Escaping RGBland: Selecting colors for statistical graphics

Achim Zeileis [a,*], Kurt Hornik [a], Paul Murrell [b]

[a] *Wirtschaftsuniversität Wien, Austria*
[b] *The University of Auckland, New Zealand*

A B S T R A C T

Statistical graphics are often augmented by the use of color coding information contained in some variable. When this involves the shading of areas (and not only points or lines)—e.g., as in bar plots, pie charts, mosaic displays or heatmaps—it is important that the colors are perceptually based and do not introduce optical illusions or systematic bias. Based on the perceptually-based Hue-Chroma-Luminance (HCL) color space suitable color palettes are derived for coding categorical data (qualitative palettes) and numerical variables (sequential and diverging palettes).

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

Color is an integral element of graphical displays in general, and many statistical graphics in particular. It is easy to create color graphics with any statistical software package and color images are therefore virtually omnipresent in electronic publications such as technical reports, presentation slides, or the electronic version of journal articles (e.g., in *Computational Statistics & Data Analysis*) and increasingly also in printed journals. However, more often than not, color choice in such displays is sub-optimal because selecting colors is not a trivial task and there is relatively little guidance about how to choose appropriate colors for a particular visualization. When selecting colors for a statistical graphic, there are three main obstacles to overcome:

1.1. The colors should not be unappealing

It is not necessary for the colors in a statistical plot to reflect fashion trends, but basic principles such as avoiding large areas of fully saturated colors (Tufte, 1990) should be adhered to. The requirement is not that the user should have a degree in graphic design, but that the software should provide users with an intuitive way to select colors and control their basic properties. Thus, it is necessary to employ a *color model* or *color space* that describes colors in terms of their perceptual properties: hue, brightness, and colorfulness.

A color model typically supported by software packages involves the specification of colors as Red–Green–Blue (RGB) triplets. However, this specification corresponds to color generation on a computer screen (see Poynton, 2000) rather than corresponding to human color perception. For humans, it is virtually impossible to control the perceptual properties of a color in this color space because there is no single dimension that corresponds to, e.g., the hue or the brightness of the color. As a consequence, various perceptually-based color spaces have been suggested, where each dimension of the color space can be matched with a perceptual property. One approach widely implemented in software packages involves Hue-Saturation-Value (HSV) triplets (Smith, 1978), a simple transformation of RGB triplets (see Wikipedia, 2007b). Unfortunately,

---

* Corresponding address: Department of Statistics and Mathematics, Wirtschaftsuniversitat Wien, Augasse 2-6, A-1090 Wien, Austria. Tel.: +43 1 31336 5053; fax: +43 1 31336 774.

E-mail address: Achim.Zeileis@R-project.org (A. Zeileis).

the dimensions in HSV space map poorly to perceptual properties and the use of HSV colors encourages the use of highly saturated colors. A perceptually-based color model that mitigates these problems involves Hue-Chroma-Luminance (HCL) triplets (see Ihaka, 2003), resulting from a transformation of CIELUV color space (see Wikipedia, 2007a). This is the color space we advocate in this article.

1.2. The colors in a statistical graphic should cooperate with each other

The typical purpose of color in a statistical graphic is to distinguish between different areas or symbols in the plot—to distinguish between different groups or between different levels of a variable. This means that there will typically be several colors, or a *palette* of colors, used within a plot and that those colors should be related to each other.

A natural solution for this task is to vary (at least) one perceptual property of the colors, e.g., the hue or the brightness, keeping other properties fixed. In a perceptually-based color space, this corresponds to selecting colors by traversing paths through this space along its dimensions. This approach is implemented in many color picker tools (Moretti and Lyons, 2002; Meier et al., 2004), however, these are typically based on the HSV model. As argued above, the dimensions in HSV space are not truly independent and hence it is not possible to vary just one perceptual property while keeping the others fixed. This means that it is relatively difficult to select sets of HSV coordinates that yield colors that are "in harmony" (see Munsell, 1905). For statistical graphics, this is important because it can introduce size distortions in the perception of shaded areas and can produce optical illusions (Cleveland and McGill, 1983). These problems can again be addressed by employing the HCL color space. In addition to the earlier point that HCL space is useful because it allows humans to understand where within the space a particular HCL triplet is located, HCL also allows us to understand motion within the color space because distances between colors have an intuitive and rational meaning.

1.3. The colors should work everywhere

The final issue to deal with is that, in an ideal situation, colors should be selected so that they continue to work in any context. For example, different areas of a plot should still be distinguishable when the graphic is displayed on an LCD projector rather than a computer screen, or when it is printed on a grayscale printer, or when the person viewing the graphic is color-blind. These goals cannot always be attained, but attention should be paid to these issues and in many situations it is also possible to resolve any problems.

An example of this approach is **ColorBrewer.org** (Harrower and Brewer, 2003), an online tool for selecting color schemes for maps. It provides a collection of prefabricated palettes, information about the suitability for printers or color-blind viewers, and guidance on how to choose a suitable palette for coding various types of information. The drawback to this tool is that it only provides a fixed set of colors for each palette and there is no way to extend the existing palettes.

Following Brewer (1999) and Harrower and Brewer (2003), we take a similar approach and distinguish three types of palettes: qualitative, sequential and diverging. The first is tailored for coding categorical information and the latter two are aimed at numerical or ordinal variables. Unlike **ColorBrewer.org**, we suggest a general principle for selecting colors by traversing paths along perceptual axes in HCL color space. Consequently, the user can decide which path exactly should be taken and how many colors should be selected. Furthermore, by matching the paths with perceptual meaning, the suitability of a palette for color-blind viewers or grayscale printing can be assessed.

The remainder of the paper is organized as follows: Section 2 provides some motivating examples, showing how typical HSV-based graphics can be enhanced by using HCL-based palettes. Section 3 gives a brief introduction to the underlying HSV and HCL color spaces before Section 4 suggests strategies for deriving HCL-based qualitative, sequential, and diverging palettes. Some more illustrations are presented in Section 5. Section 6 offers some general remarks on the implementation in statistical software as well as some details on the implementation in the R system for statistical computing (R Development Core Team, 2008). Section 7 concludes the paper with a discussion.

2. Motivation

To show what can be gained by selecting appropriate color schemes, we present two examples for typical color graphics, contrasting commonly-used HSV palettes with more suitable HCL palettes. Both examples and all HSV palettes are taken from (the electronic version of) recent publications in *Computational Statistics & Data Analysis* (Volume 51).

Our first illustration is a heatmap, a very popular display for visualizing a scalar function of two arguments. Here, we use a bivariate kernel density estimate (for the Old Faithful geyser eruptions data from Azzalini and Bowman, 1990)—other typical applications include classification maps (as in Tenenhaus et al., 2007, Fig. 2), function surfaces (as in Harezlak et al., 2007, Fig. 8), or measurements on a spatial grid (as in Gelfand et al., 2007, Fig. 3). Fig. 1 shows a heatmap bringing out the relationship between the duration of an eruption of the Old Faithful geyser in Yellowstone National Park and the waiting time for this eruption. It reveals a multi-modal distribution: short waiting times (around 50 min) are typically followed by a long eruption (around 4 min) whereas long waiting times (around 80 min) can be followed by either a long or short eruption (around 2 min).
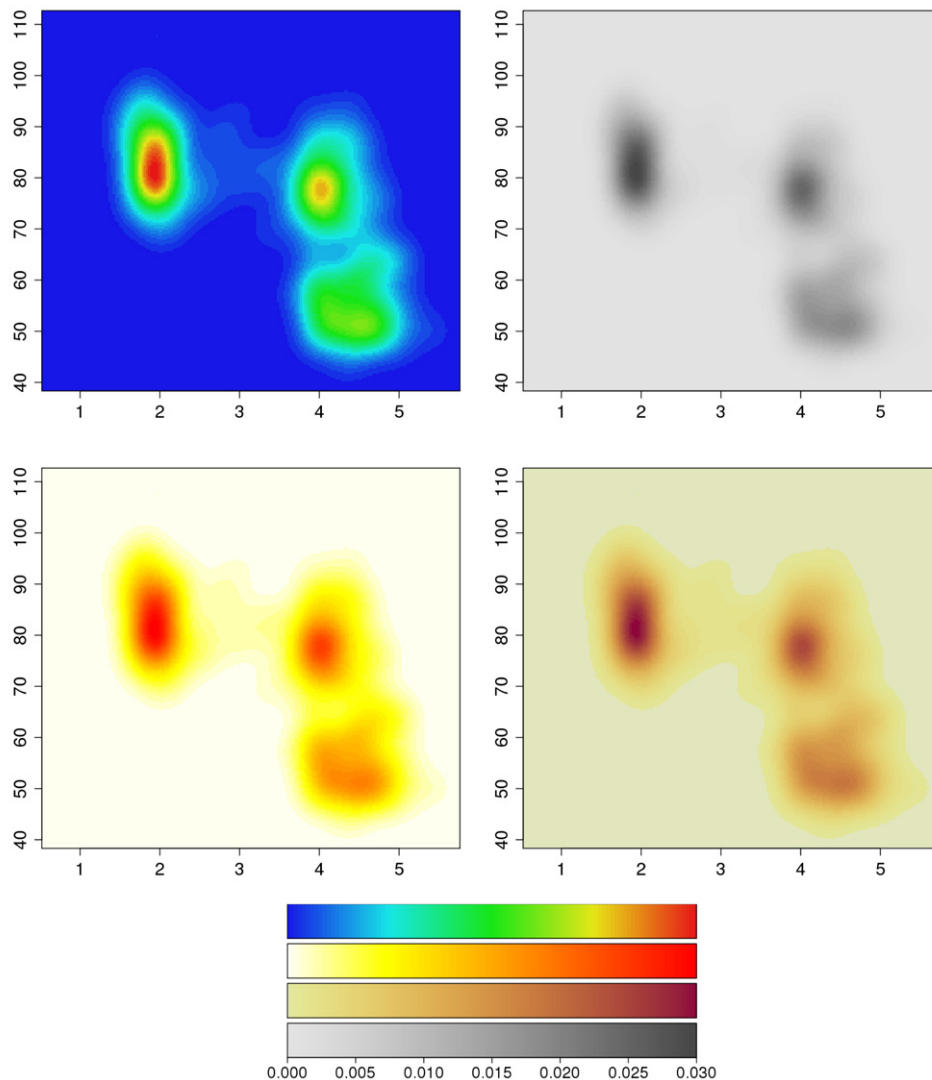
Fig. 1. Bivariate density estimation of duration (*x*-axis) and waiting time (*y*-axis) for Old Faithful geyser eruptions. The palettes employed are (counterclockwise from top left) an HSV-based rainbow, HSV-based heat colors, HCL-based heat colors and grayscales.

A simple and very effective palette for such a display is a set of gray colors as in the top right panel of Fig. 1. This is often (appropriately) used in printed papers when the journal does not offer color graphics—however, in journals that support color graphics (or on presentation slides and in interactive usage in statistical software packages), many users prefer to have colored displays and most often use HSV palettes (as in the two left two panels). These palettes code the variable of interest by varying hue in an HSV color wheel, yielding a "rainbow" of colors, as done by Harezlak et al. (2007, Fig. 8) or Tenenhaus et al. (2007, Fig. 2). The palette in the upper left panel codes increasing density by going from a blue to a red hue (via green and yellow)—a similar strategy are the "heat colors" in the lower left panel that increase from yellow to red. The latter works somewhat better than the former, however both palettes exhibit several drawbacks. The modes in the map look much more like "rings" rather than a smoothly increasing/decreasing density. The heatmap looks very flashy which—although good for drawing attention to a plot—makes it hard to hold the attention of the viewer for a longer time because the large areas shaded with saturated colors can be distracting and produce after-image effects (Ihaka, 2003).

In contrast, the gray colors used in the top right panel do not exhibit the same disadvantages, coding the variable of interest much better and without flashy colors. If, however, the user wants to increase the contrast by adding some color to the plot, this could be done by using a better balanced version of the heat colors (as shown in the bottom right panel). These colors also increase from a yellow to a red hue while using the same brightness levels as in the grayscale palette (see also Ware, 1988; Levkowitz and Herman, 1992, for similar approaches). Thus, when converted to a grayscale or printed out on a grayscale printer (as in the printed version of *Computational Statistics & Data Analysis*), the upper and lower right panel would look (virtually) identical. Both palettes have in common that they give increasing perceptual emphasis to regions with
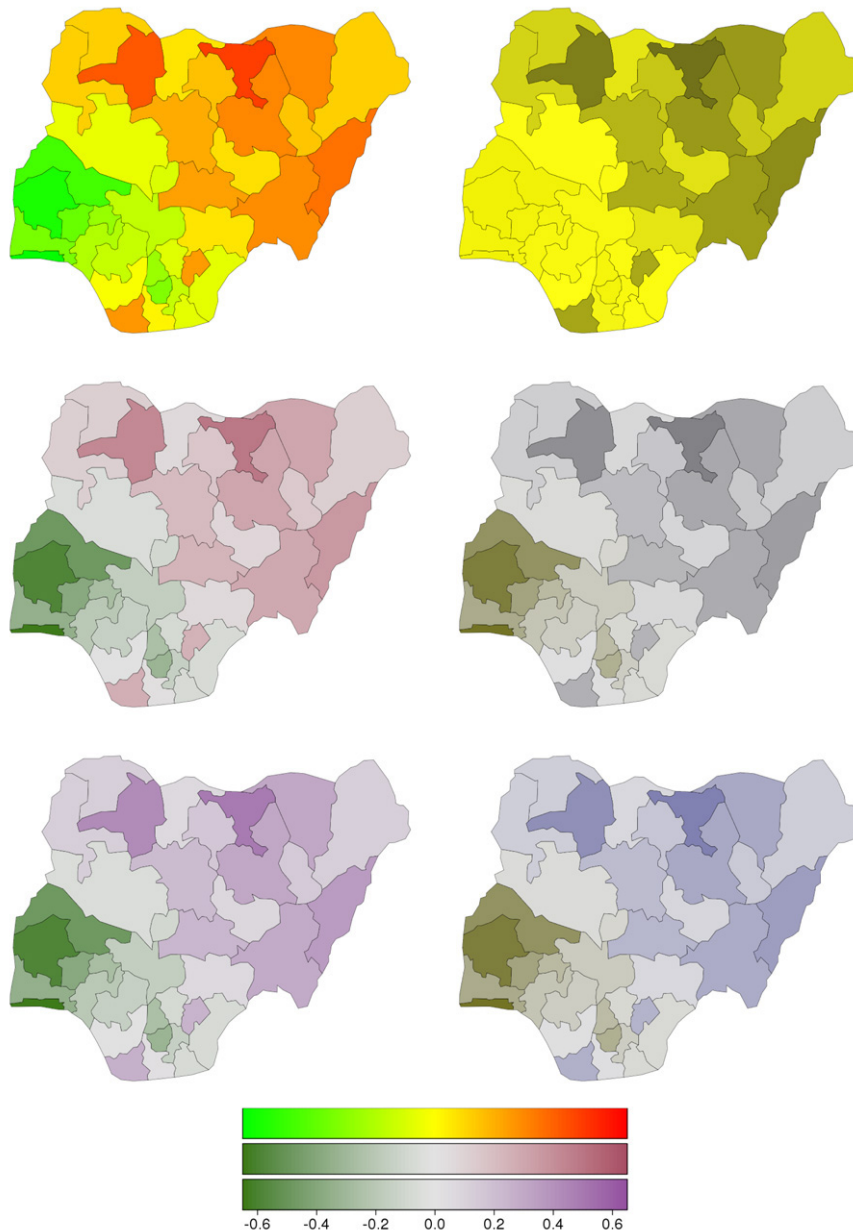
Fig. 2. Posterior mode estimates for childhood mortality in Nigeria. The color palettes employed are (from top to bottom) an HSV-based rainbow and two HCL-based diverging palettes. In the right panels red–green contrasts are collapsed to emulate protanopic vision.

increasing density, resulting in a heatmap that highlights the (small) interesting high-density regions and not to the large low-density regions surrounding them. By balancing the lower right panel with respect to its brightness levels (i.e., light/dark contrasts), it is assured that the graphic is intelligible for color-blind readers and that the same graphic works in colored electronic versions and printed grayscale versions of a publication.

The second example is presented in Fig. 2 (taken from Kneib, 2006, Fig. 5, left), depicting posterior mode estimates for childhood mortality in different regions of Nigeria. Spatial variations (not included in the model) in childhood mortality are brought out by shading a map according to the corresponding model deviations, revealing decreased mortality in the south–west and increased mortality in the north–east. Kneib (2006) uses an HSV-based palette coding the deviations by the hue, going from green via yellow to red (see the upper left panel of Fig. 2). Similar approaches are used by Harezlak et al. (2007, Fig. 8) and Tenenhaus et al. (2007, Fig. 2). Our first HCL-based palette in the middle left panel also employs green and red hues for negative and postive deviations respectively, but codes neutral values (around 0) by a neutral light gray. Compared to the HSV-based palette, this offers again a number of advantages: only the interesting areas are highlighted
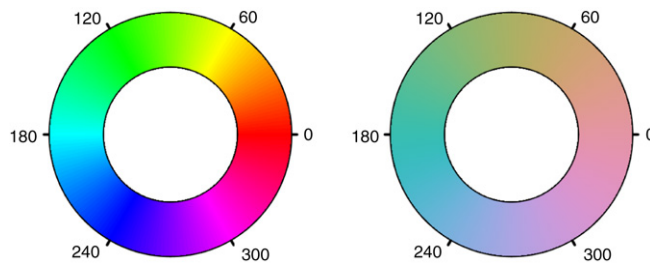
Fig. 3. HSV-based and HCL-based color wheel.

by high-chroma colors; flashy colors are avoided making it easier to look at the display for a longer time; positive and negative deviations with the same absolute size receive the same perceptual weight by being balanced with respect to their brightness.

Another advantage of the HCL-based palette becomes apparent when we collapse the red–green distinctions in the colors to emulate what a person with protanopia (a common red–green color blindness) would see. The right panel of Fig. 2 displays the respective palettes from the left panel, but transformed for protanopic vision using the tools of Lumley (2006). In the upper right panel, the low-mortality regions in the south-west just "disappear" because they are transformed to virtually the same yellow as the average-mortality regions. In contrast, in the middle right panel, regions with increased and decreased mortality remain clearly distinguishable. Only the type and amount of color is changed (red turns gray, green turns yellow/brown) but regions with large deviations are still easily indentifiable by being balanced towards darker grays. To avoid that one of the two "branches" of the palette is only gray in the protanopic version, we could also employ a purple (rather than a red) hue for positive deviations. This is used in the lower panels of Fig. 2 showing that after collapsing red–green distinctions, a useful yellow/brown–gray–blue palette remains from the original green–gray–purple palette. Thus, the lower left panel could be used ensuring that it works for both normal and protanopic vision. Similar arguments hold for deuteranopic vision.

3. Color spaces

For choosing color palettes, it is helpful to have a basic idea how human color vision works. It has been hypothesized that it evolved in three distinct stages: 1. perception of *light/dark* contrasts (monochrome only), 2. *yellow/blue* contrasts (usually associated with our notion of warm/cold colors), 3. *green/red* contrasts (helpful for assessing the ripeness of fruit). See Kaiser and Boynton (1996), Knoblauch (2002), Ihaka (2003), and Lumley (2006) for more details and references. Furthermore, physiological studies have shown that light is captured by three distinct receptors (so-called cones) in the retina and hence coded in three dimensions by the human visual system. The perception of this coding, i.e., the subjective experience of this light, is less well-understood; however, current psychophysical theories all describe perception in terms of three dimensions (Knoblauch, 2002). Therefore, colors are typically described as locations in 3-dimensional spaces.

The three dimensions used by humans to describe colors are typically:

- *hue* (dominant wavelength),
- *chroma* (colorfulness, intensity of color as compared to gray),
- *luminance* (brightness, amount of gray).

Thus, they do not correspond to the axes described above, but rather to polar coordinates in the color plane (yellow/blue vs. green/red) plus a third light/dark axis. Color models that attempt to capture these perceptual axes are also called perceptually-based color spaces.

A popular implementation of such a color space, available in many graphics and statistics software packages, are HSV colors. They are a simple transformation of RGB colors defined by a triplet (H, S, V) with $H \in [0, 360]$ and $S, V \in [0, 100]$. Although simple to specify and easily available in many computing environments, HSV colors have a fundamental drawback: its three dimensions map to the three dimensions of human color perception very poorly. The three dimensions are confounded: The brightness of colors is not uniform over hues and saturations (given value, see Fig. 3)—therefore, HSV colors are often not considered to be perceptually based.

To overcome these drawbacks, various color spaces have been suggested that more closely map to the perception dimensions, the most prominent of which are the CIELUV and CIELAB spaces developed by the Commission Internationale de l'Éclairage (2004). These color spaces are still not perfect—such an ideal is never truly attainable because color perception will always differ between individuals—but they facilitate color matching for humans (Zhang and Montag, 2006) and offer a significant improvement over HSV, as Fig. 3 illustrates.

Ihaka (2003) argues that CIELUV colors are typically preferred for use with emissive technologies such as computer screens which makes them an obvious candidate for implementation in statistical software packages. By taking polar coordinates in the UV plane of CIELUV, HCL colors are obtained, defined by a triplet (H, C, L) with $H \in [0, 360]$ and
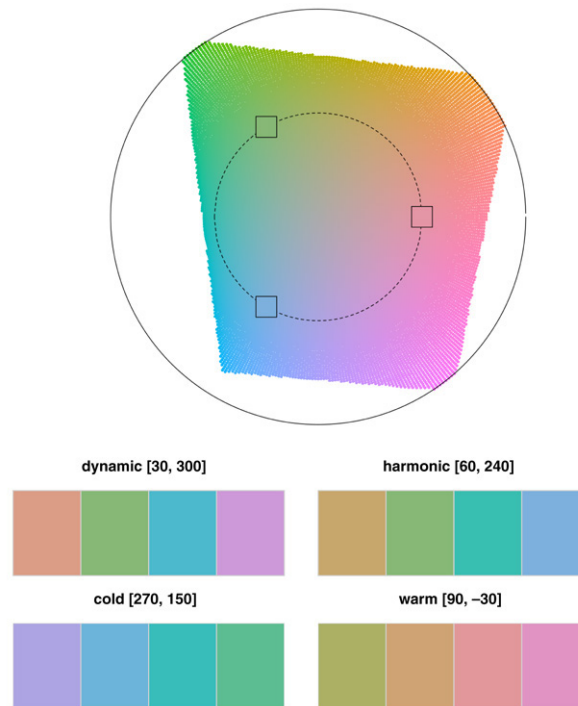
Fig. 4. Qualitative palettes: construction and examples. Top: Construction in the hue/chroma plane for $L = 70$, the dashed circle correponds to a radius $C = 50$ with chosen angles $H = 0$, 120, 240. Bottom: Examples by varying hue in different intervals for given $C = 50$ and $L = 70$.

$C, L \in [0, 100]$. Given a certain luminance L, all colors resulting from different combinations of hue H and chroma C have the same level of brightness, i.e., are balanced towards the same gray and thus look virtually identical when converted to a gray scale. However, the admissible combinations of chroma and luminance coordinates (within the space's boundaries) depend on the hue chosen. The reason for this is that some hues lead to light and others to dark colors, e.g., full chroma yellow is brighter (i.e., has higher luminance) than full chroma blue.

The balancing of HCL colors gives us the opportunity to conveniently choose color palettes which code categorical and/or numerical information by translating it to paths along the three perceptual axes. However, some care is required for dealing with the irregular shape of the HCL space which will be addressed in the following sections.
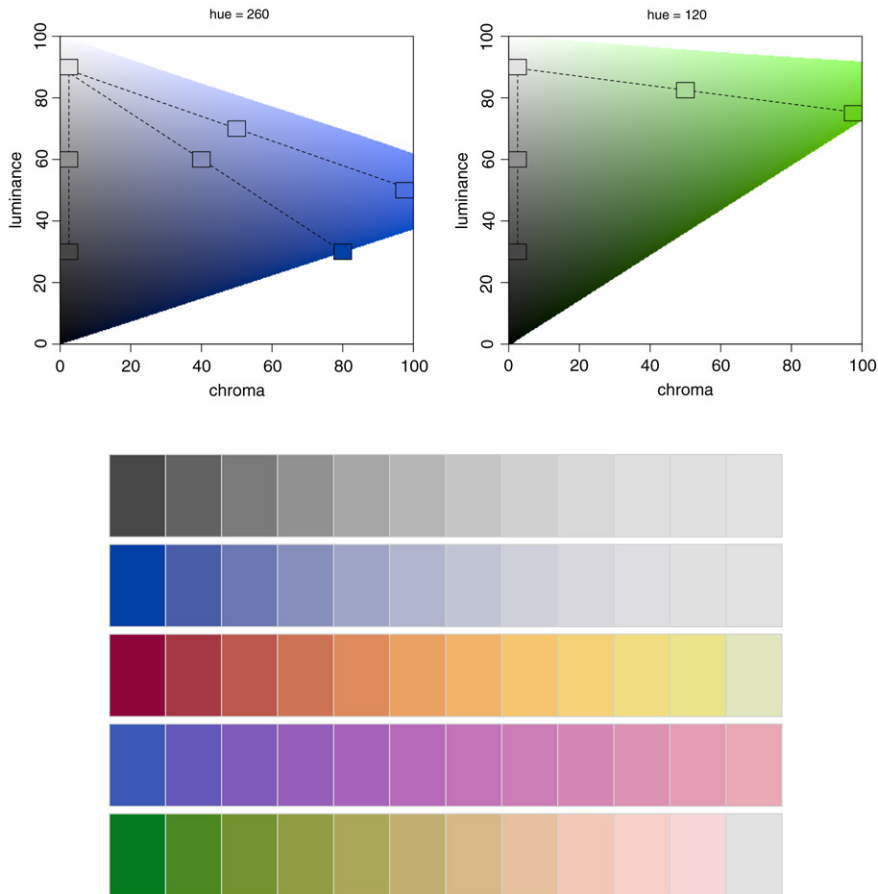
4. Color palettes

4.1. Qualitative palettes

Qualitative palettes are sets of colors for depicting different categories, i.e., for coding a categorical variable. Usually, these should give the same perceptual weight to each category so that no group is perceived to be larger or more important than any other one. Typical applications of qualitative palettes in statistics would be bar plots (see Ihaka, 2003), pie charts or highlighted mosaic displays (see Section 5).

Ihaka (2003) describes a simple strategy for choosing such palettes: chroma and luminance are kept fixed and only the hue is varied for obtaining different colors which are consequently all balanced towards the same gray. This effect is illustrated in Fig. 3 which shows a color wheel obtained by varying the hue only in HSV coordinates (H, 100, 100) and HCL coordinates (H, 50, 70). Clearly, not only the hue but also the amount of chroma and luminance varies for the HSV wheel.

Fig. 4 (top panel) depicts how the HCL-based color wheel is constructed. It shows the hue/chroma plane of HCL space given a luminance of $L = 70$. Not all combinations of hue and chroma are admissible, however for a chroma of $C = 50$ a full color wheel can be obtained. For choosing the hues in a certain palette, various strategies are conceivable. A simple and intuitive one is to use colors as metaphors for categories (e.g., for political parties), another approach would be to use segments from the color wheel corresponding to nearby or distant colors. The latter is shown in Fig. 4 (bottom panel) which depicts examples for generating qualitative sets of colors (H, 50, 70). In the upper left panel colors from the full spectrum are used ($H = 30, 120, 210, 300$) creating a 'dynamic' set of colors. The upper right panel shows a 'harmonic' set with $H = 60, 120, 180, 240$. Cold colors (from the blue/green part of the spectrum: $H = 270, 230, 190, 150$) and warm colors (from the yellow/red part of the spectrum: $H = 90, 50, 10, 330$) are shown in the lower left and right panel, respectively.

Fig. 5. Sequential palettes: construction and examples. Top: Construction is shown in the chroma/luminance plane for two hues $H = 260$ (left) and $H = 120$ (right). Colors are chosen by varying either only luminance or both luminance and chroma. Bottom: Examples by varying only luminance (first row), chroma and luminance (second row), and hue, chroma and luminance (remaining rows).

4.2. Sequential palettes

Sequential palettes are used for coding numerical information that ranges in a certain interval where low values are considered to be uninteresting and high values are interesting. Suppose we need to visualize an intensity or interestingness i which (without loss of generality) is scaled to the unit interval. A typical application in statistics are heatmaps (see Fig. 1).

The simplest solution to this task is to employ light/dark contrasts (Levkowitz and Herman, 1992), i.e., rely on the most basic perceptual axis. The interestingness can be coded by an increasing amount of gray corresponding to decreasing luminance in HCL space:

$$(H, 0, 90 - i \cdot 60),$$

where the hue H used does not matter, chroma is set to 0 (i.e., no color), and luminance ranges in [30, 90] avoiding the extreme colors white ($L = 100$) and black ($L = 0$). Instead of going linearly from light to dark gray, luminance could also be increased nonlinearly, e.g., by some function $i' = f(i)$ that controls whether luminance is increased quickly with intensity or not. We found $f(i) = i^p$ to be a convenient transformation where the power p can be varied to achieve different degrees of non-linearity.

Furthermore, the intensity i could additionally be coded by colorfulness (chroma), e.g.,

$$(H, 0 + i' \cdot C_{max}, L_{max} - i' \cdot (L_{max} - L_{min})).$$

This strategy is depicted in the top left panel of Fig. 5 for a blue hue $H = 260$ with $L_{max} = 90$ and different combinations of maximal chroma ($C_{max} = 0, 80$ and 100, respecitvely) and minimal luminance ($L_{min} = 30, 30$ and 50, respectively). The first two combinations are also shown in the first two rows of the bottom panel in Fig. 5. The top right panel of Fig. 5 shows that the exact same strategy is not possible for the green hue $H = 120$. While the gray colors without chroma can be chosen in the same way, there is a stronger trade-off between using dark colors (with low luminance) and colorful colors (with high chroma). Hence, the second path from light gray to full green ends at a much lighter color with $L_{min} = 75$.
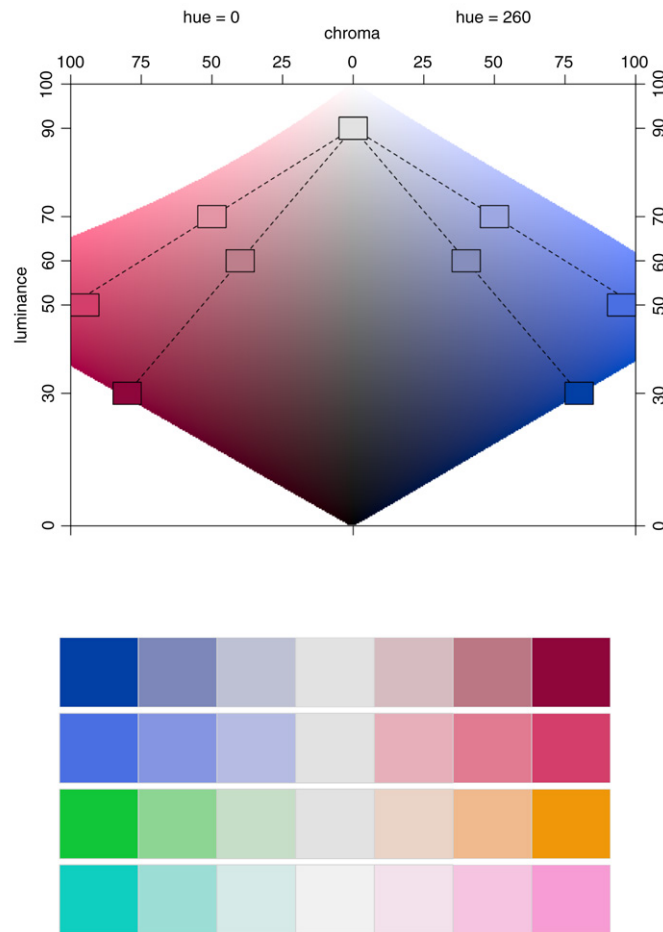
Fig. 6. Diverging palettes: construction and examples. Top: Construction is shown in the chroma/luminance plane for hues $H = 0$ and $H = 260$ (back to back). Colors are chosen by simultaneously varying luminance and chroma. Bottom: Examples with different pairs of hues and decreasing luminance contrasts.

To increase the contrast between the colors in the palette even further, the ideas from the previous sequential palettes can also be combined with qualitative palettes by simultaneously varying the hue as well:

$$(H_2 - i \cdot (H_1 - H_2), C_{max} - i' \cdot (C_{max} - C_{min}), L_{max} - i'' \cdot (L_{max} - L_{min})).$$

One application is an HCL-based version of "heat colors" that increase from a light yellow (e.g., $(90, 30, 90)$) to a full red (e.g., $(0, 100, 50)$). To make the change in hue visible, the chroma needs to increase rather quickly for low values of i and then only slowly for higher values of i. This can be achieved by choosing an appropriate transformation i' for chroma and a different transformation i'' for the luminance. Such a strategy is adopted for the palettes shown in the last three rows in Fig. 5 using different pairs of hues as well as different chroma and luminance contrasts.

4.3. Diverging palettes

Diverging palettes are also used for coding numerical information ranging in a certain interval—however, this interval includes a neutral value. Examples for this include residuals or correlations (both with the neutral value 0) or binary classification probabilities (with neutral value 0.5) that could be visualized in mosaic plots (e.g., as in Zeileis et al., 2007, Fig. 3), classification maps (e.g., as in Tenenhaus et al., 2007, Fig. 2) or model-based shading in maps such as Fig. 2. Analogously to the previous section, we suppose that we want to visualize an intensity or interestingness i from the interval $[-1, 1]$ (without loss of generality).

Given useful sequential palettes, deriving diverging palettes is easy: two different hues are chosen for adding color to the same amount of 'gray' at a given intensity $|i|$. Fig. 6 (top panel) shows the chroma/luminance plane back to back for the hues $H = 0$ and 260 with two different paths—giving slightly different emphasis on luminance or chroma contrasts—from a full red over a neutral gray to a full blue. As the top panels of Fig. 5 illustrates, the pair of hues should be chosen carefully because the admissible values in the chrome/luminance plane differ across hues. Clearly, for deriving symmetric palettes,
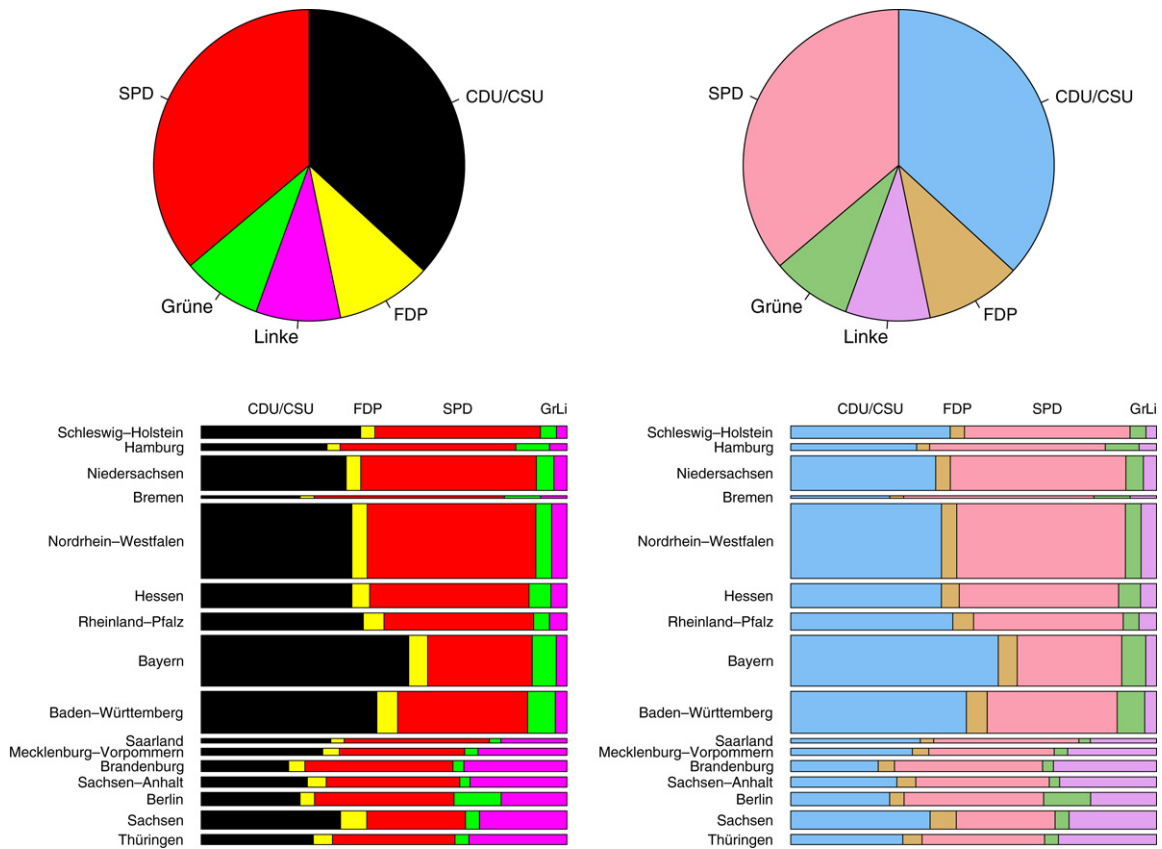
Fig. 7. German election 2005 with HSV-based (left) and HCL-based (right) qualitative palette. Top: Pie chart for seats in the parliament. Bottom: Mosaic display for votes by province.

only colors from the intersection of the admissible chroma/luminance planes can be used. The particular hues $H = 0$ and 260 used in Fig. 6 (top panel) were chosen because they correspond to similar geometric shapes in the chroma/luminance plane, allowing for both large chroma and luminance contrasts. If potential viewers of the resulting graphic might be color-blind, the pair of hues should be taken from the yellow/blue axis of the color wheel rather than the green/red axis as contrasts on the latter axis are more difficult to distinguish for color-blind people (Lumley, 2006).

Fig. 6 (bottom panel) shows various examples of conceivable combinations of hue, chroma and luminance. The first palette uses a broader range on the luminance axis whereas the others mostly rely on chroma contrasts.

5. Illustrations

This section provides several more demonstrations of how HCL-based color palettes can be used to produce statistical plots that are less "flashy" and more perceptually balanced, so that the viewer can extract information from the plots more easily and more accurately.

Although problematic for many tasks, pie charts can be useful for visualizing whether a set of pie segments constitutes a majority. A typical application is shown in the top row of Fig. 7, visualizing the distribution of seats in the German parliament "Bundestag" following the 2005 election. In this election, five parties were able to obtain enough votes to enter the Bundestag—however, neither the governing coalition of SPD and Grüne nor the opposition of CDU/CSU and FDP could assemble a majority. Given that no party would enter a coalition with the leftists "Die Linke", this lead to a big coalition of CDU/CSU and SPD. In graphical displays, the parties are usually matched by using colors as metaphors: red for the social democrats SPD, black for the conservative CDU/CSU, yellow for the liberal FDP, green for the green party "Die Grünen" and purple for the leftist party "Die Linke". The left panel shows fully saturated HSV colors as usually found in the (German) media whereas the right panel uses less flashy HCL colors with the same hues (except for the CDU/CSU where a blue hue instead of the extreme "color" black is used). The advantage of the latter is again that they are easier to hold in focus for a longer time. Furthermore, they are all balanced towards the same gray and have the same amount of color, resulting in a perceptually balanced palette that does not introduce undesired graphical distractions. While it could be argued that this pie chart is such a simple display that a well-balanced palette is not so important and that flashy colors are to be preferred, the balancing properties of the HCL-based palette are very important in more complex displays such as the mosaic display
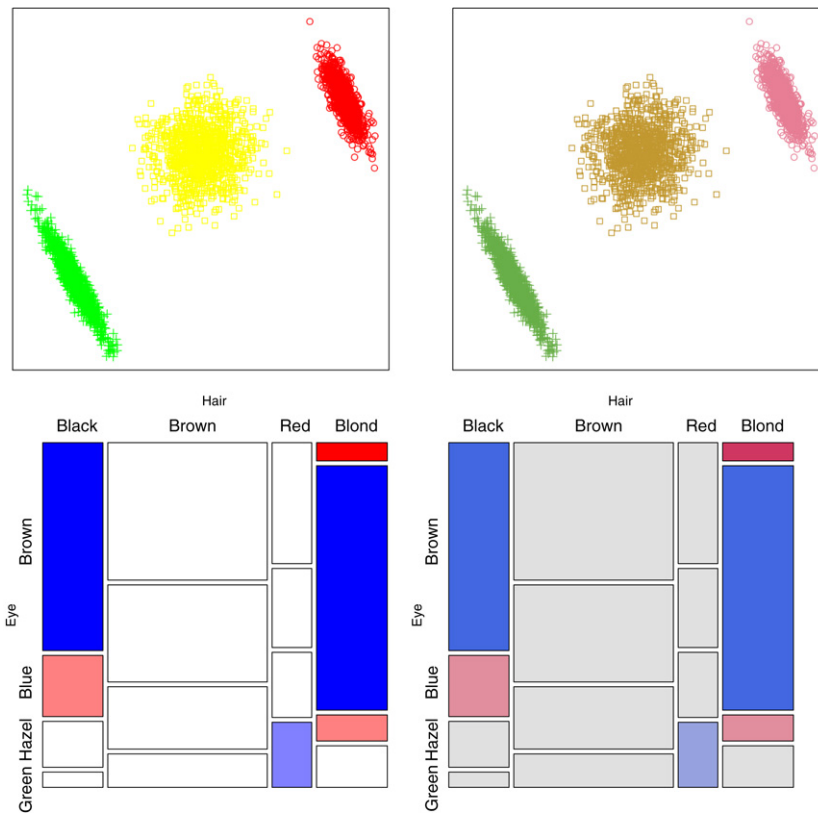
Fig. 8. Further examples for HSV-based (left) and HCL-based (right) palettes. Top: Scatter plot with three clusters and qualitative palette. Bottom: Extended mosaic display for hair and eye color data with diverging palette.

in the bottom row of Fig. 7. This re-uses the same colors and shows the distribution of votes in the 2005 election stratified by province (Bundesland). The order of provinces is from north to south, first for the 10 western provinces (the former Federal Republic of Germany, FRG), then for the 6 eastern provinces (the former German Democratic Republic, GDR). Clearly, the SPD performed better in the north and the CDU/CSU better in the south; furthermore, Die Linke performed particularly well in the eastern provinces and in Saarland.

Color choice is usually much more important in graphical displays with shaded areas compared to displays with only points or lines. However, color choice becomes much more relevant also in scatter plots when there are many points. As an example, the top row in Fig. 8 depicts a scatter plot with three clusters similar to the one shown in Celeux et al. (2000, Fig. 3). To indicate cluster membership, three different colors (green, yellow, red) are employed. The HSV colors in the left panel are again very flashy and differ strongly with respect to luminance: the yellow is much lighter and hardly visible. In contrast, the HCL colors in the right panel use the same hues, but are balanced with respect to chroma and luminance, i.e., the amount of color and gray.

As our final example, the bottom row in Fig. 8 (taken from Friendly, 2002, Fig. 11, left/middle) visualizes the cross-tabulation of hair end eye color of 592 students in a mosaic display. Clearly, hair and eye color are not independent and the pattern of association is highlighted by means of residual-based shading. Cells associated with Pearson residuals whose absolute value exceeds (2 or) 4 are shaded (light) blue/red. This shows that there are significantly more students with black hair and brown eyes, blond hair and blue eyes, red hair and green eyes than expected under independence. Conversely, fewer students than expected have blond hair and brown/hazel eyes or black hair and blue eyes. Comparing the HSV and HCL colors in the bottom row of Fig. 8, it is shown again that the HCL colors are better balanced (between the red and blue colors) and less flashy. Even the cells associated with small residuals (below 2) are somewhat easier to read when shaded in a light grey rather than white. More details and extensions to residual-based shadings can be found in Zeileis et al. (2007).

6. Software

Implementing the different color palettes suggested in the previous section is extremely easy if the software environment chosen already provides an implementation of HCL colors: from the formulas provided above the HCL coordinates for a palette can be conveniently computed. Somewhat more work is required if the software package does not yet provide an HCL implementation. In that case, additional functionality is needed for translating HCL coordinates to the software

package's color system which may vary between different packages, but standardized RGB (sRGB) is often used. The typical way of coordinate conversion is to go first from HCL to CIELUV by simply transforming the polar H and C coordinates back to the original U and V. Subsequently, CIELUV is converted to CIEXYZ which in turn is converted to sRGB. The details of these conversions are somewhat technical and tedious (and hence omitted here), however the conversion formulas are still straightfoward to implement and can, for example, be found in Wikipedia (2007a) or Poynton (2000).

The R system for statistical computing (R Development Core Team, 2008) provides an open-source implementation of HCL (and other color spaces) in the package **colorspace**, originally written by Ross Ihaka. The coordinate transformations mentioned above are contained in C code within **colorspace** that are easy to port to other statistical software systems. Version 1.0-0 of **colorspace** (Ihaka et al., 2008) also includes an implementation of all palettes discussed above. (Originally, the code for the palettes was in the **vcd** package, Meyer et al. (2006) but it was recently moved to **colorspace** to be more easily acessible.) Qualitative palettes are provided by `rainbow_hcl()` (named after the HSV-based function `rainbow()` in base R). Sequential palettes based on a single hue are implemented in the function `sequential_hcl()` while `heat_hcl()` offers sequential palettes based on a range of hues. Diverging palettes can be obtained by `diverge_hcl()`. Technical documentation along with a large collection of example palettes is available via `help("rainbow_hcl", package = "colorspace")`. Furthermore, R code for reproducing the example palettes in Figs. 4–6 (and some illustrations) can be accessed via `vignette("hcl-colors", package = "colorspace")`.

The default color palettes in the **ggplot2** package (Wickham, 2008) are also based on HCL colors, using similar ideas to those discussed in this article.

7. Discussion

Many statistical graphics—especially when displayed on a computer screen, e.g., as in interactive usage, electronic papers or presentation slides—employ colors to code information about a certain variable. Despite this omnipresence of color, there is often only little guidance in statistical software packages on how to choose a palette appropriate for a particular visualization task—auspicious tools such as **ColorBrewer.org** notwithstanding. We try to address this problem by suggesting color schemes for coding categorical information (qualitative palettes) and numerical information (sequential and diverging palettes) based on the perceptually-based HCL color space.

We provide paths through HCL space along perceptual axes of the human visual system so that colors selected along these paths match human perceptual dimensions. This gives the users the possibility to conveniently experiment with the HCL-based palettes by varying several simple and intuitive graphical parameters. For qualitative palettes, these are the coordinates on the chroma and luminance axis, respectively, controlling whether the colors are light or dark and how colorful they are. For sequential and diverging palettes, the user can decide whether contrasts in the chroma or luminance direction (or both) should be employed. In our experience (as illustrated in Section 2), chroma contrasts work sufficiently well if a small set of colors is used. However, when a larger set of colors is used (e.g., for heatmaps where extreme values should be identifiable) it is much more important to have a big difference in luminance.

Based on these conceptual guidelines and the computational tools readily provided in the R system for statistical computing (and easily implemented in other statistical software packages), users can generate palettes varying these graphical parameters and thus adapting the colors to their particular graphical display.

Acknowledgments

References

Azzalini, A., Bowman, A.W., 1990. A look at some data on the old faithful geyser. Applied Statistics 39, 357–365.

Brewer, C.A., 1999. Color use guidelines for data representation. In: Proceedings of the Section on Statistical Graphics, American Statistical Association, Alexandria, VA. pp. 55–60. URL: http://www.personal.psu.edu/faculty/c/a/cab38/ColorSch/ASApaper.html.

Celeux, G., Hurn, M., Robert, C.P., 2000. Computational and inferential difficulties with mixture posterior distributions. Journal of the American Statistical Association 95 (451), 957–970 (Fig. 3).

Cleveland, W.S., McGill, R., 1983. A color-caused optical illusion on a statistical graph. The American Statistician 37, 101–105.

Commission Internationale de l'Éclairage, 2004. Colorimetry, 3rd edition. Publication CIE, Vienna, Austria, ISBN: 3-901-90633-9, 15, 2004.

Friendly, M., 2002. A brief history of the mosaic display. Journal of Computational and Graphical Statistics 11 (1), 89–107 Fig. 11 (left, middle).

Gelfand, A.E., Banerjee, S., Sirmans, C.F., Tu, Y., Ong, S.E., 2007. Multilevel modeling using spatial processes: Application to the Singapore housing market. Computational Statistics & Data Analysis 51, 3567–3579, doi:10.1016/j.csda.2006.11.019. Fig. 3.

Harezlak, J., Coull, B.A., Laird, N.M., Magari, S.R., Christiani, D.C., 2007. Penalized solutions to functional regression problems. Computational Statistics & Data Analysis 51, 4911–4925, doi:10.1016/j.csda.2006.09.034. Fig. 8.

Harrower, M.A., Brewer, C.A., 2003. **ColorBrewer.org**: An online tool for selecting color schemes for maps. The Cartographic Journal 40, 27–37. URL: http://ColorBrewer.org/.

Ihaka, R., 2003. Colour for presentation graphics. In: Hornik, K., Leisch, F., Zeileis, A. (Eds.), Proceedings of the 3rd International Workshop on Distributed Statistical Computing, Vienna, Austria, ISSN 1609-395X, URL: http://www.ci.tuwien.ac.at/Conferences/DSC-2003/Proceedings/.

Ihaka, R., Murrell, P., Hornik, K., Zeileis, A., 2008. **colorspace**: Color Space Manipulation. R package version 1.0-0, URL: http://CRAN.R-project.org/package=colorspace.

Kaiser, P.K., Boynton, R.M., 1996. Human Color Vision, 2nd edition.. Optical Society of America, Washington, DC.

Kneib, T., 2006. Mixed model-based inference in geoadditive hazard regression for interval-censored survival times. Computational Statistics & Data Analysis 51, 777–792, doi:10.1016/j.csda.2006.06.019. Fig. 5 (left).

Knoblauch, K., 2002. Color vision. In: Yantis, S., Pashler, H. (Eds.), Steven's Handbook of Experimental Psychology – Sensation and Perception volume 1, third edition. John Wiley & Sons, New York, pp. 41–75.

Levkowitz, H., Herman, G.T., 1992. Color scales for image data. IEEE Computer Graphics and Applications 12 (1), 72–80, doi:10.1109/38.135886.

Lumley, T., 2006. Color coding and color blindness in statistical graphics. ASA Statistical Computing & Graphics Newsletter 17 (2), 4–7. URL: http://stat-graphics.org/newsletter/v172.pdf.

Meier, B.J., Spalter, A.M., Karelitz, D.B., 2004. Interactive color palette tools. IEEE Computer Graphics and Applications 24 (3), 64–72. URL: http://graphics.cs.brown.edu/research/color/.

Meyer, D., Zeileis, A., Hornik, K., 2006. The strucplot framework: Visualizing multi-way contingency tables with vcd. Journal of Statistical Software 17 (3), 1–48. URL: http://www.jstatsoft.org/v17/i03/.

Moretti, G., Lyons, P., 2002. Tools for the selection of colour palettes. In: Proceedings of the New Zealand Symposium On Computer-Human Interaction, SIGCHI 2002, University of Waikato, New Zealand.

Munsell, A.H., 1905. A Color Notation. Munsell Color Company, Boston, Massachusetts.

Poynton, C., 2000. Frequently-asked questions about color. URL: http://www.poynton.com/ColorFAQ.html (accessed 06.11.07).

R Development Core Team 2008. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN: 3-900051-07-0, URL: http://www.R-project.org/.

Smith, A.R., 1978. Color gamut transform pairs. Computer Graphics 12 (3), 12–19. ACM SIGGRAPH 78 Conference Proceedings, URL: http://www.alvyray.com/.

Tenenhaus, A., Giron, A., Viennet, E., Béra, M., Saporta, G., Fertil, B., 2007. Kernel logistic PLS: A tool for supervised nonlinear dimensionality reduction and binary classification. Computational Statistics & Data Analysis 51, 4083–4100, doi:10.1016/j.csda.2007.01.004. Fig. 2.

Tufte, E.R., 1990. Envisioning Information. Graphics Press, Cheshire, CT.

Ware, C., 1988. Color sequences for univariate maps: Theory, experiments and principles. IEEE Computer Graphics and Applications 8 (5), 41–49, doi:10.1109/38.7760.

Wickham, H., 2008. ***ggplot2***: An Implementation of the Grammar of Graphics. R package version 0.7, URL: http://CRAN.R-project.org/package=ggplot2.

Wikipedia 2007a. CIELUV Color Space — Wikipedia, The Free Encyclopedia. URL: http://en.wikipedia.org/wiki/CIELUV_color_space (accessed 06.11.07).

Wikipedia 2007b. HSV Color Space — Wikipedia, The Free Encyclopedia. URL: http://en.wikipedia.org/wiki/HSV_color_space (accessed 06.11.07).

Zeileis, A., Meyer, D., Hornik, K., 2007. Residual-based shadings for visualizing (conditional) independence. Journal of Computational and Graphical Statistics 16 (3), 507–525, doi:10.1198/106186007X237856.

Zhang, H., Montag, E.D., 2006. How well can people use different color attributes? Color Research & Application 31 (6), 445–457, doi:10.1002/col.20257.