of a qualitative nature, rather than simply reading a single value at one point. As one might expect, different types of colormaps are optimal for these different visualization tasks, and so our system should somehow take this into account. Short of actually asking a user an explicit question about how one desires to utilize the visualization, we must resort to various heuristics to discern this. Some previous visualizations, such as IBM's PRAVDAColor system [Trenish95] have determined the rough spatial frequency of the currently-displayed data range, and then select a colormap optimized for either a high (form visualization) or low (value visualization) frequency context.

Note that in bivariate visualizations, we will usually be more interested in correlation between the two variables, rather than the absolute values of the variables themselves, since in the latter case, two separate visualizations of each variable would be more effective anyways. Therefore, we do not need to place as high a priority on color discriminability for quantitative colormaps as in the univariate case.

Rules of thumb for conveying meta-information

It is apparent that from a data visualization standpoint, there are multiple tasks that can be performed using color, and the requirements of these tasks dictate the type of colormap that needs to be employed for "optimal" perception.

The lightness attribute is most often used in visualizations where the perception of detailed forms is important. Our eyes have significantly higher frequency sensitivity to luminance data than other qualities of color. Physiologically, we have on the order of 120 million rods (which sense achromatic lightness), while we have about 2 million cones in each eye (which are responsible for color vision).[X]

On the other hand, hue is a natural parameter for conveyance of a certain ordering, as we have discussed earlier. Saturation can also be strategically used to increase a color scale's potential to convey form.

Guidelines for choosing and generating effective color maps

With a understanding of the type of information we want to convey in our data, and what types of communication the various perceptual attributes of color most readily communicate, we are now ready to set some concrete guidelines for which types of color mapping schemes work best in various situations.

Trumbo's principles:

The general guidelines for choosing color that we have discussed so far can largely be put into a functional framework of four concepts, first put forth together by Trumbo [X]. The first two apply to both univariate and bivariate visualizations, while the latter two deal with mapping two independent variables to a single color.

- *Order*: If the levels of a statistical variable are ordered, then the colors chosen to represent them should be perceived as preserving the order.
- *Separation*: Important differences in the level of a statistical variable should be represented by colors clearly perceived as different.
- *Independence of mappings*: If preservation of univariate information or display of conditional distribution is a goal, then the levels of the component variables should not act to obscure one another.
- *Trichotomy of colors*: If display of positive association or correlation is a goal, scheme elements should resolve themselves visually into three classes: those on or near the principal diagonal which would illustrate correlation between the two variables, those above it, and those below.

In terms of Trumbo's first principle of order, a main desiderata of a univariate color map is the accuracy in which a given perceived color can be quantitatively related to where it lies within our specified range.

Example Color Scales

The gray scale can be considered to be the basis for which all color scales are derived. There is only one attribute that varies along the scale, namely lightness. It is, as one might expect, the best at conveying pure form information. However, as Levkowitz points out, it suffers from limited "dynamic range" – a relatively short path length in CIELUV space.

The rainbow scale is commonly used in information visualization, where it is often the default color scale. However, it has several serious shortcomings: since it only varies the hue parameter, it is not perceptually uniform with brightness and saturation, as mentioned earlier. In visualization, this can create possibly misleading artifacts. For instance, the yellow region of the spectrum mav serve to draw attention and "highlight" an area which in reality, simply happens to fall in the value interval that maps to the yellow range and has no special semantic significance.

Other frequently used univariate color scales include the heated object scale, and the magenta color scale. They are described in more detail in [Levkowitz97].

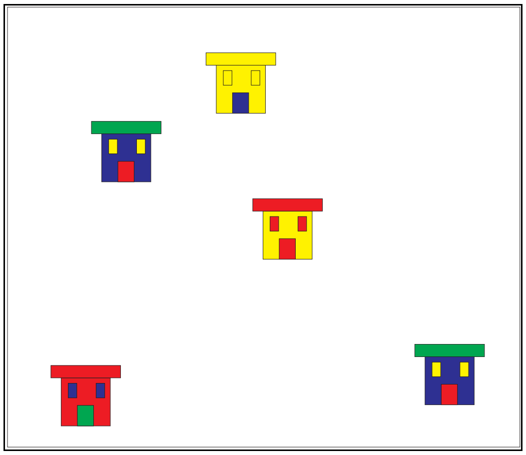

Figure 9: A typical visualization of nominal data, based on Levkowitz's color icons

Nominal Visualizations

We now examine how visualization of various metatypes of data has varying needs, and how these requirements build upon one another.

With nominal values, color serves the purpose of identification, and the ability to differentiate a color from other points is key. Often times, we will use variation of visualization techniques such as Chernoff faces [X]. Multi-dimensional nominal spaces can be represented with largely similar color sets by adding search features (such as doors and windows in the "houses" above.) Those sets need not be orthogonal, as long as other search considerations (such as linear separation) are taken into account.

Ordinal Visualizations

Hue now plays a more direct role in the visualization of any type of ordinal data, since order is the key attribute that we want to communicate to the viewer.

Interval Visualizations

Here, distance between two data points becomes a meaningful notion, so we will generally use a single hue or a pair of complementary ones, while relying on luminance or saturation if form information is to be more effectively conveyed.

Ratio Visualizations

Finally, to progress from interval to ratio visualization, we need a midpoint (or zero) which can be easily distinguised from the rest of the color space. Often, this is represented by gray, with values being greater or lesser than the midpoint/zero value being represented by complementary hue pairs, to facilitate knowing where relative to the midpoint a point is (ie, lesser or greater than it).

Trenish, et. al. [Trenish95] summarize some these guidelines in a table.

Data Type	Spatial Frequency	Representation Task		
		Isomorphic	Segmentation	Highlighting
Ratio (true zero)	Low	*Luminance:* uniform *Hue:* opponent or complementary pairs *Saturation:* monotonically increasing from gray	Even number of segments Many segments OK	Larger range for highlighted features
	High	*Luminance:* monotonically increasing *Hue:* opponent or complementary pairs *Saturation:* monotonically increasing from gray	Even number of segments Fewer segments	Smaller range for highlighted features
Interval	Low	*Luminance:* uniform *Hue:* opponent pairs *Saturation:* monotonically increasing from gray	Many segments OK	Larger range for highlighted features
	High	*Luminance:* monotonically increasing *Hue:* uniform or small hue variation *Saturation:* monotonically decreasing	Fewer segments	Smaller range for highlighted features
Ordinal	Low	*Luminance:* uniform *Hue:* variation around hue circle *Saturation:* monotonically decreasing	Fewer segments	Increase luminance of highlighted area
	High	*Luminance:* monotonically increasing *Hue:* variation around hue circle *Saturation:* uniform	More segments	Increase saturation of highlighted area
Nominal	Low High	*Luminance:* uniform *Hue:* variation around hue circle *Saturation:* uniform	Fewer segments than 7	Increase luminance or saturation of highlighted area

Figure 10: Some rules of thumb for assigning colors to visualizations. [Table from Treinish, et al.]

Univariate Path Generation in Color Space

To generate a color ramp for continuous quantitative data, we ideally would like to interpolate in a straightforward geometric way in color space between two chosen endpoints, and from that interpolation, generate a 1-D manifold in color space that captures the scale.

For a continuous color scale, in order to maximize the number of JNDs (ie, length in our color space), we will usually start with black, and end with white, although this is by no means necessary or even desirable (for instance, in the case of using complementary hues in ratio data visualizations).

As an aside, Van Overvelt takes path length extension one step further, by making the color map cyclic (that is, allowing the path traced out in color space to cross) at the expense of ambiguity when considering a local color[X], depicted below in the figure.

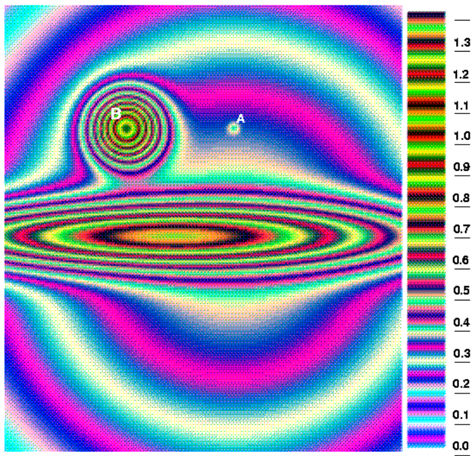

Figure 11: Repeating colors

In general, we will start with a set of guidelines (such as the need for hue ordering), which translates into geometric constraints in our chosen color space (going clockwise around the hue space, for instance), and insert color points as we go along a regular path designed to fit whatever constraints are present.

Levkowitz outlines various iterative methods for ensuring perceptual linearity and uniformity for generated color maps, as well as guidelines for experimentally evaluating the perceptual effectiveness of maps. These perceptual steps should be an integral and systematic part of the post-processing of any generated colormaps before they are made known to the user, and should be generally automatic in their operation. We will not further discuss them here, since they are only ways to perceptually refine maps chosen or generated *a priori* to conform to our rough guidelines (for hue order, etc).

The above, which generates continuous maps, is also natural for discretized (for ordinal data) maps. For purely nominal data, we can simply choose a set of colors in our color space that satisfy some criteria, such as linear separability. The latter can be achieved by taking the colors as the vertices of some planar polygon in hue space.

Surfaces in Color Space for Bivariate color maps

The space that we employ is a modified cylindrical parameterization of (L*, u*, v*) space, with L being represented by the vertical axis, and u*, v* being the other two axes which, taken together with L*, form an orthonormal basis.

Robertson[X] and others have proposed various geometric map arrangements in this space. Pham[X] extends this work by using B-spline curves and patches to represent the maps in color space. We outline some of these arrangements below, which are described in more detail in [Robinson]:

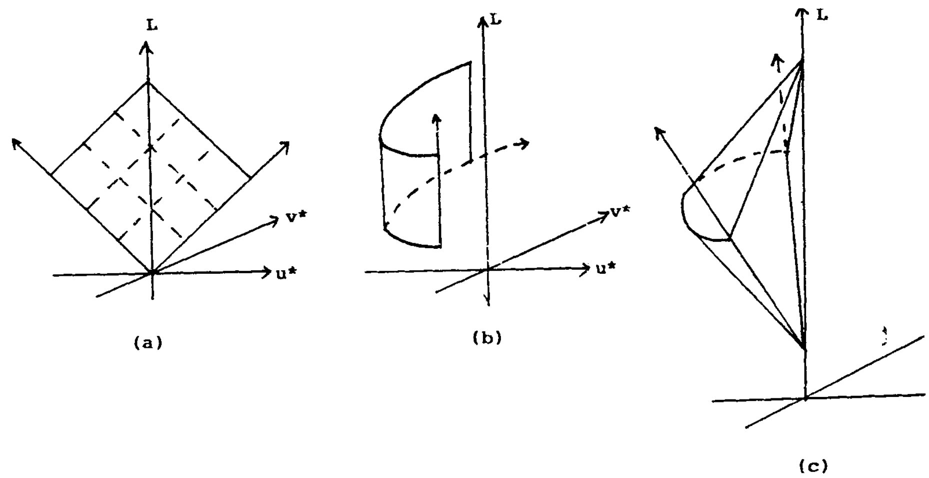

Figure 12: Some sample geometric parameterizations for bivariate color maps

A basic parameterization, shown in (a) above, is that of a square, with the achromatic lightness axis forming one of its main diagonals. We can see that this provides a natural separation of the surface into two areas of a complementary hue pair, one above and one below the main diagonal. Thus, we have three readily identifiable regions, two regions of constant hue, and one without hue (which would indicate data points correlated with each other).

Another scheme involves one variable being assigned to the vertical lightness axis, while the other one is linked to hue. In our color space, this results in a partial-cylindrical shape, as shown in (b) above. The use of such a scheme would be ideal, for instance, if one variable was ordinal and could thus be best expressed using the dimension of hue. However, a disadvantage is that correlation is not so easily seen with this kind of scheme.

While each of the above-depicted parameterizations has their merits, perhaps the most versatile map is a "double" cone parameterization (depicted in (c) above.). It allows one to see both positive and negative correlations, as illustrated in the figure above. Specifically, positive correlation is illustrated by a progression from black, to a fully-saturated hue halfway, that then decreases in saturation until the maximum value, represented by white (graphically, a planar slice through the cone perpendicular to the u*, v* plane). Negative correlation is illustrated as a hue "arc" from one end of the cone to the other, of medium lightness.

Analogous to the univariate case, we can start by systematically generating a surface in 3-D color space, or more practically, choosing from a "library" of basic geometric shapes, such as the three presented above, and transforming them (eg, moving vertices, or rotating them) as needed.

We consider origin and ending points in color space (this time for two variables, where there will need to be two separate ending points, and possibly even two separate origins). Additionally, we will also have certain guidelines, such as wanting to show the diagonal of positive correlation as a pure lightness map, and turn them into geometric features of our surface (eg, having the vertical lightness axis be the diagonal of our figure).

A Framework for Colormap Generation and Exploration

The space of perceptually-based color maps for a given univariate or even bivariate map is potentially large, even in light of the guidelines that we put down here. Furthermore, we want to view colormaps as an entity of the visualization that can be changed to emphasize various aspects of the data being visualized.

Therefore, we would like a method for generating colormaps which conform to guidelines that we have put forth, which is both driven by the current visualization function at hand, as well as the characteristics of the data being visualized. Furthermore, we also want to be able to quickly be able to come up with, and explore other colormaps which are related but different (i.e., geared towards showing different characteristics of the data, such as order versus form), rather than just going through a one-time process of generating a static map.

We now describe a framework which satisfies these needs: it allows us to generate colormaps in a largely automatic, yet targeted, manner, given some high-level user input on the goals of the visualization task at hand. Furthermore, it provides a structure for exploring various rules, guidelines, and heuristics for color mapping, including new ones which might be under consideration.

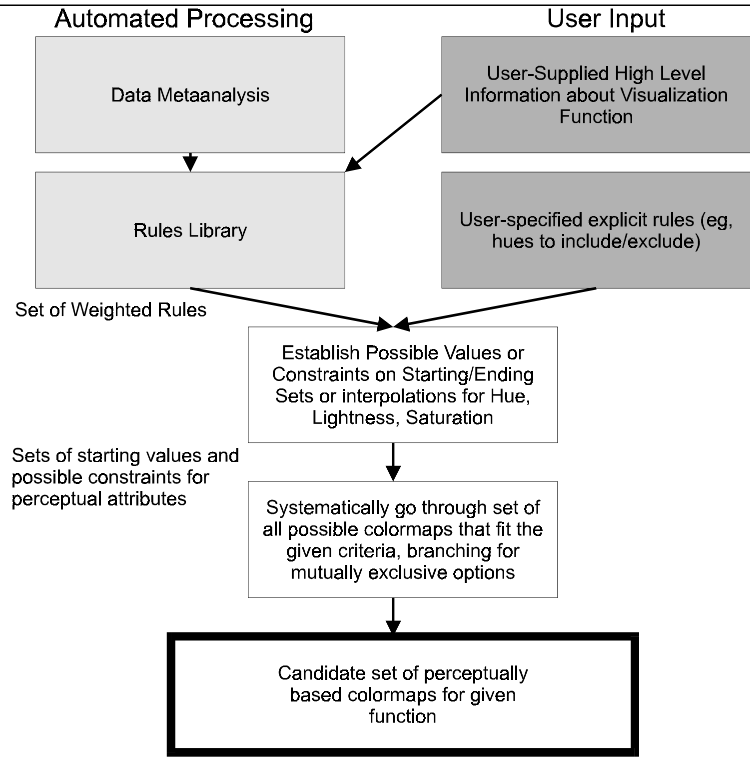

Figure 13: Functional block diagram for our color generation system

Motivation

The interface we propose is organized around a rules-based colormap generator. Central to this part is a library of rules encapsulating guidelines and constraints for color usage in visualization. Zhou and Feiner [Zhou96] have done some previous work in characterizing data for display.

```
If(SEPARATE_INTO_TWO_CLASSES ∈ Visualization.Goals) then HueSet ←
    ComplementaryHueGenerator [1.0]
```

If one of the goals of the visualization (as explicitly specified by the user, set as a default by an application, or decided by a higher-level process) includes the separation of quantitative data into two classes, then we will probably want our colorscales to incorporate a pair of complementary hues.

```
If(Var1.SpatialFrequency = LOW) then {
    Interpolator ← MidpointInterpolator [1.0]
    Interpolator ← LuminanceInterpolator [0.5]
    }

If(Var1.SpatialFrequency = HIGH) then {
    Interpolator ← LuminanceInterpolator [1.0]
    Interpolator ← MidpointInterpolator [0.5]
    }
```

Here, this set of two rules illustrates the role that the weights, denoted in our notation by [brackets], play in our system. We see that for low spatial frequencies, a hue-based scale is preferred, while for high-frequency data sets, a luminance scale is usually better. Yet, at times, we may want to change to another scale to see other characteristics of the data. Thus, we do not exclude the less-favored option, but merely give it a lower weight. This will automatically be reflected in the higher weights given to the "preferred" option, depending on the frequency parameter.

```
If(Var1.Type = NOMINAL) then HueSet ← LinearlySeparableHueGenerator [1.0]
```

This rule simply states that if our data type is nominal, then we should have a set of linearly separable colors (which will be automatically generated by the LinearlySeparableHueGenerator mini-module.)

```
If(BiVariate & Var1.Type = CONTINUOUS & Var2.Type = CONTINUOUS) then {
    If(SHOW_POSITIVE_CORRELATION ∈ Visualization.Goals) then Interpolator ←
    {Interpolator ∩ PlaneInterpolator} [1.0]

    If(SHOW_POSITIVE_CORRELATION ∈ Visualization.Goals |
    SHOW_NEGATIVE_CORRELATION ∈ Visualization.Goals) then Interpolator ←
    {Interpolator ∩ DoubleConeInterpolator} [1.0]
    }
```

We believe that this automated colormap generation, followed by interactive user selection and subsequent modification, is more advantageous than having users generate maps by hand the initial phase, or otherwise choose from a default "library." The initial set of maps generated will hopefully be closer to perceptually optimal for the given visualization context than what a user would come up with, or a "one-size-fits-all" library. At the same time, it allows for interactive exploration by varying color as a visualization parameter to highlight, filter, or emphasize certain subsets of the data.

In order for the color manager to generate color maps that are suitable for the visualization task at hand, we need to know something about our maps' desired characteristics regarding the way they use color. For instance, visualizing a continuous scalar field and assigning color maps to Chernoff faces are two tasks with significantly different requirements in terms of color usage. The visualization module itself, therefore, should have a standardized interface that the color manager can query to get information, or meta-characteristics, on the type of visualization being performed which it needs to properly generate suitable colormaps.

Such meta-characteristics that can be specified are:

- Metatypes of variables being visualized
- Text/Icons/Continuous distributions
- Statistical characterizations (high frequency/low frequency/high variance/low variance)

Awareness of such meta-information in the visualization environment is also a part of some other systems, including [Trenish95].

A part of our framework, therefore, will incorporate a standardized way of describing the meta-characteristics of a given visualization metaphor, including its usage of color resources. Such a method of description can take on one of several different forms -- for example, a XML DTD -- and initial work is being done in the context of specific visualization architectures [RIVET]. Of course, in order to meaningfully create such a standard, we would first have to come up with some sort of taxonomy to characterize these visualizations. This opens up another research question in and of itself far beyond the scope of color management and the scope of this paper.

What are some sample rules? For instance, in a visualization of continuous univariate scalar data, we would want a colormap which itself was continuous, with changes in the brightness channel. A scatter plot of discrete values might have hues separated by as much distance (in a perceptual metric space) as possible..

The rules would determine the high-level characteristics of the desired colormap itself: whether the map would be continuous or discrete, the number of independent variables, as well as any colors that must be included or excluded from the final result.

The following attributes are also examples of those which would be specified by rules:

- Color characteristics to vary in the colormap – hue, lightness, or saturation
- Geometric parameterization of the color space (cone, double cone, plane)

The *perceptual heuristic evaluator* takes a map and returns a numerical heuristic value on the perceptual effectiveness of that map in this visualization context, which is then used by the selection UI to initially present the most "promising" maps based on this metric. A straightforward and trivial scheme which could be used is a simple path (or surface, in the case of bivariate maps) integration for continuous dimensions – i.e., determining arc length or surface area, which corresponds to a "dynamic range" for a color map. Of course, there are many other characteristics that can potentially be a factor in perceptual effectiveness of a given colormap in a visualization, but this modular design allows us to experiment with various heuristic functions as necessary, in a relatively easy manner.

Since it is likely that such an evaluator would need knowledge of meta-data, and visualization meta-chracteristics on a global level, it should be able to have knowledge of the visualization and its data as well (taken from the Map Generator) . However, the goal of the evaluator is not to *eliminate* bad colormaps, but rather to simply *prioritize* generated colormaps that already conform to our rules for our current visualization. Hence, the calculation of the heuristic should be approximate and quick, rather than slow and "accurate" – for instance, analyzing the entire dataset is probably overkill.

Future Work

In this report, we have noted the usefulness of developing a taxonomy of visualization metaphors and techniques, from which we can develop a standard way to specify to various modules the graphical metacharacteristics of each. A systematic rule-based system to consolidate the guidelines for effective color selection is the key to the initial generation of colormaps.

Color is only one part of the visualization equation. Much work has been done on using texture and shape glyphs to add more dimensionality and variables to the range of typical visualizations. With advances in computer audio processing capability, recent research has also focused on developing sound as a "visualization" technique. A promising future direction then will be to integrate these approaches in the context of the total human perception and cognitive experience. At the higher levels, are the various percepts orthogonal to each other?

Animation of color maps is also something that could be potentially addressed. The color maps described in this report are temporally static and unchanging. While our interface encourages experimentation and interaction in switching between various color maps in the process of exploratory visualization, having animated color maps which have an inherent temporal, dynamic component to it may be useful, to aid for instance in perceiving gradients in scalar field data.

Another will be continued development on perceptual metrics, including psychophysical energy-based methods for consideration of perceptual effects on the global level, rather than on just the local level. The modular design of this system provides an excellent test bed for the design and evaluation of perceptual metrics and heuristics.

Finally, accurate psycho-mathematical models of human perception will be another field of research which will be fruitful in understanding human perception.

Conclusion

Is automatic color map generation even feasible? Given boundary conditions which are based on semantics, we can generate a effective, perceptually optimal color map which can represent multiple types of data. Furthermore, by shielding the user from the complexities of working directly in perceptual color space, we can make the dimension of color more accessible and usable as a tool for exploring data sets being visualized.

As for the question of fully automatic color map generation, the initial motivation, the philosophy and the requirements of exploratory visualization seem to preclude fully-automatic generation. Given a notion of what kind of quantitative data to look for, a computer could apply various analysis techniques to the data to automatically generate color maps to accentuate features it deems to be salient, such as maxima, minima, and other patterns. For instance, the PRAVDAColor module does rudimentary frequency-space analysis of the data to aid in its color schemes, having a set of maps for low and high-spatial frequency data.

Of course, one has to answer the question that if it is possible to do fully automatic data analysis and feature extraction, then what is the purpose of humans in the visualization pipeline anyways?

Since we are dealing with data for which we do not *a priori* know the characteristics of, we believe that it is safe to say that there will always be much work required to make cognitive conclusions about the data, in terms of high-level patterns that go beyond basic mathematical phenomena such as maxima, minima, or gradients. The role of color is not that of a panacea, but rather, to add the expression of more variables to the overall visualization, and leveraging the human mind's cognitive capability to quickly search, filter, or otherwise zero in on color, generally at a pre-attentive level. If we approach the user of color taking into advantage these special abilities that we have to process information which is unique to color, through a framework that keeps perceptual principles in mind, we can more effectively harness those abilities and increase the efficiency of visually exploring data.

[Albers63]

[D' Zmura91]

[Kawai95]

[Levkowitz91] H. Levkowitz. Color Icons: Merging Color and Texture Perception for Integrated Visuzliation of Multiple Parameters. In *Proc. IEEE Visualization 1991.*, pages 164-170.

[Levkowitz92] H. Levkowitz and G. T. Herman. Color scales for image data. *IEEE Computer Graphics and Applications*, 12(1):72-80, January 1992.

[Levkowitz97]

[Rheingans92] Reinegans, Penny. Color, change, and Control for Quantitative Color Display. In *Proc. Visualization 1992*, pages 252-259.

[Robertson]

[RIVET98]

[Treinish95] PRAVDAColor

[Tufte]

[Uselton]

[Ware88] Ware, Colin. Color Sequences for Univariate Maps: Theory, Experiments, and Principles. *IEEE Computer Graphics and Applications,* September 1988, pages 41-49.

[VanOvervelt97] http://www.acm.org/jgt/papers/vanOverveld97/

[Zimbardo96] Philip Zimbardo, Psychology and Life 14th ed. HarperCollins, New York. 1996.