Energy Aware Color Sets

Johnson Chuang¹, Daniel Weiskopf² and Torsten Möller¹

¹School of Computing Science, Simon Fraser University, Burnaby, Canada ²VISUS, Universität Stuttgart, Germany

Abstract

We present a design technique for colors with the purpose of lowering the energy consumption of the display device. Our approach is based on a screen space variant energy model. The result of our design is a set of distinguishable iso-lightness colors guided by perceptual principles. We present two variations of our approach. One is based on a set of discrete user-named (categorical) colors, which are analyzed according to their energy consumption. The second is based on the constrained continuous optimization of color energy in the perceptually uniform CIELAB color space. We quantitatively compare our two approaches with a traditional choice of colors, demonstrating that we typically save approximately 40 percent of the energy. The color sets are applied to examples from the 2D visualization of nominal data and volume rendering of 3D scalar fields.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.3]: Picture/Image Generation—Computer Graphics [I.3.6]: Methodology and Techniques—Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism—

1. Introduction

An increasing use of electrical energy in today's technologically-driven world is leading to a growth in greenhouse gas emissions [Mur08]. As the world becomes increasingly aware of its environment, there is also a trend for IT industries to look into greener computing [WC08]. Green computing encourages the practice of using computing resources efficiently in such a way as to create minimal impact on the environment. A key green objective, therefore, is to reduce the energy consumption of using computers and their related subsystems.

The mobile device industry, one of the fastest growing in consumer electronics, is also seeing a trend to go green. The display is one of the major power consumers in modern computer systems, using up to 38 percent of the total system power in PCs and up to 50 percent in mobile devices [MM05, She08]. Thin film transistor liquid crystal displays (TFT LCDs) currently dominate the display market. One of the main components of a TFT LCD screen is the display backlight, which is often regarded as the primary consumer of display energy. Previous studies show that TFT LCDs are not very energy efficient as the electrical-to-light energy conversion has losses exceeding 80 percent [IP05]. Energy efficiency especially is a priority for mobile devices because they are typically powered by batteries and have limited energy resources. Furthermore, there is a trend for battery capacities to be increasing at a much slower rate than the energy requirements for mobile applications [MLC06]. Emerging display technologies such as organic light-emitting diode (OLED) displays promise to be more energy efficient than their TFT LCD counterpart [For03]. With the increased sophistication and rising energy demands of multimedia applications, it is advantageous to explore energy saving techniques for these upcoming displays.

Compute-intensive applications like volume visualization, previously feasible only on desktop machines, are slowly making their way to mobile devices [MW08]. One of the fundamental ways to visualizing and communicating complex information is by labeling data using colors. For example, in volume visualization, users often assign colors to voxels in tasks like grouping and labeling of data. Further, labeling different regions on a map using color is common in order to study and communicate weather data, census data, or a number of other geospatial information. Previous work does not consider the energy cost of colors for displays.

^{© 2008} The Author(s)

Journal compilation © 2008 The Eurographics Association and Blackwell Publishing Ltd. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

Color, however, plays an important role in energy consumption for emerging display technologies.

The goal of our paper is to spark awareness of energy consumption of emerging color displays. Although there are various applications that we can consider, from graphical user interfaces to mobile games, we focus on color design for data visualization. To the best of our knowledge this is the first attempt in the visualization and graphics community to make energy consumption a design constraint. In this paper we specifically study the effect of color on energy cost in scientific applications, like volume visualization, to more general applications like color labeling of maps. We present two design techniques for colors that lower energy consumption. Both approaches apply an optimization process to minimize energy, and they take into account different perceptual aspects such as perceptual color differences or categorization and naming of colors. Our work could benefit emerging display technologies, for example, by extending the battery life of mobile devices and lowering the energy cost of desktop machines.

2. Background

2.1. Power Management and Energy Awareness

In modern desktop computers and mobile devices, the display subsystem is one of the major consumers of electrical energy. For example, TFT LCDs, which dominate the current display market, typically consume 30 to 50 percent of battery power on mobile phones [She08]. Most of that power is needed by a cold cathode fluorescent lamp (CCFL) which provides the necessary background lighting to illuminate the LCD. Most previous work, therefore, reduces energy consumption by decreasing the CCFL backlight, and compensate for the reduced LCD luminance by increasing its transmittance [TLC07, CC06, CSC02].

The key to enabling an energy efficient display design is to facilitate variable power output across individual subportions of the screen [HVG*04]. The main stream of current TFT LCD technology does not support that, as they typically have a single backlight illuminating the whole screen. In fact, backlit TFT LCDs are not very efficient because polarizers and color filters absorb most of that backlight. One emerging technology, the OLED display, promises to be more energy efficient than LCDs. It is emissive, requires no backlight, and allows for variable display power in different regions of the screen [ILMR03]. Researchers have already started looking into energy-aware application interfaces for such energy-adaptive displays [RGMN06]. The technological role model for our approach is a line of emerging energy adaptive displays-e.g., Sony's OLED displays, Samsung's 9-series "local dimming" TVs, and even high dynamic range (HDR) displays that use variable, space-variant background lighting in combination with a high-resolution TFT LCD panel [SHS*04]. This kind of HDR display is expected to hit the markets in the near future, for example, in

the form of Dolby HDR Video technology, which originated from BrightSide technologies [Dol08].

2.2. Color Maps and Human Visual Perception

Mapping scalar values to color (facilitated via color maps) is one of the most common approaches to data visualization. Two types of color maps can be distinguished: discrete and continuous. A discrete color map is a one-to-one mapping of discrete data values to color, and is often used for the visualization of nominal (categorical) data. Continuous color maps, on the other hand, define a continuous color range.

Color information arriving at the human eye is split into 3 channels: 1 achromatic and 2 chromatic [KSJ95]. The achromatic channel carries luminance information and, amongst others, is responsible for motion detection and shape perception. In contrast, the chromatic channels carry color information useful for visual grouping and labeling. Since we focus on discrete color maps for labeling and grouping, we target colors that are constant in lightness and only vary in color. At this point, we make the distinction between lightness and luminance. Lightness (typically denoted L^*) is the perceived brightness of a color subject to the human perceptual system. On the other hand, luminance (denoted L) is the measurable physical quantity (in cd/m^2) which most closely corresponds to lightness. Iso-lightness colors avoid unwanted effects related to non-iso-lightness color maps for nominal data, such as emphasizing certain color-mapped regions, which might be caused by different lightness, and introducing a natural ordering according to lightness.

The design of effective color maps is a difficult task and often requires multiple iterations. In most applications, users have the freedom to manually specify their own color maps or choose from an existing set of predefined ones. Although using a predefined color map may be convenient, it may lead to faulty interpretations of the underlying data by introducing artifacts or non-existent features. For example, the popular rainbow color map can be confusing, obscuring, and actively misleading [BT07]. There are tools to ease the creation of effective color maps. For example, the ColorBrewer system uses a set of guidelines to generate a collection of color maps for discrete data [HB03]. Earlier approaches to color map design include effective color maps based on color distance and color category [Hea96], dynamic color map creation for data exploration [RT90], and guidelines for univariate colors [War88]. The particular problem of specifying iso-lightness color maps, which are normally difficult to create without a calibrated monitor, can be addressed using a face-based approach [KRC02].

All of the previous work on color map design and color perception neglects energy requirements of the display. This comes as no surprise because older display technologies (CRTs, LCDs, etc.) have approximately constant energy cost for displaying any color. Our work specifically applies to



Figure 1: The energy required to display this tooth visualization (left) is calculated by using the maxima of the R, G, and B values within each tile. The right image illustrates a partitioning of the screen into 10×10 tiles, with shades of gray going from black to white representing the maximal R, G, B portion within that tile, estimating the energy cost.

emerging display technologies that support variable energy consumption across the screen. Therefore, this paper combines existing knowledge of perceptually motivated color map design with new low-energy requirements.

3. Screen Space Variant Energy Model

Displays with space variant control of light energy on screen, such as HDR displays [SHS*04] or OLED displays, are the basis for our model of energy estimation. Emerging display technologies using non-uniform backlights (such as HDRs) make it possible for varying power consumption across the screen, and as a side effect lowered energy cost. We model the non-uniformity of this kind of backlighting as a set of non-overlapping tiles covering the screen; the tiles are arranged in a rectangular grid. Each tile can be seen as an independent backlight responsible for providing adequate lighting for only the pixels within it. Although other grid layouts such as hexagonal configurations for LED-driven displays or even overlapping "tiles" might appear in physical devices, the qualitative results will be similar as for the rectangular grid approach.

Now let us look at the energy requirements for one tile. The white background lighting is the main consumer of energy. White light is filtered through the TFT LCD panel to produce different amounts of red, green, and blue colors. In general, the energy requirement is proportional to the number of "on" pixels and the brightness of their R, G, and B components [ILMR03, VZJ06]. The optimal background lighting only needs to be as large as the maximum of the linear, non-gamma-corrected R, G, or B values within that tile. We use linear RGB (instead of gamma-corrected RGB, or sRGB) because they are a good approximation of the actual radiative power emitted from the display [CC06].

Figure 1 illustrates our approach of calculating the optimal background lighting. We partition the screen into a lowresolution $w \times h$ grid of tiles. Within each tile *i* we have pix-

© 2008 The Author(s) Journal compilation © 2008 The Eurographics Association and Blackwell Publishing Ltd.

els *P* with red, green, and blue values in the range of 0 to 1. The maximum color components from each tile are summed up and normalized, to obtain a relative measure of energy E_{display} between 0 and 1:

$$E_{\text{display}} = \frac{1}{wh} \sum_{i=1}^{wh} \max_{P \in \text{Tile } i} (\max(P_{\text{red}}, P_{\text{green}}, P_{\text{blue}}))$$
(1)

Equation 1 applies to displays that have several independent backlight panels (e.g., HDR displays as modelled in Figure 1). For an OLED display with separate R, G, and B primaries, a more appropriate measure of energy may be the sum of the RGB components.

4. Discrete Optimization of Energy Aware Colors

The goal of this paper is to determine a set of iso-lightness colors that are (1) easy to distinguish and (2) associated with low energy when displayed. We formulate this goal as the optimization of display energy under the constraint of good perceptual distinguishability.

Our first optimization approach is discrete: from a set of M iso-lightness and distinguishable colors, a subset of N colors ($N \le M$) is chosen so that the sum of the energies associated with the chosen colors is the minimum of any subset of N colors. The optimization algorithm is simple. First, the energy of each color is computed. Second, colors are sorted according to the energy in ascending order. Third, the first N colors are picked from the sorted list.

The open question is how the set of M adequate input colors is chosen. We propose to adopt named (categorical) colors because they are sufficiently distinct and they exhibit proven perceptual benefits. For example, Kawai et al. [KUU95] suggest that the time it takes to distinguish multiple colors depends partly on their named color region. In a user study, Healey [Hea96] confirms that the name of the color region, or color category, indeed affects the performance of color target identification. Our discrete design technique for low energy colors makes use of this knowledge and examines color categories and their energy performance. We further constrain our design to iso-lightness colors, which are not automatically guaranteed by named colors.

For the color computations, we employ a perceptual color model that has the achromatic and chromatic color dimensions decoupled. This allows us to use the achromatic part for lightness, and the chromatic part for color category. The CIELAB color space [Fai06] is a widely-used color space designed to be perceptually uniform. It has an achromatic axis for lightness L^* and two chromatic axes for red-green and yellow-blue color components a^* , b^* . CIELAB is the basic color model of this paper—despite a few shortcomings compared with more advanced color appearance models [Fai06].



Figure 2: Top: fully-saturated HSL colors before transformation. Bottom: similar to top, except now the HSL lightness has been transformed to match the CIELAB lightness. The colors with HSL lightness > 0.5 are less saturated because otherwise they would be out of the monitor's gamut.

While CIELAB can aid in choosing iso-lightness colors, it lacks support for directly specifying hue, making it difficult to choose specific color tones. LCH_{ab}, which specifies CIELAB colors using cylindrical coordinates (chroma and hue), can be used to predict color hues in CIELAB [Fai06]. However, because of the lack of hue uniformity, some hue slices in LCH_{ab} actually contain slightly varying color hues. For example, a slice from the blue hue plane may contain tints of purple. Therefore, as an alternative, the HSL color space can be applied to select hues. The only downside is that it is not perceptually uniform, so iso-lightness color picking in HSL is not straightforward. Our solution is to transform the lightness axis in HSL to match the L^* axis in CIELAB. This gives us HSL_{LAB}, which has 3 control parameters: hue, lightness, and saturation. Figure 2 shows a slice from a fully-saturated HSL before and after transformation.

The number of easily distinguishable colors is small. For example, in his study of categorical colors, Healey [Hea96] finds that 7 distinct iso-lightness colors is the maximum number of colors that can be displayed at one time without lowering the response time and accuracy of target color identification. He provides percentages of how likely observers will give a particular color name to a displayed Munsell color. Based on his findings, the color names with high percentages of being assigned to a Munsell color are ranked (in decreasing order): green, blue, orange, purple, red, yellow, aqua, pink, brown, and magenta. We choose the 6 most highly ranked colors (green, blue, orange, purple, red, yellow) with HSLLAB hue angles 120, 240, 30, 292, 0, and 60 degrees, respectively. Saturation is set to maximum to facilitate good distinguishability. Lightness may be changed according to user preferences. Figure 3 shows the energy con-



Figure 3: Energy plot of the hues of 6 categorical colors (from left to right: blue, red, purple, orange, green, yellow).



Figure 4: Categorical colors of varying lightness sorted by increasing energy cost. Please note that the images of this paper are designed for sRGB monitor displays, not for good reproduction in print.

sumption of these 6 colors with respect to increasing lightness. The energy of a color is the maximum of its linear R, G, or B values, as described in Section 3. Finally we sort the colors by increasing energy cost to obtain the energy aware categorical color palette in Figure 4. Users can pick distinguishable iso-lightness colors with increasing energy cost by choosing colors from bottom to top.

While the named colors provide a reasonable choice of energy aware colors, they are not optimal when the whole available color space is considered. First, the hues of the categorical colors are fixed and, therefore, the optimization cannot shift the colors around in color space to achieve better perceptual distinguishability or lower energy consumption if not all M colors are needed. Second, the categorical colors are not completely evenly distributed along the boundary of the color gamut. Figure 5(a) illustrates that, for example, the perceptual distance between blue and green is much larger than between green and yellow. Even color placement would lead to better distinguishability. Third, the issue of the fixed hues is aggravated for high lightness-when colors are close to the upper lightness limit; then, colors tend to become less saturated (shifted towards white, due to the restricted gamut of the monitor), as demonstrated in Figure 5(b).



Figure 5: In both CIELAB iso-lightness color slices, the white lines are representative LCH_{ab} hue angles for (clockwise) green, yellow, orange, red, purple, and blue. (a) Medium lightness green, yellow, and orange, although low in energy, are not very distinguishable from each other because of their relative closeness. (b) High lightness orange, red, purple, and blue have very low perceptual difference and tend towards white.

5. Continuous Optimization of Color Energy in Perceptual Color Space

In the search for energy aware categorical colors, we constrained ourselves to fully-saturated colors from nonoverlapping named color regions. The constraints essentially restricted us to a single level of saturation and a small subset of all possible hues. We now loosen those constraints on saturation and hue, and proceed to find a set of energy aware, distinguishable, iso-lightness colors.

A perceptual color space like CIELAB fits our needs for distinguishability and iso-lightness. CIELAB is designed to be perceptually uniform, meaning that a perturbation to a component value produces an equal magnitude of change in visual difference. The perceived color difference between any two colors in CIELAB can therefore be computed by treating the $L^*a^*b^*$ components as points in 3D space and taking the Euclidean distance between them [Fai06]. Although CIELAB is not the most advanced color system and does not incorporate sophisticated color-appearance effects, it is a well-established perceptually uniform color system. The following discussion is not restricted to CIELAB and could be immediately adopted for any other color system that supports the computation of a measure of lightness and of perceptual differences.

We formulate our goal of finding energy aware, distinguishable, iso-lightness colors as an optimization problem with 3 input parameters: the number of colors N, the level of lightness L^* , and a minimum perceptual color distance dthat must be enforced between every pair of colors. The rationale is that this minimum distance is a means of ensuring distinguishability of the colors. Typically, d is provided by the user who wants to achieve a certain separation of colors.

The goal is to minimize the maximum display energy un-

© 2008 The Author(s) Journal compilation © 2008 The Eurographics Association and Blackwell Publishing Ltd



Figure 6: (a) $L^* = 65$ iso-lightness slice from CIELAB. (b) An energy heightfield showing the energy consumption of the colors; the green-yellow region tends to have lower energy.

der the constraint that colors are at least a distance *d* apart from each other. Following a soft constraint implementation (i.e. penalizing colors that are closer than distance *d*), we propose the following cost function E_{cost} to measure the relative energy of *N* colors chosen from the iso-lightness color slice. Let $\mathbf{C} = \{C_1, C_2, ..., C_N\}$ be the *N* colors, then the cost is

$$E_{\text{cost}}(\mathbf{C}) = E_{\max}(\mathbf{C}) + k \sum_{\substack{C_i, C_j \in \mathbf{C} \\ C_i \neq C_i}} E_{\text{penalty}}(C_i, C_j) \quad (2)$$

where

$$E_{\max}(\mathbf{C}) = \max_{1 \le i \le N} (\max(C_{i, \text{red}}, C_{i, \text{green}}, C_{i, \text{blue}})) \quad (3)$$

and

$$E_{\text{penalty}}(C_i, C_j) = \begin{cases} 1 - \frac{1}{d} dist(C_i, C_j) & \text{if } dist(C_i, C_j) < d \\ 0 & \text{otherwise} \end{cases}$$
(4)

 E_{max} finds among all the colors the maximum linear R, G, or B value (recall that the maximum RGB value is a measure of energy). The terms $C_{i,\text{red}}$, $C_{i,\text{green}}$, and $C_{i,\text{blue}}$ denote the R, G, and B components of color C_i , respectively. The function E_{penalty} penalizes two colors if they are too close, i.e. less than *d* apart. The perceptual distance is denoted $dist(C_i, C_j)$. Here, a linear penalty function is used, although nonlinear functions may work as well. The constant *k* determines the relative importance E_{max} and E_{penalty} , describing the relative increase of the penalty energy with decreasing color distance. In our optimization, we let *k* be equal to the maximum allowable RGB value, or k = 1.

Figure 6 illustrates the $L^* = 65$ color slice from CIELAB, and its energy heightfield. The optimization problem can be considered as finding N points on the surface of the energy heightfield, such that the 3D Euclidean distance between all points is at least d, and the maximum height among all points is minimized. From Figure 6, it can be expected that green and yellow colors use less energy because they are lower



Figure 7: (a) Optimizing for N = 7 colors, iso-lightness level of $L^* = 65$, and minimum color distance of d = 30. (b) Optimizing for N = 7, $L^* = 45$, d = 50. The top-row images show the a^* - b^* plane, the bottom-row shows the color sets.

in the heightfield. Figure 3 shows the same qualitative behavior. Moreover, the energy heightfield and its boundaries, which correspond to the gamut of displayable colors, are quite smooth; thus, optimization is expected to pose no particular difficulties.

The optimization process minimizes E_{cost} with respect to **C**. Since L^* is fixed, each C_i has two components a_i^* and b_i^* (from the CIELAB representation of C_i), for a total of 2N variables. With this choice of variables, the hard constraint of constant L^* is always guaranteed. For optimization, we apply the Nelder and Mead method [NM65], also known as the downhill simplex method or amoeba method. Other optimization methods for non-linear multi-variable functions may be used alternatively. We apply the optimization method to minimize $E_{\text{cost}}(a_1^*, b_1^*, ..., a_N^*, b_N^*)$. The downhill simplex method works iteratively, starting from an initial configuration. We randomly generate the initial configuration, making sure that all the colors are within gamut. During optimization, colors are also checked to lie within the gamut. If a color were to move outside the gamut, it will be projected back to the closest point on the gamut boundary. In the illustration of Figure 6, the gamut boundaries appear as the boundaries between colored regions and background. Finally, to avoid finding only local minima, we run the optimization process several times with random initializations and keep the best solution. Figure 7 shows the result of optimization for two different example sets of input parameters. As expected, optimal colors tend to come from the green and yellow regions, which lead to low energy consumption.

6. Applications and Results

Sets of discrete colors may be used in various visualization applications that aim at showing nominal data or membership and grouping information. In this section, we present two typical examples. The first example is a 2D mapping that may appear in cartography, geospatial information systems, or other 2D data visualization. The second example is from direct volume visualization, where discrete colors are employed to identify differently classified materials. We intend to demonstrate that energy aware colors can be chosen for such typical visualization applications without sacrificing perceptual distinguishability. Please note that the result images of this paper are designed for monitor display, not for good reproduction in print. Therefore, the images are best viewed on screen.

Figure 8 shows an example of a 2D map with discrete colors. The map is taken from ColorBrewer [HB03]. We have chosen ColorBrewer as a ground for comparison because ColorBrewer is specifically designed for creating color mappings for 2D maps, facilitating good visual perception. Figure 8(a) shows the 4-class qualitative Dark2 coloring scheme from ColorBrewer. In our first approach (Section 4), the average CIELAB lightness of the original 4 ColorBrewer colors, $L^* = 54$, is used to look up 4 energy aware colors from our discrete optimization result. Figure 4 suggests picking yellow, green, orange, and purple in that order. We apply our colors to the map to obtain Figure 8(b). In our second approach (Section 5), we optimize for 4 colors in CIELAB space using the input parameters N = 4, $L^* = 54$, and d = 67. The minimum distance d = 67 is chosen identical to the smallest CIELAB color distance between the original 4 ColorBrewer colors. The resulting colors are applied to the map to obtain Figure 8(c). From visual inspection, the three images of Figure 8 provide visual encoding and distinguishability on similar levels. Furthermore, Figure 8(c) guarantees the same CIELAB distinguishability as the original color map from Figure 8(a); i.e., based on CIELAB, these two color maps provide the same quality of perceptual distinguishability.

Figure 9 documents the energy consumption of the three images of Figure 8. Energy consumption is measured according to our energy model from Equation 1, using various grid sizes. The most prominent observation is the substantial energy savings achieved by colors chosen according to our continuous optimization approach. Across the different grid resolutions, we save on average 41 percentage of energy compared to using ColorBrewer colors. In the discrete optimization case, we actually see an increase in energy when the grid resolution is low (less than 288 tiles). This is due to our purple color having a large RGB value. In general, the energy consumption drops with increasing number of tiles because the energy control is becoming more fine-grained. A typical number of tiles for practical applications is in the range of 500 to 1000; for comparison, 760 LEDs (i.e., "tiles") are reported for the HDR display developed by Seetzen et al. [SHS*04]. On OLED displays with one or more LEDs per pixel, a higher number of tiles can be used.

The second visualization application is from volume vi-







Figure 8: (a) Colors from the 4-class qualitative Dark2 scheme in ColorBrewer. (b) Four energy aware colors from our discrete optimization. (c) Four energy aware colors from our continuous optimization in perceptual CIELAB space.

sualization. Figure 10 shows the example of the volume rendering of a tooth dataset. The images were generated by GPU-based 3D texture slicing with post-classification via 1D transfer functions [EHK*06]. The tooth dataset has 3 distinct regions: the crown, outer encasing, and interior. The 3 regions are classified by assigning 3 different colors to the

© 2008 The Author(s) Journal compilation © 2008 The Eurographics Association and Blackwell Publishing Ltd.



Figure 9: Energy comparison for colors applied to the 2D ColorBrewer map.

respective scalar value ranges within the transfer function. We want to choose a fair, traditional set of 3 base colors to compare our energy aware colors against. This proved to be difficult, as there is no rule-of-thumb for picking traditional base colors. We decide to use shades of red, green, and blue each at 75% RGB intensity for comparison because color pickers in typical volume-visualization software are based on RGB and users tend to apply pure RGB colors. The average lightness of those RGB colors is $L^* = 45$. In the first approach using discrete optimization, we choose yellow, green, and orange at $L^* = 45$. In the second approach, we use the parameters N = 3 and $L^* = 45$. We estimate an appropriate d by dividing the length of the gamut boundary of the $L^* = 45$ iso-lightness slice (see Figure 5(a)) by 7, the maximum number of recommended iso-lightness colors [Hea96]. This gives us a minimum distance d = 70 for the continuous optimization. The tooth dataset was automatically rotated to simulate typical volume interaction, and 11 energy measurements were taken at regular intervals, then averaged. Figure 10 shows the tooth colored using the traditional colors and energy aware colors. From visual inspection, colors appear to be similarly discernible in Figure 10(a) (i.e., the original RGB choice) and Figure 10(c) (i.e., continuous optimization). However, the discrete optimization in Figure 10(b) makes it hard to distinguish the outer encasing and interior. As already discussed in Section 4, the colors from discrete optimization are not designed to optimize perceptual distance. The lack of perceptual distinguishability is apparent in the volume-rendering example because only 3 different colors are used here. In contrast, the continuous optimization approach allows us to achieve a much larger perceptual difference between the 3 colors. Figure 11 documents the energy consumption for the volume-rendered images, with an average saving of 37 percent in energy for both of our approaches.



Figure 10: The tooth dataset colored with (a) traditional colors, (b) energy aware colors from our discrete optimization result, and (c) energy aware colors found through continuous optimization in CIELAB space. The tooth images (upper row) show the actual volume renderings. The colored squares (bottom row) represent the discrete colors used for volume classification.



Figure 11: Energy comparison for colors applied to the volume visualization of tooth data.

7. Conclusions

The main goal of this paper is to trigger awareness for the energy consumption associated with emerging energy adaptive color displays. We have shown that substantial energy savings can be achieved by changing the color mapping in data visualization—without sacrificing good visual perception. On a technical level, we have presented two design methods for distinguishable, iso-lightness colors with the goal of lowering the energy consumption of the display device. Our methods are based on a screen space variant energy model, using a low resolution grid of tiles to estimate energy consumption. The first approach is an optimization of discrete colors based on color category. We restrict ourselves to using a small set of color hues, as previous work suggests that having too many colors slows down the speed of identifying colors. The second approach optimizes over a continuous iso-lightness color space by minimizing a cost function that takes into account the energy of the colors and the distances between them. Both methods address color distinguishability and iso-lightness. We have compared our color mappings to those chosen by the ColorBrewer system for 2D maps and to a traditional set of colors for volume rendering. Typical energy savings are around 40 percent for these examples.

There are several challenges in designing energy aware colors. The discrete optimization approach might suffer from the problem that colors are not chosen according to perceptual distances. In particular, modifying lightness may lead to similar looking colors. For example, dark shades of orange and yellow look very similar because brown is considered a dark shade of yellow or orange alike. Therefore, additional control over perceptual distances would be beneficial for the discrete optimization approach. In contrast, the continuous optimization of colors provides explicit control over perceptual color distances and, thus, does not run into the above issue of the discrete approach. However, the continuous optimization is not aware of color names or categories. Therefore, we can end up having more than one color of the same category (see Figure 7(a) for an example with 3 greenish colors), which might be problematic if visualization is used for visual communication that requires naming of image regions. A solution might be a hybrid approach that combines the discrete and continuous optimization methods of this paper.

There are further challenges that could also be addressed in future work. Our color schemes work well on 2D nominal data because the colored regions do not overlap. However, color blending is a typical phenomenon in volume rendering, and can cause colors to merge and change dramatically. Additional constraints or optimization criteria might be applied to improve volume-rendered images, so that the original colors are more distinguishable. Another line of future research might explore alternative color (appearance) models. Perceptual effects like simultaneous contrast and surround luminance certainly have not been addressed in this paper. Similarly, visual inspection of our results ought to be followed by a full validation study, including actual power measurements and color naming tests, in order to assess the effectiveness of our color designs for specific visualization applications.

Our methods can also be used in the design of user interfaces, allowing developers to create low energy color schemes for their applications and still maintain content readability. Low energy color user interfaces can be especially beneficial for battery powered mobile devices. Furthermore, we are not only limited to iso-lightness colors; we only choose to use iso-lightness colors in this paper as a design constraint for our color maps. One can imagine the use of non-iso-lightness colors for the design of user interfaces. An optimization can be performed in a CIELAB subspace, bounded by user-specified L_{\min}^* and L_{\max}^* , to explore colors with varying lightness. Finally, the long-term goal is a general method to automatically convert arbitrary input images to images with energy aware colors.

References

- [BT07] BORLAND D., TAYLOR R.: Rainbow color map (still) considered harmful. *IEEE Computer Graphics and Applications* 27, 2 (2007), 14–17.
- [CC06] CHENG W., CHAO C.: Perception-guided power minimization for color sequential displays. In ACM Great Lakes Symposium on VLSI (GLSVLSI) (2006), pp. 290–295.
- [CSC02] CHOI I., SHIM H., CHANG N.: Low-power color TFT LCD display for hand-held embedded systems. In *International Symposium on Low Power Electronics and Design (ISLPED)* (2002), pp. 112–117.
- [Dol08] DOLBY: Technology behind Dolby HDR video, 2008. http://www.dolby.com/professional/video/ hdr-video-tech-overview.html, [last access: 29 Sep 2008].
- [EHK*06] ENGEL K., HADWIGER M., KNISS J. M., REZK-SALAMA C., WEISKOPF D.: *Real-Time Volume Graphics*. A K Peters, 2006.
- [Fai06] FAIRCHILD M. D.: Color Appearance Models, 2nd ed. Wiley & Sons, 2006.
- [For03] FORREST S. R.: The road to high efficiency organic light emitting devices. Organic Electronics 4, 2–3 (2003), 45–48.
- [HB03] HARROWER M. A., BREWER C. A.: ColorBrewer.org: an online tool for selecting color schemes for maps. *The Carto*graphic Journal 40, 1 (2003), 27–37.
- [Hea96] HEALEY C. G.: Choosing effective colours for data visualization. In *Proceedings of the IEEE Conference on Visualization* (1996), pp. 263–270.

© 2008 The Author(s)

Journal compilation © 2008 The Eurographics Association and Blackwell Publishing Ltd.

- [HVG*04] HARTER T., VROEGINDEWEIJ S., GEELHOED E., MANAHAN M., RANGANATHAN P.: Energy-aware user interfaces: an evaluation of user acceptance. In ACM SIGCHI Conference on Human Factors in Computing Systems (CHI) (2004), pp. 199–206.
- [ILMR03] IYER S., LUO L., MAYO R., RANGANATHAN P.: Energy-adaptive display system designs for future mobile environments. In *MobiSys '03: Proceedings of the 1st International Conference on Mobile Systems, Applications and Services* (2003), pp. 245–258.
- [IP05] IRANLI A., PEDRAM M.: DTM: dynamic tone mapping for backlight scaling. In DAC '05: Proceedings of the 42nd Annual Conference on Design Automation (2005), pp. 612–617.
- [KRC02] KINDLMANN G., REINHARD E., CREEM S.: Facebased luminance matching for perceptual colormap generation. In *Proceedings of the IEEE Conference on Visualization* (2002), pp. 299–306.
- [KSJ95] KANDEL E. R., SCHWARTZ J. H., JESSELL T. M. (Eds.): Essentials of Neural Science and Behavior. Appleton & Lange, Norwalk, 1995.
- [KUU95] KAWAI M., UCHIKAWA K., UJIKE H.: Influence of color category on visual search. In Annual Meeting of the Association for Research in Vision and Ophthalmology (Fort Lauderdale) (1995). Paper #2991.
- [MLC06] MOCHOCKI B., LAHIRI K., CADAMBI S.: Power analysis of mobile 3D graphics. In DATE '06: Proceedings of the Conference on Design, Automation and Test in Europe (2006), European Design and Automation Association, pp. 502–507.
- [MM05] MOSHNYAGA V. G., MORIKAWA E.: LCD display energy reduction by user monitoring. 94–97.
- [Mur08] MURUGESAN S.: Harnessing green IT: Principles and practices. *IT Professional 10*, 1 (2008), 24–33.
- [MW08] MOSER M., WEISKOPF D.: Interactive volume rendering on mobile devices. In Vision, Modeling, and Visualization VMV '08 Conference Proceedings (2008), pp. 217–226.
- [NM65] NELDER J. A., MEAD R.: A simplex method for function minimization. *The Computer Journal* 7, 4 (1965), 308–313.
- [RGMN06] RANGANATHAN P., GEELHOED E., MANAHAN M., NICHOLAS K.: Energy-aware user interfaces and energyadaptive displays. *Computer 39*, 3 (2006), 31–38.
- [RT90] RHEINGANS P., TEBBS B.: A tool for dynamic explorations of color mappings. In *Symposium on Interactive 3D Graphics* (1990), pp. 145–146.
- [She08] SHEARER F.: *Power Management in Mobile Devices*. Newnes, Burlington, MA, 2008.
- [SHS*04] SEETZEN H., HEIDRICH W., STUERZLINGER W., WARD G., WHITEHEAD L., TRENTACOSTE M., GHOSH A., VOROZCOVS A.: High dynamic range display systems. ACM Transactions on Graphics 23, 3 (2004), 760–768.
- [TLC07] TSAI P., LIANG C., CHEN H. H.: Image quality enhancement for low backlight TFT-LCD displays. In *IEEE International Conference on Image Processing (ICIP)* (2007), vol. 3, pp. 473–476.
- [VZJ06] VALLERIO K. S., ZHONG L., JHA N. K.: Energyefficient graphical user interface design. *IEEE Transactions on Mobile Computing* 5, 7 (2006), 846–859.
- [War88] WARE C.: Color sequences for univariate maps. IEEE Computer Graphics and Applications 8, 5 (1988), 41–49.
- [WC08] WILLIAMS J., CURTIS L.: Green: The new computing coat of arms? *IT Professional 10*, 1 (2008), 12–16.