Perceptually Driven Visibility Optimization for Categorical Data Visualization

Sungkil Lee, Mike Sips, and Hans-Peter Seidel

Abstract—Visualization techniques often use color to present categorical differences to a user. When selecting a color palette, the perceptual qualities of color need careful consideration. Large coherent groups visually suppress smaller groups and are often visually dominant in images. This paper introduces the concept of class visibility used to quantitatively measure the utility of a color palette to present coherent categorical structure to the user. We present a color optimization algorithm based on our class visibility metric to make categorical differences clearly visible to the user. We performed two user experiments on user preference and visual search to validate our visibility measure over a range of color palettes. The results indicate that visibility is a robust measure, and our color optimization can increase the effectiveness of categorical data visualizations.

Index Terms—Color design, visualization, visibility, user interface

1 Introduction

Graphical representation of data is an effective technique in assisting a user to recognize patterns, to detect trends, and to identify relationships hidden in data sets. Visualizing data is an algorithmic process of mapping data onto two planar variables (x and y positions on the screen space) and a limited number of retinal variables, such as color, size, value, texture, shape, and orientation [2]. Color is an effective retinal variable for distinguishing groups of items; a typical visualization task would be to locate a particular categorical group in the visual display.

Selection of categorical colors is a challenging problem, and often require insight into aesthetic and perceptual aspects of color. Unfortunately, many visualization packages provide little guidance on how to choose proper categorical colors for a particular visualization technique. More importantly, without an in-depth understanding of how color works in categorical data visualization, it is difficult for the packages to suggest a color palette that avoids difficulties in detecting coherent categorical groups.

Color appearance changes, depending on the spatial variations of data items in an image; this is because the visual sensitivity of the human visual system to spatial frequencies differs for various hue/saturation and luminance components of color. Choosing categorical color without considering these perceptual mechanisms can lead to misleading interpretation of data. Two problems generally occur: 1) visually dominant structures locally suppress small inhomogeneous groups; 2) colors of different groups are barely discriminable without maintaining sufficient perceptual distances among them.

These problems motivated us to explore the possibility of an automated, yet perceptually significant approach that optimizes the composition of a given categorical color mapping. Fig. 1a shows a visualization using a ColorBrewer palette and Fig. 1b its improvement with our optimization algorithm. In what follows, we present the challenges we faced and our efforts to solve them.

2 Our Contributions

First, we define a perceptually driven metric, called *class visibility*, to measure the perceptual intensity of a group. To address the first challenge, we propose class visibility that takes the spatial variation of groups into account. We associate the perceptual intensity of a group with its *visual saliency*. While visual saliency is typically used in detecting regions that pop out from an image, here we utilize visual saliency to find high spatial-frequency groups visually suppressed by their surrounding neighborhood. The aim of this research is to make these groups visible via optimization. While the estimation of visual saliency commonly relies on spatial averages [15], averaged colors do not correctly label a categorical group; this is particularly true for "center" image regions (mapped to the narrow retinal area of the highest visual acuity). To tackle this limitation, we propose using *point saliency* to measure the perceptual intensity of a *single* class in pixels of the visual display. *Class visibility* is then computed by integrating the saliency of points over the center areas. Further, our class visibility takes the structural coherence (fine-grained structure versus large homogeneous structure) of categorical groups into account, to differentiate between levels of distraction.

Second, we present a *visibility*-based optimization algorithm, to improve the color of categorical data visualizations. Our optimization aims to balance the contrast (of both

- *S. Lee is with the Department of Computer Science and Engineering, Sungkyunkwan University, Suwon 440-746, Republic of Korea. E-mail: sungkil@skku.edu.*
- *M. Sips is with the Helmholtz Centre Potsdam, German Research Centre for GeoScience GFZ, Sektion 1.5, Geoinformatik Telegrafenberg, A20 303, D-14473 Potsdam, Germany. E-mail: sips@gfz-potsdam.de.*
- *H.-P. Seidel is with the Max-Planck-Institut Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany. E-mail: hpseidel@mpi-inf.mpg.de.*

Manuscript received 30 Nov. 2011; revised 22 Aug. 2012; accepted 1 Nov. 2012; published online 29 Nov. 2012.
Recommended for acceptance by K. Mueller.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCG-2011-11-0301.
Digital Object Identifier no. 10.1109/TVCG.2012.315.

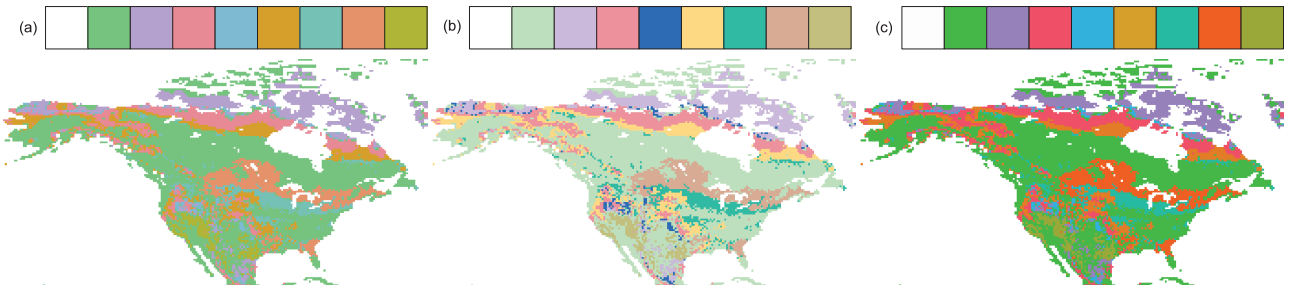

Fig. 1. While the categorical colors of the input (a) are easily distinguishable in the large color bar, their mapping onto the complex structure produces a limited visibility of categorical differences. Our visibility optimization (b) significantly improves the visibility of categorical differences in contrast to (a). Note that an enhancement of global contrast does not necessarily achieve such a result (c). The three visualizations present the distribution of eight vegetation classes, induced by climate change in North America.

luminance and saturation) of the image between small and large groups. The algorithm improves the composition of a color palette, using colors that increase the perceptual intensity of small groups and reduce the visual weight of large groups. This optimization approach allows us to improve a color palette chosen by existing tools or artists, yet not optimized for the given data. We derive a cost functional that virtually equalizes *class visibility* scores of distinct groups. The optimization procedure tackles the second problem we mentioned by applying *discrimination constraints*, which allow categorical colors (of low spatial frequencies) to keep perceptually distinguishable distances from each other. We explicitly control the dimensions of saturation and luminance, while simultaneously preserving the hue of categorical colors. This enables us to preclude the possibility that optimized categorical colors significantly differ from the user's preferred colors. Furthermore, local optima of our cost metric are likely to capture categorical color candidates that match the user's preference.

Lastly, we performed user experiments to subjectively assess our approach in terms of user preference and visual search performance. We measured the preference of participants for the source image as well as the two optimized images. In the visual search task, we measured task performance when searching for a particular class. The results show that our approach greatly improves the effectiveness of categorical colors, and most users prefer the optimized color palette.

3 Background and Related Work

The study of how to use color effectively in visualizations has commanded considerable research effort. This section reviews previous studies and neurological background underlying our optimization scheme.

3.1 Color Selection in Visualization

3.1.1 Principles of Effective Color Palettes

Trumbo [30] introduced *order* and *separation* as basic principles when selecting colors to visualize statistical variables. *Order* requires that colors chosen to present an ordered statistical variable should be perceived as preserving that order. *Separation* requires that colors chosen to present the differences of a variable should be clearly perceived as different. We focus on *separation* by considering the spatial coherence of categorical groups and the contribution of the local neighborhood.

Our optimization approach follows three additional principles provided by Zeileis et al. [37]; "colors should not be unappealing," "colors should cooperate with each other," and "colors should work everywhere."

3.1.2 Algorithms for Color Palettes

Many approaches have been proposed to realize *order* and *separation* for univariate color palettes [21], [23]. They were designed to select colors with uniform contrast by sampling along a smooth and continuous path in a perceptual color space such as CIE $L^*a^*b^*$ or CIE $L^*u^*v^*$ [7] (in short, Lab or Luv).

3.1.3 Predefined Color Palettes

ColorBrewer.org [13], [34] is a widely used tool for selecting customized color palettes. A drawback to such tools is that the spatial coherence and the spatial context of categorical groups are not considered in the selection process. Many users can easily select their favorite colors, but they need guidance to improve the local contrast of categorical groups; this motivated our study into the possibility of an automated enhancement of categorical color mapping for casual users.

3.1.4 Rule-Based Systems

Rule-based systems offer better choices by incorporating well-established rules used by professional designers and basic psychophysical properties of color [1], [32]. In particular, Bergman and collaborators incorporated the varying sensitivity of the human visual system in spatial variations, as a rule to restrict the choice of predefined color palettes for mapping continuous variables (see Section 3.2 for further discussion). They classify categorical groups to be of either low or high spatial frequency, by examining the difference between the input and its low-pass filtered image. However, this simple metric can be limited in the dichotomy between low and high spatial frequencies; we made the experience that many real data sets contain a mixture of coarse and finer spatial structures.

Our algorithm improves such a metric with a structured finer-scale analysis. Our visibility analysis allows us to *locally* distinguish the perceptual intensities of coarse and finer spatial structures, enabling more context-dependent categorical data visualization.

3.1.5 Perceptually Optimized Color Palettes

Study of the perceptual qualities of color in graphical data representations has captivated less interest amongst researchers in the visualization community. One of the basic findings of previous studies is that perceptually well-separated colors significantly reduce time required for searching visual targets in visual displays [8], [17], [35]—the results of our experiments on visual search empirically support this. Linear separation [8] or isoluminant colors in Luv were often used [14], [17]. These strategies are based solely on separating colors in Luv. On the contrary, our method considers the spatial coherence and spatial characteristics of categorical groups in a visualization.

3.2 Effect of Structure on Color Appearance

Visual sensitivity to spatial frequencies (or size) differs for both hue/saturation and luminance components of color (e.g., see [1], [6], [9], [27], [33]).

The hue mechanism of the human visual system is tuned to lower spatial frequencies, implying that hue and saturation are adequate for conveying large-scale spatial variation. As the target being viewed decreases in size, color appearance changes dramatically. For instance, as seen in Fig. 1a, the colors of the large uniform boxes are easily perceived as distinct, but small targets on the map become nearly indistinguishable, whereas the same colors are used to label them. This phenomenon is called "spreading" (e.g., see [9], [27]), and refers to the blending of a color stimuli with its surrounding colors. As a consequence, color discrimination is likely to be weak for small stimuli. Here, the contrast in saturation needs to be sufficiently increased to reliably label groups. Our optimization tries to increase the contrast for small regions and reduce it for larger regions via the saturation component; this allows us to produce images with balanced colors, according to the associated size of the categorical group.

On the other hand, the luminance mechanism is tuned to higher spatial frequencies. When high frequencies are not visually aggregated into clusters, they are usually invisible [18]. Hence, their luminance contrast needs to be increased to make labels distinguishable; discrimination via color becomes weaker in this case. Our optimization takes this into account. The enhancement of contrast in luminance and saturation channels improves the visibility of inhomogeneous small groups.

3.3 Feature-Driven Attention Model

The reflexive bottom-up capture of visual features has been proposed, to operate as a two-stage process in the "feature integration theory" [16], [25], [29]. In the first stage, preattentive primitives, e.g., color and lightness, are detected and separately encoded into feature maps. This stage is rapidly performed in parallel, without conscious attention directed to particular items. In contrast, the second stage is a slow serial conjunctive search to integrate the feature maps into a single topographical *saliency map*, which is largely mediated by focused attention [29].

The neuronal mechanisms of early vision provide us with important insights when detecting and enhancing the utility of categorical colors. Unlike large homogeneous structures, fine structures are visible only if their screen region stands out from its local surroundings. Substantial neurological evidence show that these findings result from the organization of the neural hardware, called the "center-surround [10], [20], [25], [28], [36]. The peripheral retinal zone (surround) suppresses neuronal activation in narrow receptive fields of the highest spatial acuity (center).

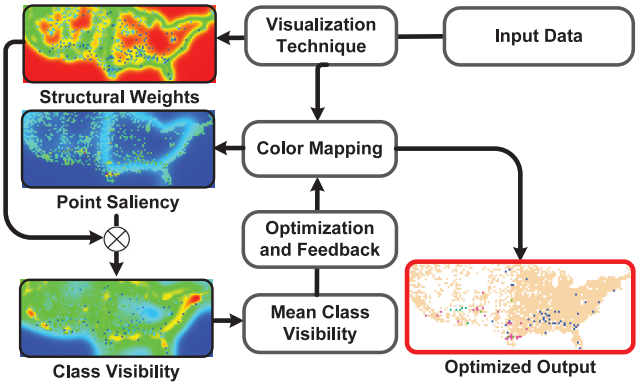

Fig. 2. Overview of our visibility optimization approach.

Typically, the computational model of the center-surround difference has been realized using the subtraction of multiscale image pyramids [15], [19]. In general, the center-surround difference is useful in the analysis of a complex scene, to detect salient spots that are likely to persist prior to the later recognition stage of visual objects [22].

While most previous studies used the center-surround difference to find salient spots that attract the user's attention in a complex scene, our approach utilizes saliency to quantify how many visual stimuli are visible against surrounding distractors, and to improve the effectiveness of categorical color encoding. This approach is likely to result in enhanced local contrast of invisible stimuli, which will also facilitate an improved discrimination among categorical colors.

4 Class Visibility

This section describes the concept of class visibility and an algorithm to compute class visibility levels of categorical data. Given labeled data points, a mapping of these data points to screen-space positions, and their associated labels to color, we compute the visual saliency of each point against its *surrounds* in the image (which we call *point saliency*). Then, we integrate the point saliency of each class over neighboring regions of different *center* scales, to estimate the visibility of individual classes without losing categorical differences. Fig. 2 shows the overview of our framework.

4.1 Preliminary Definitions

Let $D \subset \mathbf{Z}^2$ be the screen space. We define $D = \{1, \ldots, W\} \times \{1, \ldots, H\}$, where W and H are the dimensions of the screen window. Let A be the set of *data points* $A \subseteq D$ with $A = \{\mathbf{a}_1, \ldots, \mathbf{a}_N\}$ where $\mathbf{a}_i$ is the screen position of each point. We assume that each data point has a unique screen position with $i \neq j \Rightarrow \mathbf{a}_i \neq \mathbf{a}_j \ \forall i, j \in \{1, \ldots, N\}$; we do not

allow an overlay of different data points on the same pixel. Additionally, let $B = \{\mathbf{b} \in D | \mathbf{b} \notin A\}$ be the set of *background* (unoccupied) pixels in the screen space.

We assume that each point is uniquely labeled as belonging to a particular class. Let τ be an external source that labels data points by using a priori semantic information. Let $\tau(A) = \{c^1, \ldots, c^M\}$ be the set of M classes. The background pixels are assigned to a single class label $\tau(B) = \{c^0\}$. Thus, each point $\mathbf{p} \in D$ is assigned to one of the $M+1$ classes, where $c_\mathbf{p}$ is the associated class label. Let $C : \tau(D) \rightarrow \{C^0, C^1, \ldots, C^M\}$ be a color mapping function. To visualize the classes in the screen space, the set of unique class labels $\tau(D)$ is mapped to a set of $M+1$ unique colors. We denote $C_\mathbf{p}$ as the associated color of a data point $\mathbf{p} \in D$.

4.2 Point Saliency

Point saliency at a particular screen position represents how much a data point's color stands out from the local surround at that position. We formally define the point saliency as follows:

Definition 4.1 (Point Saliency). *Given a point $\mathbf{p} \in D$, the point saliency $S_\mathbf{p}$ of $\mathbf{p}$ is the single-scale difference between $C_\mathbf{p}$ and the mean color of the neighboring surround $\mathcal{N}_\mathbf{p}$ of $\mathbf{p}$ with*

$$S_\mathbf{p} = \Delta\mathcal{E}\left(C_\mathbf{p}, \frac{1}{|\mathcal{N}_\mathbf{p}|}\sum_{\mathbf{q}\in\mathcal{N}_\mathbf{p}} C_\mathbf{q}\right), \tag{1}$$

where $\Delta\mathcal{E}$ denotes a color distance metric, to measure the difference between two colors (see Section 4.3). We explain how to determine the scale $\mathcal{N}_\mathbf{p}$ in Section 4.4.

Our point saliency is similar to center-surround differences. The center is defined for a single point/pixel, which means that the point saliency at pixel $\mathbf{p}$ directly maps to one of the classes at $\mathbf{p}$. This definition allows us to weight the contribution of a single categorical color to the visibility of the categorical groups whose labels cannot be averaged spatially.

In contrast, conventional center-surround differences define center by averaging colors over narrow spatial regions [15]. Thus, the resulting color of each pixel does not represent a single class; the resulting color is likely to lose its categorical distinction. Hence, the conventional definition cannot be directly applied when estimating the visibility of individual classes.

4.3 Measurement of Color Distance

Point saliency requires us to measure the perceptual distance $\Delta\mathcal{E}$ between two colors. Many attempts have been made in the past to establish metrics to measure perceptual color distance. One fundamental option is the CIE 1976 metric [7], based on the perceptually uniform CIE *Lab* space. *Lab* space was developed as a color space that allows users to determine color differences. Given a white-point reference, an euclidean distance between two *Lab* colors represents a device-independent difference between the two *Lab* colors to the human eye.

As already alluded to, we keep initial hue values constant during the optimization (Section 5). To exclude the hue component from the optimization, it is more convenient to represent the distance in terms of CIE

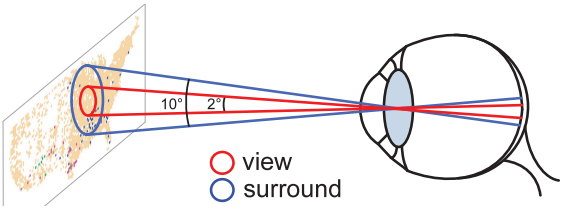

Fig. 3. The visual angles of the view and surround.

$L^*C^*H^*$ (in short, LCH) that is compatible with *Lab* space—C^* (chroma) and H^* (hue) are calculated by taking the radius and angle in the plane of two chromacity components a^* and b^*. We define the distance $\Delta\mathcal{E}(x, y)$ between two color values in LCH space $x = [L^*_x, r_x, \theta_x]^\top$ and $y = [L^*_y, r_y, \theta_y]^\top$ as

$$\Delta\mathcal{E}(x, y) = \sqrt{(\Delta L^*)^2 + (\Delta a^*)^2 + (\Delta b^*)^2}, \tag{2}$$

where

$$\begin{aligned} \Delta L^* &= L^*_x - L^*_y \\ \Delta a^* &= r_x \cos\theta_x - r_y \cos\theta_y \\ \Delta b^* &= r_x \sin\theta_x - r_y \sin\theta_y. \end{aligned} \tag{3}$$

4.4 Scales of Center and Surround

Class visibility is derived by integrating point saliency over a narrow region of the highest visual acuity in the retina (fovea region). We call such a region *view* to distinguish it from the traditional term "center." The view is formally defined as follows:

Definition 4.2 (View). *The view $\Omega_\mathbf{p}$ is a ball with radius ϵ centered at $\mathbf{p} \in D$*

$$\Omega_\mathbf{p} = \{\mathbf{q} \in A | \ \|\mathbf{p} - \mathbf{q}\| \leq \epsilon\}. \tag{4}$$

ϵ, which is fixed over the entire image, controls the extent of the highest spatial acuity in an image.

While the majority of previous studies [15] used an empirical size to indicate the center, we determine the size of the center in a manner dependent on the viewing angle. This strategy reduces manual adjustment to gain better performance. We use the visual angle of 2 degree to the display, which corresponds to the angle from the pupil to the rod-free fovea region [9]. The radius of the center ϵ can be easily calculated in terms of pixels, a given resolution and the size of a screen. Likewise, we specify the visual angle for the surround as used in (1). We use a visual angle of 10 degree, which approximates suppressive zones found in neurophysiological studies (8-10 degree) [26]. Fig. 3 illustrates the concept of the view and surround.

In practice, it is more convenient to handle diverse viewing conditions; for example, an observer is located slightly closer or further from the screen. We handle this with multiple levels of the center and surround. We use two scales for the center and three for the surround—in our implementation, 1 and 2 degree for the center and 5, 10, and 20 degree for the surround—resulting in six pairs of center-surround scales. Our design benefits from its grounding in the typical viewing scenarios of applications, because

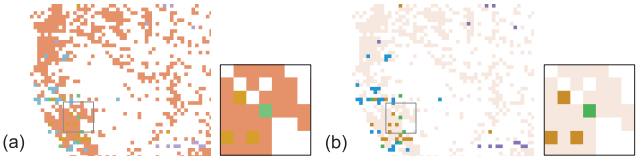

Fig. 4. Example of data points with weak visibility given the used colors ((a) see the green point) and its improvement in visibility (b).

typical multiscale representations do not pinpoint exactly which of the scales dominates the subsequent processing.

4.5 Class Visibility

Our *class visibility* is a measure of how color and the spatial distribution of each class affect its perceptual intensity to the human visual system. We estimate the *class visibility* of a point by weighting its point saliency with its structural weight. The structural weight at the point is computed by the contribution of neighboring points. Unlike the conventional filtering approach, we only collect the contribution of points whose labels are the same as the center point. By precluding the contribution of points of different labels, we can avoid losing categorical differences, while effectively approximating the spatial integration of the center areas. We formally define class visibility as follows:

Definition 4.3 (Class Visibility). *The visibility $V_{\mathbf{p}}$ of the class c^m at a pixel $\mathbf{p} \in D$ with $c_{\mathbf{p}} = c^m$ is*

$$V_{\mathbf{p}} = w_{\mathbf{p}} S_{\mathbf{p}}, \tag{5}$$

where $w_{\mathbf{p}}$ is a structural weight of $\mathbf{p}$ that measures the contribution of $c_{\mathbf{p}}$ to the visibility of c^m in $\Omega_{\mathbf{p}}$.

In our implementation, $w_{\mathbf{p}}$ is simply the fraction (or number) of data points inside $\Omega_{\mathbf{p}}$ that belong to the same group as the data point at $\mathbf{p}$. Our weighted metric effectively captures the difficulty of a user to observe a specific class at narrow local regions. For instance, imagine a region around a couple of green points in which the majority of the points are red (see Fig. 4a for an illustration). In this case, an observer might not easily notice the presence of the green points, because homogeneously distributed red points are spread around the green points, and distract the observer. Such difficulties can be reflected by small structural weights. One could improve this definition by using a better metric such as "visual clutter" [3] to consider the shape of the data distribution. Our attempts to employ visual clutter have not yielded great improvements, however.

Finally, we need a representative measure of each class through the whole image beyond a single point or view. Since class visibility is defined for each point, we average the class visibility for all the points in the image whose labels are identical. We call it *mean class visibility*, and its formal definition is as follows:

Definition 4.4 (Mean Class Visibility). *Let $c^m \in \tau(D)$ be a group of data points. The mean class visibility V^m of c^m is defined as*

$$V^m = \frac{1}{|c^m|} \sum_{\mathbf{p} \in c^m} V_{\mathbf{p}}, \tag{6}$$

where $|c^m|$ is the number of its data points.

Fig. 5 shows example images obtained with our visibility analysis for an input image (Fig. 5a) and an optimized output image (Fig. 5d); our optimization attempts to equalize the visibilities of different classes (detailed in Section 5). In Figs. 5b and 5e, the point saliency of small-size classes appears relatively high in comparison with the large classes. However, it is observed that the small classes are not actually clearly visible in Fig. 5a or at least similarly visible to the large classes in Fig. 5d. Structural weighting with Fig. 5g captures such trends; overall, the visibility of the small classes is estimated lower than the larger classes (Fig. 5c) or at least similar to the larger classes (Fig. 5f), respectively.

5 Color Optimization

This section describes our algorithm for perceptually optimizing a given color palette, based on our class visibility metric.

5.1 Cost Functional

The objective of our optimization is to *virtually equalize* the visibility of all classes. Our optimization makes almost

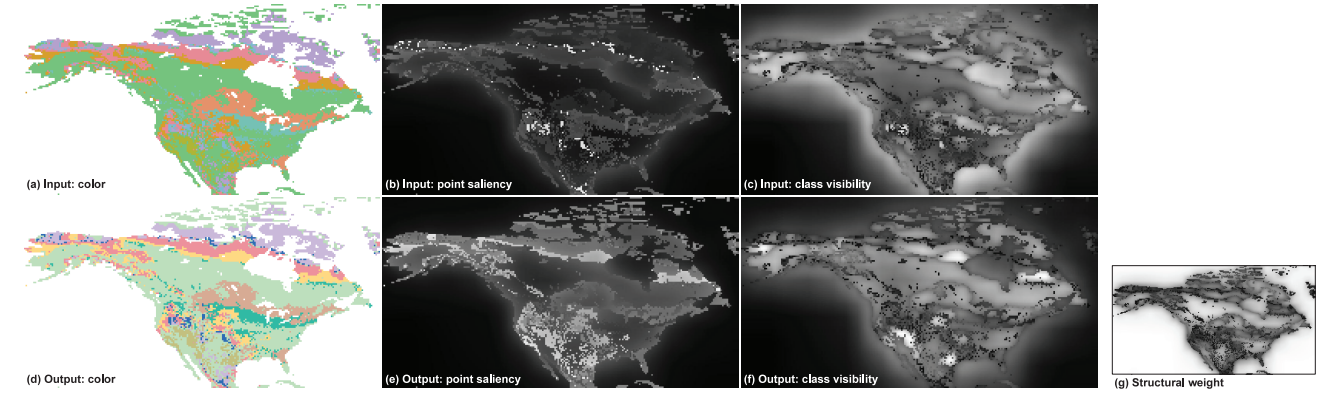

Fig. 5. The analysis and optimization of an example input image (a). The overall visibility of all the classes was significantly improved in the output image (d). The output was obtained with the color optimization procedure. The point saliency, class visibility, and structural weight images are shown aside for each input and output color image. Note that the intermediate images are normalized into $[0, 1]$ (black: 0, white: 1).

invisible classes (particularly those with a small number of data points) pop out from the surround, while keeping visually dominant classes visible.

One possible approach is to define a cost functional for each data point and to minimize differences among the class visibility levels in its local surrounds. Although such a local optimization might find optimal colors for each point, this is infeasible for categorical data visualization, because the color of a particular class should be consistent through the whole image.

To preserve the color of a particular class throughout the image, we define the variable to optimize in the global scope. Let $\mathbf{X}$ be the vector to be optimized. Given a set of categorical colors, $\mathbf{X}$ is defined as

$$\mathbf{X} = [L^0,\ \ldots,\ L^M,\ r^0,\ \ldots,\ r^M]^\top, \tag{7}$$

where L^k and r^k are the lightness and chroma of C_k in LCH space. Note that hue values (θ) will stay constant during the optimization. Also, note that $\mathbf{X}$ includes the color C_0 of background points, which often plays a key role in the detectability of main colors against a colored background (see Fig. 8 for an example).

Our cost functional for class visibility equalization is a mean-squared sum over the differences from all mean class visibilities to a single target constant (T; see Section 5.2 for details), which is defined as follow:

Definition 5.1 (Cost Functional). *Let $c^m \in \tau(D)$ be a group of data points, and let V^m be the mean visibility of a class c^m. The cost is*

$$E(\mathbf{X}) = \frac{1}{M}\sum_{m=0}^{M}(V^m(\mathbf{X}) - T)^2. \tag{8}$$

The minimization of this cost will virtually equalize class visibility levels, taking the spatial characteristics of the image into account, and thus, improving the visibility of structures less visible to the user.

5.2 Target Visibility

Our optimization equalizes the scores of class visibility to a single level T. One plausible way to choose such a target level is to use the mean over all V^m. Prior to optimization, we initially analyze the input image, and use the average value of the class visibility as the target visibility T. This target value is kept constant throughout the optimization process. Such a T improves the visibility of groups that are hard to detect, while visually dominant groups, such as large homogeneous classes, show a decrease. Such changes make the effects of categorical colors more balanced in their perceptual intensities.

Another alternative is to use the maximum among all V^ms. We may encounter hard-to-optimize visualizations when a mean visibility is used as a target visibility. For example, when dim or indistinguishable colors are used, all the classes in an image may exhibit a very low visibility. Then, the optimization with mean-target visibility causes all classes to stay at low visibility levels. Such difficulties can be surmounted by enforcing globally higher visibility levels (i.e., using the maximum or even higher).

In what follows, the target visibility that uses the mean visibility is denoted as T_{mean} and those with the maximum visibility are denoted as T_{max}. In our user study (Section 6), we compare the effects of T_{mean} and T_{max} in terms of preference and task performance.

5.3 Bounding Constraints

The domain of the cost functional (8) is unbounded, and accordingly, the optimized parameters can be located outside the set of physically meaningful color values. We avoid this problem by applying the barrier function, for each color channel y, given as

$$g_{[y_{\min},y_{\max}]}(y) = e^{-(y-y_{\min})} + e^{-(y_{\max}-y)}, \tag{9}$$

where $[y_{\min}, y_{\max}]$ is the acceptable range of y.

For instance, each channel of an RGB color should lie within $[0,1]$, in normalized scale. To facilitate this, we convert the resulting LCH color to RGB and evaluate the barrier functions at every iteration step of the optimization. We also clamp the range of luminance and chroma within $[0,1]$ in the normalized scale. When very dark or desaturated colors are not preferred, we can avoid them by more tightly bounding the luminance and chroma (e.g., $[0.1,1]$).

All those constraints are summarized in a single cost functional

$$B(\mathbf{X}) = \frac{1}{M}\sum_{m=0}^{M}(g(R^m) + g(G^m) + g(B^m) + g(L^m) + g(r^m)), \tag{10}$$

where R^m, G^m, B^m, L^m, and r^m are red, green, blue, luminance, and chroma values of C_m, respectively.

5.4 Color Discrimination Constraints

One additional important constraint is discrimination between colors that are perceptually close. The color of classes should be perceived as different by the user. To apply this constraint, we use a similar form of the barrier function, which defines a repulsive force term between all possible pairs of class colors, $R(\mathbf{X})$. Given $M(M+1)/2$ pairs of class colors, $R(\mathbf{X})$ can be formulated as

$$R(\mathbf{X}) = \frac{2}{M(M+1)}\sum_{i,j} e^{J-\Delta\mathcal{E}(i,j)}, \tag{11}$$

where i and j indicate the color pair selected from all the classes. J is a particular constant used to control the amount of the repulsive force. We scale J along with the just noticeable difference (JND) of Lab color.

In theory, when a pairwise color distance is greater than JND (known as 2.3 in Lab), this should produce distinguishable colors. However, the hue problem of the CIE76 metric (bluish colors have higher detection thresholds) does not guarantee a sufficiently noticeable difference. We tried to find a general value, yet there seemed to be no such a rule. Our rule of thumb is maintaining color distances at moderately higher JND (3-7 JND in our examples). Larger J is generally preferable for better distinguishing colors. However, when too many classes are defined, colors are not likely to move during the optimization, due to the strong repulsion force. In such cases, we may decrease J to enable the colors to move in the optimization.

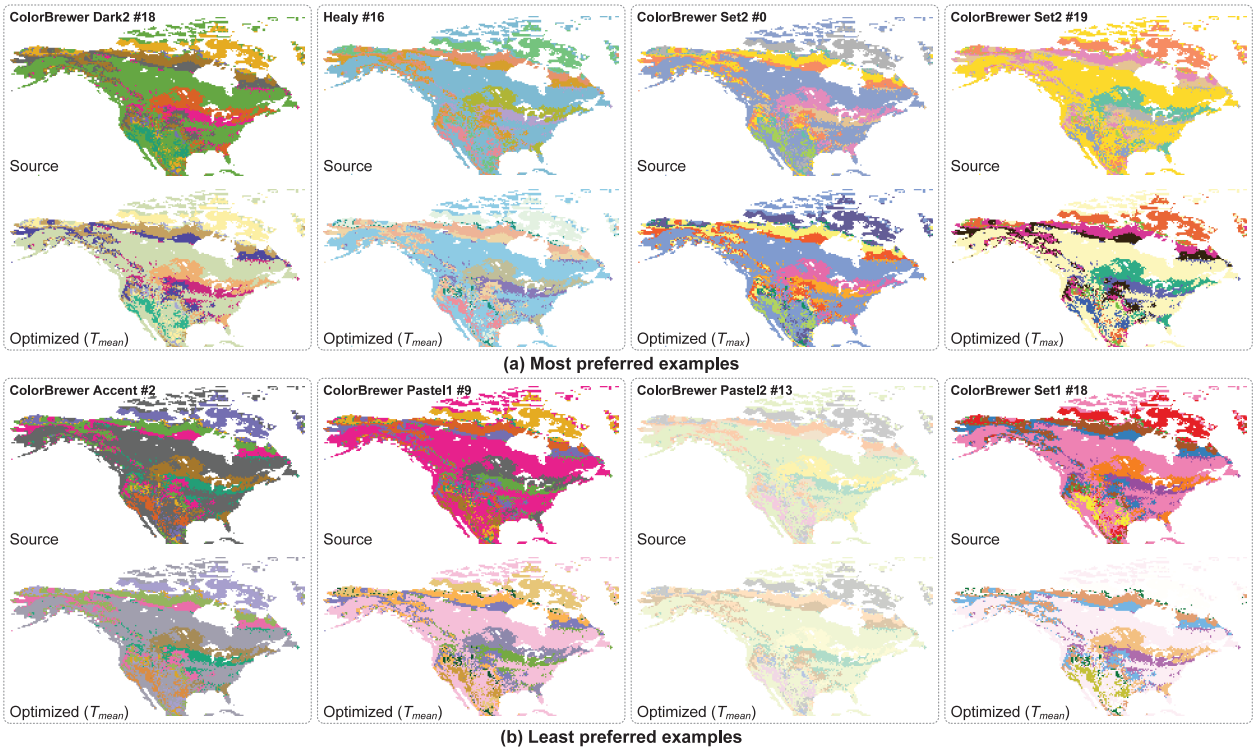

Fig. 6. Examples of most and least preferred output images used for our user experiment. "Vegetation class change" data were used ($M = 7$).

5.5 Optimization

An appropriate method of optimizing the cost functional (8) with the set of constraints (10) and (11) is to iteratively minimize

$$F(\mathbf{X}) = E(\mathbf{X}) + \lambda\{B(\mathbf{X}) + R(\mathbf{X})\} \tag{12}$$

and schedule the weight parameter λ to move toward infinity. As λ approaches infinity, our constraints will behave like hard constraints. This is a costly process, however. In practice, we found that a fixed λ value of 10 produces satisfactory results anyway. Since the resulting *soft constraints* do not completely rule out the possibility of obtaining infeasible color values, we project again each color value into the feasible region after optimization.

Since (12) is continuously differentiable, the nonlinear conjugate gradient-based minimization was adopted. The color of every single pixel can be represented as a linear combination of $\mathbf{X}$, and thus, we precompute the linear combination of weights prior to the iteration to save computational time in the evaluation step. The iteration for minimization stops when the norm of the gradients is small enough not to move any more.

An initial color might not reside in the valid extents given by our constraints, in particular for given luminance and chroma. To make sure that initial colors are projected into the valid range, we first run an optimization to minimize $B(\mathbf{X})$ prior to the main optimization of (12). Since $B(\mathbf{X})$ is convex, it can be easily minimized. We stop the iteration and evaluate T when all colors are within the valid extents.

Note that (12) may not be convex and, accordingly, finding a global optimum is difficult. However, we have observed that local optimal solutions are already visually plausible in many situations. When runtime performance is not critical, one could try a global optimization to find a better solution.

6 User Study

This section reports our user experiments, evaluating our visibility optimization in terms of user preference and performance of visual search tasks.

6.1 Methods

6.1.1 Participants and Stimuli

Twelve participants with normal or corrected-to-normal vision took part in the experiment. One-half of participants had professional training in visualization (visualization scientists and cartographers), while the other half had no experience in working with visualizations.

A 20-inch LCD display with a pixel resolution of $1{,}920 \times 1{,}200$ was used for the presentation of the stimuli. The display was calibrated using test images for faithful color reproduction [12]. The participants were seated 57 cm from the display in a lit room.

The experiment covered nine input color palettes chosen from ColorBrewer [34] and a perceptual color palette proposed by Healey and collaborators [14].

We used two types of geospatial map to present the stimuli, the ethnic distribution of the population of the United States ("us-ethnic") [31] and the vegetation class changes induced by climatic change in North America ("vegetation class change") [24]. "us-ethnic" presented five ethnic groups (see Fig. 7a for an example), and "vegetation class change" presented seven climate groups (see Figs. 6 and 7b). Each group was assigned 20 different permutations

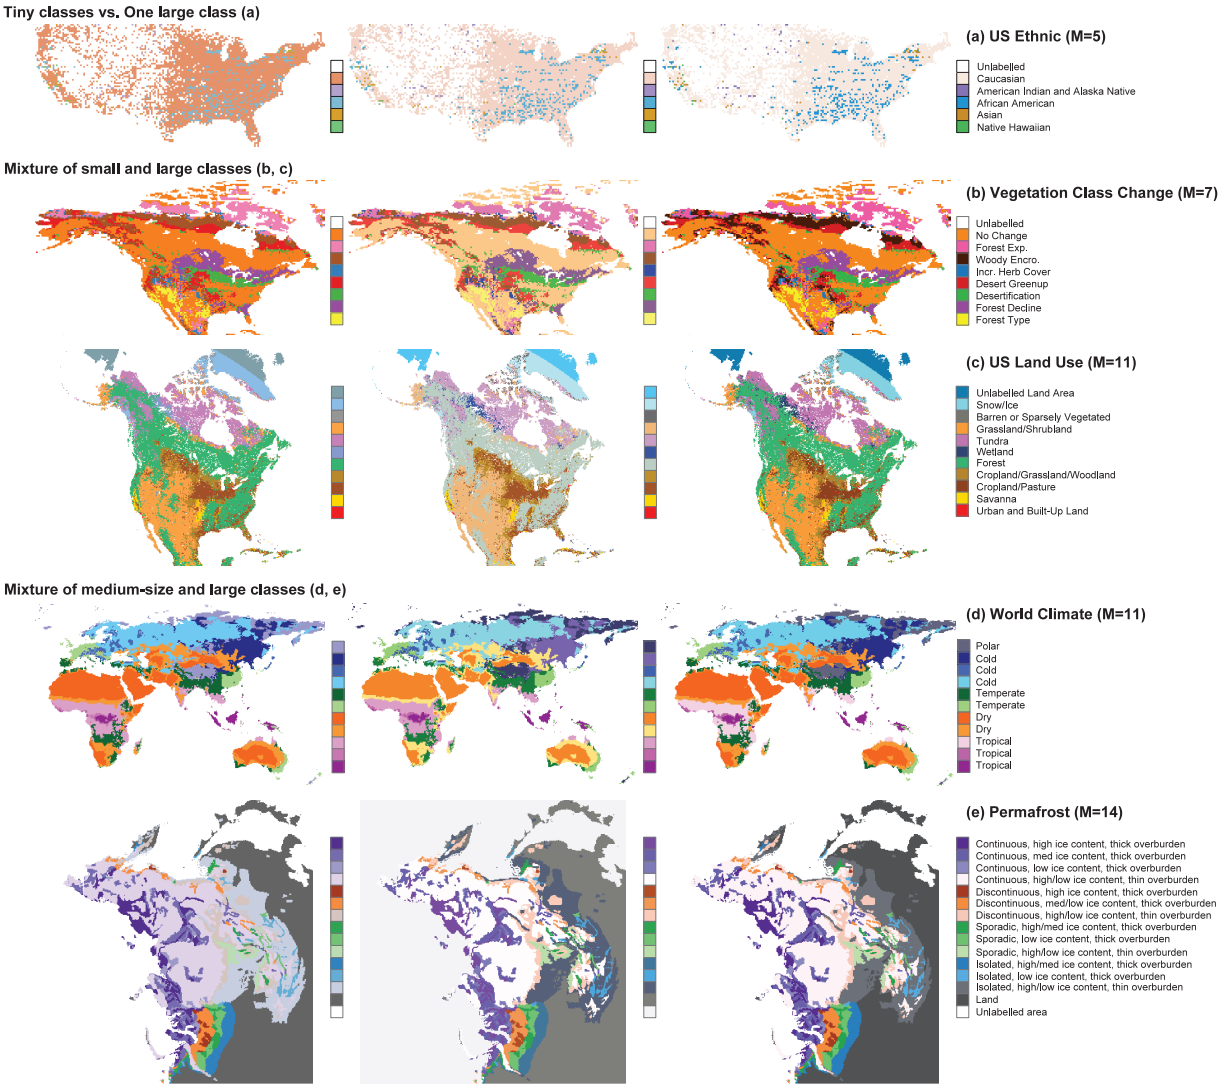

Fig. 7. Five source images and their optimized outputs (left: input, middle: T_{mean}, right: T_{max}). Our approach improves the visibility of small classes significantly, while the visibility of large classes is likely to decrease (the middle column) or to be maintained (the right column). *US Land Usage*, *World Climate*, and *Permafrost* images are provided through the courtesy of National Center for Earth Resources Observation and Science (US), FAO-SDRN Agrometeorology Group, and National Snow and Ice Data Center (US), respectively.

of a color palette; for each color palette, we generated 20 different images.

The input resolutions of the "us-ethnic" and "vegetation class change" maps for optimization were 720×304 and 464×248. To compare the effects of different target visibility, each input map produced two types of optimized outputs; the average and the maximum of input visibility (T_{mean} and T_{max}).

6.1.2 Design

The experiment consisted of two parts: a user-preference test and a visual search test. In the user preference test, participants were asked to rate their subjective preferences on a given input and its two optimized images. In the visual search, participants were asked to find a target color on a map. The visual search was chosen as a good task to test detectability and discrimination of different categorical colors.

In the first experiment, the user preference was rated in a 100-point scale (0: no preference, 100: most preferred). At the beginning of an experimental session, we briefly introduced each participant to the experimental procedure, and the number of classes for each data set. A session consisted of nine independent trials. We randomly selected a color palette for each trial from the database. We placed the three images—one with the original color palette and the other two with the optimized color palettes—side by side on the screen. We balanced the placement of image couples on the screen based on a randomization procedure to reduce learning effects. We applied no time constraint to this experiment.

In the visual search experiment, the participants were asked to find a target color in a color-mapped image. In each trial, we presented a target color and an image to the user. We balanced the sequence of images; one at the time, a single image, among an input or two optimized images, was randomly selected and presented. Then, we randomly chose a target color from the image's categorical color palette. The participants were instructed to find and click on the target color. Also, we asked the participants to press a

TABLE 1
Summary of Mean Preference Ratings

Map	source	output (T_{mean})	output (T_{max})
us-ethnic	30.4 (σ=16.5)	45.8 (σ=15.9)	50.1 (σ=12.7)
veg. class change	40.1 (σ=17.2)	52.2 (σ=14.3)	65.1 (σ=10.0)

σ denotes the standard deviation.

button labeled "cannot decide," when it was not possible to find the target color; in fact, no participant chose the "cannot decide" button at any time. After the participant clicked on the screen, we recorded the corresponding color and task completion time for each trial.

6.2 Experimental Results and Discussion

6.2.1 User Preference

We collected 324 ratings ($= 9$ palettes $\times 12$ participants $\times$ 3 types of images) for each input map. Table 1 summarizes the results of the first experiment. We can see a higher user preference for our optimized categorical color coding (both for T_{mean} and T_{max}). As for the target visibility, the outputs generated with T_{max} were more preferable to those with T_{mean}. We inspected these images and found that visibility optimization with T_{mean} often produces less distinguishable colors (see Fig. 6), while the visibility optimization using T_{max} tries to keep high contrasts among classes, even for large classes.

We applied one-way within-subject analysis of variance (ANOVA) on the preference ratings, to check the statistical significance of differences between source and optimized categorical palettes. The differences are statistically significant with $F(2, 22) = 15.03,\ p < 0.0001$ for "us-ethnic" and $F(2, 22) = 25.05,\ p < 0.0001$ for "vegetation class change," which confirms our observation. Tukey's multiple comparison test confirmed statistically significant differences for both input-output pairs ($p < 0.0001$). A significant difference was found between average and maximum visibility optimization only for the "vegetation class change."

6.2.2 Visual Search

We measured the task performance of the participants in the visual search experiment in terms of accuracy (i.e., correct hits on a given color) and task completion time. Table 2 summarizes our findings. Our optimization did not lead to significant increases in the mean hit ratios for the "us-ethnic" map, in comparison to the source maps, but yielded some improvements (up to 4 percent for outputs generated with T_{max}) with "vegetation class change" images. Regarding task completion time, while the outputs optimized with T_{mean} resulted in an increase of 2.9 second for "us-ethnic" map, both output types for the "vegetation class change" map led to a decrease in search time of 0.7 and 0.3 second. Again, we applied one-way within-subject ANOVA and found no statistical significance within the accuracy and task completion time.

The results indicate a positive effect from optimization on the accuracy and task completion time for "vegetation class change" maps representing a mixture of small and large classes; yet, we did not find a statistical significance. The outputs generated with T_{max} performed slightly better than T_{mean}, which implies that a globally higher contrast is better in task performance, and often preferred by users. Nonoptimized input palettes are often already well organized, and hence, there is little confusion when searching for a target in many general purpose color palettes. However, there are situations in which the user must make a considerable effort to detect a target color. The task completion timings of the optimized categorical color also support this observation.

Conversely, optimized outputs for the "us-ethnic" maps did not elucidate a trend in either the task performance measure. Tiny-sized classes in the "us-ethnic" map were much harder to observe in comparison with comparable-sized classes from the "vegetation class change" map. Hence, participants still had trouble in searching and distinguishing categorical colors. Improving the visualization of such tiny classes would be interesting future work; we expect that a higher target visibility for luminance discrimination would be necessary.

7 Results

This section reports the optimization performance of our system and presents quality improvements.

7.1 Optimization Performance

The time complexity involved in setting up and solving our optimization problem (12) is dominated by the number of classes M and the resolution $|D| = W \times H$ of the source. The cost for enforcing constraints is negligible in most cases, because the constraints are applied on the resulting colors

TABLE 2
Summary of Mean Accuracy and Task Completion Time in the Visual Search Experiment

Map	Accuracy			Task Completion Time		
	source	output (T_{mean})	output (T_{max})	source	output (T_{mean})	output (T_{max})
us-ethnic	78.3 % (σ=0.30)	81.7 % (σ=0.30)	68.3 % (σ=0.31)	5.3 s (σ=3.3)	7.2 s (σ=6.9)	5.0 s (σ=2.9)
veg. class change	93.2 % (σ=0.18)	95.8 % (σ=0.11)	97.1 % (σ=0.07)	3.2 s (σ=1.9)	2.5 s (σ=0.6)	2.9 s (σ=0.9)

σ denotes the standard deviation.

TABLE 3
Summary of the Optimization Performance

Data	$\|D\|$	Classes	Func Eval.	Timing
US Ethnic	360×152	5	83	0.19 s
	720×304	5	68	0.32 s
Veg. Class Change	232×124	7	45	0.10 s
	464×248	7	45	0.18 s
World Climate	349×245	11	66	0.33 s
	698×490	11	46	0.60 s
US Land Use	352×290	11	68	0.30 s
	704×579	11	41	0.46 s
Permafrost	310×345	14	69	0.37 s
	620×690	14	48	0.63 s

only. Given M and D, the time complexity of the optimization is $O(M|D|)$.

We measured the performance of our system (Table 3), implemented on an Intel Core i7 3.07-GHz machine with NVIDIA GTX 580 and Direct3D 10 API. The repeated evaluation of the cost function and gradients was significantly accelerated by the implementation on the GPU. These evaluations accounted for most of the computational cost in computing center-surround differences.

7.2 Quality Improvements

The findings of our user study serve as a basis for the following discussion on quality improvements. Fig. 6 compares the four top-rated and four worst rated examples of "vegetation class change" data used in the user preference test. In Fig. 6a, the visibility of small classes, which are hard to detect in the input, is significantly improved, and larger classes become weaker in their visibility. Meanwhile, different classes are still clearly distinguishable despite their sizes. Thereby, the visualization of small and large classes are well balanced in their visibility. In our user experiment, most participants preferred such optimized outputs in comparison with the original input.

However, we occasionally encountered less preferred outputs (see Fig. 6b). We inspected these images and it became apparent that they can be categorized into two types: those where strong contrast between large classes is lost in the optimized outcomes, and those where there is globally low visibility and discrimination in all classes (e.g., ColorBrewer Pastel2 #13). We observed good responses in both cases for T_{max} in contrast to T_{mean}; here, T_{max} is trying to keep the contrast of large classes high but also tries to increase the visibility of all classes at a global scale.

We further tested our algorithm to examine visibility improvements on more diverse sets of images. Fig. 7 compares five input images selected within the geographical visualization domain against the two optimized outputs for each. The left column presents the original images, the middle column their optimized images using T_{mean}, and the right column their optimized images using T_{max}. The first two examples were generated with *US Ethnic* and *Vegetation Class Change* data. The other three examples (*World Climate*, *US Land Use*, and *Permafrost*) were generated using the given predefined color palette—this implies that the initial color mappings are relatively well tuned for these data sets. Since the number of colors can range up to 15 in most practical visualizations, we merged similar classes, grouping them into less than 15 classes for (c), (d), and (e). The data sets are placed in the increasing order of classes M. We further categorized them into four types according to the spatial characteristics of tiny, small, medium-size, and large groups.

Overall, the input images had problems in presenting small classes or in discriminating different classes from each other. The composition of all categorical palettes was significantly improved using our optimization algorithm. The minimization of class visibility differences effectively improves the visibility of small classes, and the discrimination constraint ensures that one can distinguish the colors at a global level. Tiny/small/medium-size classes had higher contrasts in terms of luminance and saturation, while large classes had reduced saturation and visibility. This indicates our algorithm improves data sets of different sizes well. When utilizing T_{max}, large classes are close to their input colors, while small classes are improved to some extent. This causes small classes to appear less distinct than those with T_{mean}, but it is still very useful for inputs with weak color cues.

Although our algorithm includes the background color for optimization, the resulting background color did not change substantially when using a typical white background, except for Fig. 7e. Our discrimination constraints kept the background color far away from any other colors, to ensure a better discrimination of colors. However, when the background color is not around the extrema of grayscale colors (i.e., white or black), inclusion of the background color is still useful; in such cases, the background color can be also improved through optimization for better visibility, similar to the other nonbackground colors.

Fig. 8 shows an example. Here, the inclusion of the background color into the optimization (c) significantly improves the overall visibility of all classes (both small and large classes), in contrast to the result of the optimization without considering the background color (b). This result can be explained by "simultaneous contrast" [4], [5], where the darker background adds complementary colors to neighboring nonbackground colors.

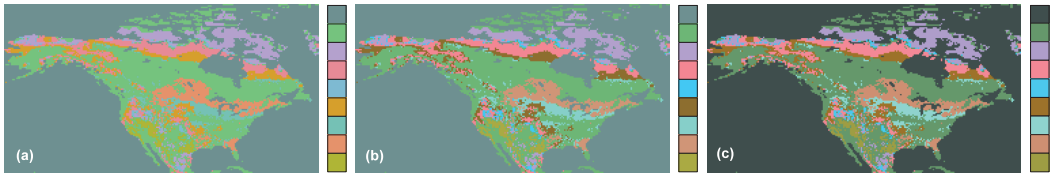

Fig. 8. The effect of optimization over background color: (a) input image, (b) output image without background color optimization, and (c) output image with background color optimization. T_{mean} was used for optimization.

8 Discussion and Limitations

The user study and the examples presented in the previous sections demonstrate that our optimized categorical color palettes yield good visibility in all classes and, thus, are preferred by users. One noteworthy aspect is that well-optimized visibility in all classes does not necessarily correlate with visually pleasing results. Particularly, maximum visibility in the optimization should be carefully used. Based on our experience, it should only be used for images that exhibit little contrast over all classes. Otherwise, we can often encounter exaggerated results that are visually displeasing.

We have experimented with the two different target visibilities T_{mean} and T_{max} for optimization, but did not find a clear conclusion on which strategy is better in general. Though, one observation is that optimization with T_{mean} often results in less distinguishable color for weak input color cues. One potential solution to this issue would be to utilize the characteristics of input images such as variance of categorical colors; according to the variance, we can use different target visibilities ranging between T_{mean} and T_{max}.

Our visibility analysis was inspired by the antagonistic organization of the center and surround in the human visual system. This was useful for detecting weak contrasts in luminance, yet still limited in detecting of hue/saturation components (sensitive to lower spatial frequencies). Consideration of the interactions of different color channels with regard to spatial frequency (or object size), such as the *spatial contrast sensitivity function* [9], would make our system more comprehensive for both chromatic and luminance channels.

We used fixed hue values throughout the optimization process. The resulting categorical color palettes match with the user's initial color preference or the suggestions made by visualization packages. However, this weakens the discrimination of large-scale color cues. The contrast enhancement of saturation components (via discrimination constraints) alleviates this problem, but a better solution is the use of a hue component, as noted in [1]. Considering hue components in the optimization process without losing the aesthetic aspect of a color palettes could be the subject of future research.

Another limitation of our system is grounded in the well-known limitation of the CIE Lab space. The distance in CIE Lab space might not be perceptually uniform for relatively long distances [11], and bluish colors have a higher threshold to get distinguished from similar neighboring colors. A simple remedy would be the minimization of the use of bluish colors or the use of alternative color spaces such as CIE1994 or CIEDE2000. However, in particular for CIEDE2000, the computation of gradients to be used in the optimization might be difficult, because it only guarantees $C0$-continuity. A nongradient-based technique might be used to facilitate optimization.

The use of CIE Lab space does not guarantee the optimal performance in all display configurations. In particular, it is well known that high-dynamic-range (HDR) display does not work well with CIE Lab space. To the best of our knowledge, studies on HDR perceptual color space are still ongoing.

Our algorithm is incompatible with a grayscale image as an input, leading to unpredictable results. This is because grayscale images have zero chroma values, which means hue values cannot be defined.

9 Conclusions and Future Work

We presented an automated color optimization algorithm based on the principles of human visual perception to make categorical differences clearly visible in a graphical representation of categorical data. We introduced the concept of visibility to quantitatively measure the utility of a color palette to present coherent structure to the user. Our measure detects regions in which visibility is inhibited by a visually dominant structure. We performed a user experiment over a range of general-purpose color palettes and two real-world data sets. The results indicate that visibility is a robust measure, and our color optimization can help to increase the effectiveness of a visual representation.

Acknowledgments

The authors would like to thank the anonymous reviewers for their constructive comments, Kwang In Kim at Max-Planck-Institut Informatik for his insightful advises on optimization, and the scientists at Potsdam Institute for Climate Impact Research for providing the (potential) vegetation class changes induced by climatic change in North America data set. They also thank Nayden Nachev for his help to implement a prototype, and the all scientists of the Sections 1.5 and 1.3 of the German Research Centre for GeoSciences GFZ for their helpful comments and continuing support of this research project. This work was in part supported by the German Federal Ministry of Education and Research, grant 03IS2191A, and by the Basic Science, Mid-career, and Global Frontier (on *Human-centered Interaction for Coexistence*) R&D programs through the NRF grants funded by the Korea Government (MSIP) (Nos. 2011-0014015, 2012R1A2A2A01045719, 2012M3A6A3055695).

References

[1] L.D. Bergman, B.E. Rogowitz, and L.A. Treinish, "A Rule-Based Tool for Assisting Colormap Selection," *Proc. IEEE Conf. Visualization,* pp. 118-125, 1995.
[2] J. Bertin, *Graphische Semiologie.* Walter de Gruyter, 1974.
[3] M.J. Bravo and H. Farid, "A Scale Invariant Measure of Clutter," *J. Vision,* vol. 8, no. 1, pp. 1-9, 2008.
[4] C.A. Brewer, "Review of Colour Terms and Simultaneous Contrast Research for Cartography," *Cartographica: The Int'l J. Geographic Information and Geovisualization,* vol. 29, no. 3, pp. 20-30, 1992.
[5] C.A. Brewer, "Evaluation of a Model for Predicting Simultaneous Contrast on Color Maps," *Professional Geographer,* vol. 49, no. 3, pp. 280-294, 1997.
[6] C.A. Brewer, "Color Use Guidelines for Data Representation," *Proc. Section Statistical Graphics,* pp. 55-60, 1999.
[7] *Recommendations on Uniform Colour Spaces, Colour Difference Equations, Psychometric Colour Terms.* CIE publication, 1978.
[8] M. D'Zmura, "Color in Visual Search," *Vision Research,* vol. 31, no. 6, pp. 951-966, 1991.
[9] M.D. Fairchild, *Color Appearance Models,* vol. 3, Wiley, 2005.
[10] C.L. Folk, R.W. Remington, and J.C. Johnston, "Involuntary Covert Orienting Is Contingent on Attentional Control Settings," *J. Experimental Psychology Human Perception and Performance,* vol. 18, pp. 1030-1030, 1992.

[11] E.M. Granger, "Is CIE L* a* b* Good Enough for Desktop Publishing?" *Proc. SPIE,* vol. 2170, p. 144, 1994.
[12] H.-K. Nienhuys, "The Lagom LCD Monitor Test Pages, 2012," http://www.lagom.nl/lcd-test/, Aug. 2012.
[13] M. Harrower and C.A. Brewer, "Colorbrewer.org: An Online Tool for Selecting Colour Schemes for Maps," *Cartographic J.,* vol. 40, pp. 27-37, 2003.
[14] C.G. Healey, "Choosing Effective Colours for Data Visualization," *Proc. IEEE Conf. Visualization,* pp. 263-271, 1996.
[15] L. Itti, C. Koch, and E. Niebur, "A Model of Saliency-Based Visual Attention for Rapid Scene Analysis," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 20, no. 11, pp. 1254-1259, Nov. 1998.
[16] J. Jonides, "Further Toward a Model of the Mind's Eye's Movement," *Bull. Psychonomic Soc.,* vol. 21, no. 5, pp. 247-250, 1983.
[17] M. Kawai, K. Uchikawa, and H. Ujike, "Influence of Color Category on Visual Search," *Proc. Ann. Meeting Assoc. for Research in Vision and Ophthalmology,* p. 2991, 1995.
[18] D.A. Keim, "Designing Pixel-Oriented Visualization Techniques: Theory and Applications," *IEEE Trans. Visualization and Computer Graphics,* vol. 6, no. 1, pp. 59-78, Jan.-Mar. 2000.
[19] C. Koch and S. Ullman, "Shifts in Selective Visual Attention: Towards the Underlying Neural Circuitry," *Human Neurobiology,* vol. 4, no. 4, pp. 219-227, 1985.
[20] S.W. Kuffler, "Discharge Patterns and Functional Organization of Mammalian Retina," *J. Neurophysiology,* vol. 16, no. 1, pp. 37-68, 1953.
[21] H. Levkowitz and G.T. Herman, "Color Scales for Image Data," *IEEE Computer Graphics and Applications,* vol. 12, no. 1, pp. 72-80, Jan. 1992.
[22] D. Parkhurst, K. Law, and E. Niebur, "Modeling the Role of Salience in the Allocation of Overt Visual Attention," *Vision Research,* vol. 42, no. 1, pp. 107-123, 2002.
[23] P. Robertson and J. O'Callaghan, "The Generation of Color Sequences for Univariate and Bivariate Mapping," *IEEE Computer Graphics and Applications,* vol. 6, no. 2, pp. 24-32, Feb. 1986.
[24] S. Sitch, B. Smith, I.C. Prentice, A. Arneth, A. Bondeau, W. Cramer, J. Kaplan, S. Levis, W. Lucht, M. Sykes, K. Thonicke, and S. Venevsky, "Evaluation of Ecosystem Dynamics, Plant Geography and Terrestrial Carbon Cycling in the LPJ Dynamic Vegetation Model," *Global Change Biology,* vol. 9, pp. 161-185, 2003.
[25] D. Sagi and B. Julesz, ""Where" and "What" in Vision," *Science,* vol. 228, no. 4704, pp. 1217-1219, 1985.
[26] J.D. Schall, T.R. Sato, K.G. Thompson, A.A. Vaughn, and C.-H. Juan, "Effects of Search Efficiency on Surround Suppression during Visual Selection in Frontal Eye Field," *J. Neurophysiology,* vol. 91, pp. 2765-2769, 2004.
[27] M. Stone, "In Color Perception, Size Matters," *IEEE Computer Graphics and Applications,* vol. 32, no. 2, pp. 8-13, Mar./Apr. 2012.
[28] J. Theeuwes, "Perceptual Selectivity for Color and Form," *Perception and Psychophysics,* vol. 51, no. 6, pp. 599-606, 1992.
[29] A.M. Treisman and G. Gelade, "A Feature-Integration Theory of Attention," *Cognitive Psychology,* vol. 12, pp. 97-136, 1980.
[30] B.E. Trumbo, "A Theory for Coloring Bivariate Statistical Maps," *Am. Statistician,* vol. 35, no. 4, pp. 220-226, 1981.
[31] United States Dept. of Commerce, Census Bureau Website, http://www.census.gov/, Mar. 2011.
[32] L. Wang, J. Giesen, K.T. McDonnell, P. Zolliker, and K. Mueller, "Color Design for Illustrative Visualization," *IEEE Trans. Visualization and Computer Graphics,* vol. 14, no. 6, pp. 1739-1754, Nov./Dec. 2008.
[33] C. Ware, *Information Visualization—Perception for Design,* vol. 3, Morgan Kaufmann, 2004.
[34] M. Wijffelaars, R. Vliegen, J.J. van Wijk, and E.-J. van der Linden, "Generating Color Palettes Using Intuitive Parameters," *Computer Graphics Forum,* vol. 27, no. 3, pp. 743-750, 2008.
[35] J.M. Wolfe, "Visual Search," *Attention,* vol. 1, pp. 13-73, 1998.
[36] J. Xing and D.J. Heeger, "Center-Surround Interactions in Foveal and Peripheral Vision," *Vision Research,* vol. 40, no. 22, pp. 3065-3072, 2000.
[37] A. Zeileis, K. Hornik, and P. Murrell, "Escaping RGBland: Selecting Colors for Statistical Graphics," *Computational Statistics and Data Analysis,* vol. 53, no. 9, pp. 3259-3270, 2009.

Sungkil Lee received the BS and PhD degrees in materials science and engineering and computer science and engineering at POSTECH, Korea, in 2002 and 2009, respectively. He is an assistant professor of computer science and engineering at Sungkyunkwan University, Korea. He was a postdoctoral researcher at the Max-Planck-Institut Informatik (2009-2011). His research interests include real-time GPU rendering, perception-based rendering, information visualization, GPU algorithms, and human-computer interaction.

Mike Sips received the PhD degree in computer science from the University of Konstanz in 2005. He is a research scientist at the German Research Centre for GeoSciences. His research interests include the development of novel visual analytics methods for spatial and temporal data. He was a visiting assistance professor at Stanford University between 2006 and 2008. During his research visit at Stanford University, he was the leader of the Independent Research Group Visual Exploration of Space-Time Pattern in Multi-Dimensional and Heterogeneous Data Spaces.

Hans-Peter Seidel is the scientific director and chair of the computer graphics group at the Max-Planck-Institut Informatik and a professor of computer science at Saarland University. He has published and lectured widely. He has received grants from a wide range of organizations, including the German National Science Foundation (DFG), the German Federal Government (BMBF), the European Community (EU), NATO, and the German-Israel Foundation (GIF). In 2003, he received the "Leibniz Preis," the most prestigious German research award, from the German Research Foundation (DFG). He is the first computer graphics researcher to receive this award.