

Modeling human color categorization

E.L. van den Broek^{a,*}, Th.E. Schouten^b, P.M.F. Kisters^c

^a Center for Telematics and Information Technology, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

^b Institute for Computing and Information Science, Radboud University Nijmegen, P.O. Box 9010, 6500 GL Nijmegen, The Netherlands

^c GX Creative Online Development B.V., Wijchenseweg 111, 6538 SW Nijmegen, The Netherlands

Available online 22 September 2007

Abstract

A unique color space segmentation method is introduced. It is founded on features of human cognition, where 11 color categories are used in processing color. In two experiments, human subjects were asked to categorize color stimuli into these 11 color categories, which resulted in markers for a Color LookUp Table (CLUT). These CLUT markers are projected on two 2D projections of the HSI color space. By applying the newly developed Fast Exact Euclidean Distance (FEED) transform on the projections, a complete and efficient segmentation of color space is achieved. With that, a human-based color space segmentation is generated, which is invariant for intensity changes. Moreover, the efficiency of the procedure facilitates the generation of adaptable, application-centered, color quantization schemes. It is shown to work excellently for color analysis, texture analysis, and for Color-Based Image Retrieval purposes.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Eleven color categories; Human color categorization; Color space segmentation; Fast Exact Euclidean Distance (FEED) transform

1. Introduction

Digital imaging technology is more and more embedded in a broad domain. As a consequence, digital image collections are booming, which creates the need for efficient data-mining in such collections (Petrushin and Khan, 2007). An adequate model of human visual perception would facilitate data-mining (Petrushin and Khan, 2007; van den Broek et al., 2006). Our approach, hereby, is to utilize human cognitive and perceptual characteristics.

For a broad range of data-mining applications, various image processing techniques have been adopted; e.g., the cultural domain (Flickner et al., 1995; van den Broek et al., 2006). However, the current paper focuses on a generic image processing technique: a color quantization scheme based on human perception. This unique color space segmentation is both relevant and suitable for the development and study of Content-Based Image Retrieval

(CBIR) in its variety of contexts (Schettini et al., 2001; van den Broek et al., 2005; van den Broek et al., 2006).

In general, we argue that color should be analyzed from the perspective of human color categories. Both to relate to the way people think, speak, and remember color and to reduce the data from 16 million or more colors to the limited number of 11 color categories: black, white, red, green, yellow, blue, brown, purple, pink, orange, and gray. Research from diverse fields of science emphasize their importance for human color perception (Derefeldt et al., 2004; Hansen et al., 2007; Kay and Regier, 2006). The use of this knowledge can possibly provide a solution for problems concerning the accessibility and the availability of knowledge, where color analysis is applied in data-mining. In addition, such a human-centered approach can tackle the computational burden of traditional (real-time) color analysis (Flickner et al., 1995; Hansen et al., 2007; Jou et al., 2004).

The 11 color categories are applicable for a broad range of CBIR domains (van den Broek et al., 2005), where in specific domains, other sets of colors might be more appropriate. In this paper, we regard the 11 color categories as

* Corresponding author. Tel.: +31 53 489 3604; fax: +31 53 489 2849.

E-mail addresses: vandenbroek@acm.org (E.L. van den Broek), T.Schouten@cs.ru.nl (Th.E. Schouten), p.kisters@gx.nl (P.M.F. Kisters).

they are used in daily life. These color categories are constructed and handled by methods that are presented in this paper. However, in the same way, it is possible to incorporate another set of colors, which is user, task, or application specific.

This paper presents a line of research starting with psychophysical experiments in Section 2. This provided us with color markers for a Color LookUp Table (CLUT) in the RGB color space. The boundaries between the color categories in the RGB space are expected to be too complex to be determined, using the limited number of CLUT markers. Therefore, we describe in Section 3 how the RGB space is transformed into two 2D projections of the HSI color space in which the boundaries are less complex. In Section 4, we present the newly developed Fast and Exact Euclidean Distance (FEED) transformation, adapted such that it can handle multi class data; in our case: the 11 color categories. Section 5 describes how the CLUT markers are fed to FEED to find the boundaries and how this is used to segment the complete color space. In Section 6, the segmented color space is validated through a comparison with human color categorization. We finish the paper with a discussion on the work presented in Section 7.

2. Validation of the 11 color categories

Twenty-six subjects with normal or corrected-to-normal vision and no color deficiencies participated voluntarily in two tasks. The first task was to write down the first 10 colors that came to mind. This task was embedded in the research because the 11 color categories are still a topic of debate; e.g., see Hansen et al. (2007) and Kay and Regier (2006). It enabled us to verify its validity for our research. The second task consisted of both a color memory experiment and a color discrimination experiment.

2.1. Method

For the color memory experiment, the subjects were instructed to categorize the stimulus into one of the color categories, represented by buttons with their names (task 2a; see Fig. 1a). In the color discrimination experiment,

the subjects were asked to press one of the 11 focal color buttons (i.e., a typical color for a color category) that best resembled the stimulus (task 2b; see Fig. 1b). Both experiments consist of four blocks of repetitions of all stimuli (in a randomized order).

The experiments ran in an average office environment on a PC with an Intel Pentium II 450 MHz processor, 128 mb RAM, a Matrox Millennium G200 AGP card, and with a Logitech 3-button Mouseman (model: M-S43) as pointing-device. The experiments were conducted in a browser-environment with Internet Explorer 6.0 as browser and Windows 98SE as operating system, using 16-bit colors, where respectively 5, 6, and 5 bits are assigned to the red, green, and blue channel.

The stimuli were the full set of 216 web-safe colors (W3 Schools, 2007). These are defined as follows: The R, G, and B dimensions (coordinates) range from 0 to 255 and are treated equally. For each dimension, six values are chosen: 0 (0%), 51 (20%), 102 (40%), 153 (60%), 204 (80%), and 255 (100%). Each of these six values is combined with each of the six values of the two other dimensions. This results in $6^3 (= 216)$ triple of coordinates in the RGB-space. These RGB-values result for both Internet Explorer and Netscape under both the Windows and the Mac operating system, in the same (non-dithered) colors, under the prerequisite that the operating system uses at least 8-bit (256) colors.

2.2. Results

The main results confirm the existence of the 11 color categories, as is illustrated in Table 1. Noteworthy is that all 26 participants wrote down the colors red, green, blue, and yellow; all belong to the 11 focal colors or color categories (Derefeldt et al., 2004; Hansen et al., 2007; Kay and Regier, 2006). With 11 occurrences, pink was the least mentioned focal color. Followed by the most frequently mentioned non-focal color: violet, which was mentioned by six of the participants. The complete results of task 1 are presented in Table 1. In the classification of the web-safe colors by the subjects, three sets of color markers can be distinguished: (i) non-fuzzy color markers: web-safe colors




Fig. 1. Screenshot of the user interface of: (a) the color memory experiment (gray buttons, labeled with a color name) and (b) the color discrimination experiment (colored buttons without a label). The stimulus did have a size of 9.5×6.0 cm, the buttons measured 1.8×1.2 cm with 0.6 cm between them. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
Frequency and confidence-intervals (p at 99%) of color names mentioned

Color name	Frequency (min–max)
Red, green, blue, yellow	26 (84.4–100.0%)
Purple	24 (74.5–98.8%)
Orange	22 (65.7–94.3%)
Black, white, brown	20 (57.5–89.2%)
Gray	15 (38.9–74.4%)
Pink	11 (25.6–61.1%)
Violet	06 (10.8–42.5%)
Beige	04 (5.7–34.3%)
Ocher	03 (3.3–30.0%)
Turquoise, magenta, Indigo, cyan	02 (1.1–25.5%)
Silver, gold, bordeaux-red	01 (0.7–20.7%)

categorized by a magnitude of at least 10 subjects to one color category; (ii) fuzzy color markers: web-safe colors that are assigned to two different color categories by at least 10 subjects; (iii) colors of which the subjects did not agree to what color category they belong. The set of non-fuzzy color markers are used in further processing since the color category they belong to is undisputed. The fuzzy color markers will be used in a later stage to validate the final color segmentation, based on the non-fuzzy color markers. The third and remaining category is excluded as data for further research.

In general, color matching using a Color LookUp Table (CLUT), based on the color markers derived from the experimental results, could enhance the color matching process significantly and may yield more intuitive values for users (Derefeldt et al., 2004; Hansen et al., 2007). In addition, such a coarse color space quantization of 11 color categories reduces the computational complexity of color analysis drastically, compared to existing matching algorithms of image retrieval engines that use color quantization schemes (c.f. PicHunter Cox et al., 1996: HSV $4 \times 4 \times 4$ and QBIC (Flickner et al., 1995): RGB $16 \times 16 \times 16$). The coarse 11 color categories quantization makes it also relatively invariant with respect to intensity changes (Hansen et al., 2007). For more detailed discussions concerning color quantization, we refer to Schettini et al. (2001).

3. The segmentation framework

The experiments presented in Section 2 provided us with categorized color markers for a Color LookUp Table (CLUT). These color markers are considered scarce data for segmenting a color space into 11 categories. Therefore, a framework is needed to that provides means to maximize the efficiency. In this section, we explain the framework that is used for the segmentation process.

3.1. HSI: segmentation color space for scarce data

The color markers are RGB coordinates; however, the RGB color space is not perceptually intuitive. Hence, the

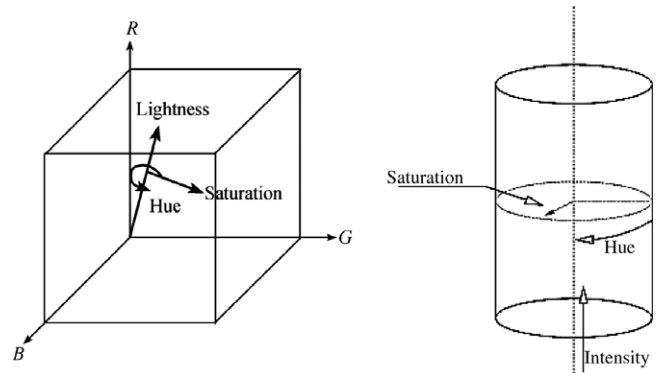


Fig. 2. Left: The relation between the RGB and the HSI color space, from the perspective of the RGB color space. Right: The cylinder shaped representation of the HSI (Hue, Saturation, and Intensity) color space, as used in this research.

position and shape of the color categories within the RGB color space are complex. Therefore, for the full color space categorization, the HSI color space is used, which is (i) perceptually intuitive, (ii) performs as good as or better than perceptual uniform color spaces such as CIE LUV (Lin and Zhang, 2000), and (iii) the shape and position of the color categories are less complex functions of location and orientation than with the RGB color space. See Fig. 2 for a visualization of the RGB and HSI color spaces as well as their relation.

Prosperously, the 216 web-safe colors are clearly distinct for human perception. As a consequence, in a perceptual intuitive color space some distance is present between them. Moreover, the perceptual intuitive character of the HSI color space results in an orientation of adjacent colors such that the web-safe colors are spatially arranged by color category.

Let us now briefly discuss the HSI color space, an intuitive color space that proved to work good under varying circumstances (Lin and Zhang, 2000; Schettini et al., 2001). The axes of the HSI space represent hue (i.e., basic color index), saturation (i.e., colorfulness or chromatic purity), and intensity (i.e., amount of white present in the color). The shape of HSI color space can be displayed as a cylinder: intensity is the central rod, hue is the angle around that rod, and saturation is the distance perpendicular to that rod; see also Figs. 2 and 3. The color categories' orientation is as follows: Around the intensity axis, the achromatic categories are located, as is shown in Fig. 3. The achromatic region has a conical shape and is described with small saturation values, the complete range of intensity, and the complete range of hue values. Around this achromatic region, the chromatic categories are located. Chromatic categories have high saturation values and occupy a part of both the total hue and the total intensity range.

3.2. From 3D HSI color space to two 2D representations

Since the HSI color space is a 3D space, the boundaries between color categories consist of 3D functions. However,

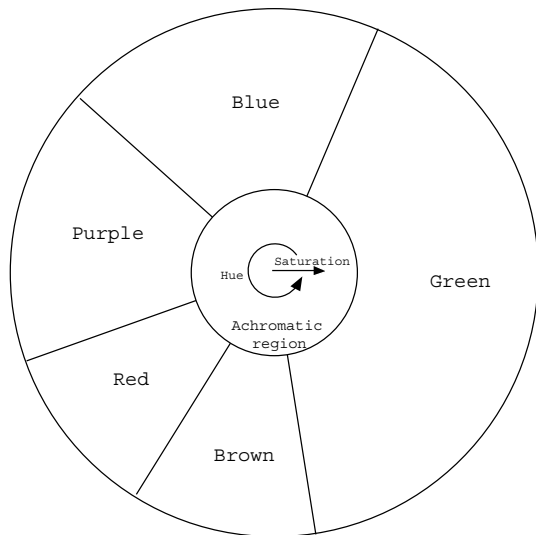


Fig. 3. Separation of the various regions in the hue-saturation plane with intensity value 300. Shown is both the separation between chromatic and achromatic regions and the separation between different hues (or colors), within the chromatic region.

the amount of HSI CLUT markers is too limited to determine the exact boundaries through a 3D segmentation, which would evolve in very weak estimations of the shape of color categories in color space. However, the perceptually intuitive axes of the HSI color space do allow a reduction in the complexity of boundary functions without losing essential features of the boundaries. The intuitive values that the axes represent facilitate the separation of chromatic and achromatic categories, using two 2D projections. Thereby, we use three assumptions:

- (1) The boundaries between achromatic categories and chromatic categories do not excessively change over the hue range; see also (Hansen et al., 2007).
- (2) The boundaries between chromatic categories do not excessively change over the saturation axis and can be approximated by a linear function toward the central rod of the color space; i.e., the intensity axis (Hansen et al., 2007). The intuitive features of the HSI space provide strong arguments for the latter assumption: Consider a chromatic point of the outer boundaries of the HSI space (with maximum saturation). When the saturation value is lowered, the perceived color becomes ‘decolorized’ or pale. Nevertheless, in general the color is perceived as belonging to the same color category.
- (3) The two boundaries between achromatic categories can each be expressed with a single intensity value.

Given the latter assumptions, segmentation can be done in three steps. First, the separation of achromatic categories and chromatic categories through a 2D plane leaving out the hue axis resulting in a saturation-intensity plane. Second, the segmentation of the chromatic colors by leav-

ing out the saturation axis: the hue-intensity plane. Third, the segmentation of the individual achromatic categories is performed in a saturation-intensity plane.

4. Feeding the color markers to a distance transform

The previous section presented a framework that utilized two HSI 2D planes. In these 2D planes, the categorized color data will occur as grouped clouds of data points. However, the information that is available about humans’ color categorization through the experiments (see Section 2) does only assign a limited number of points in these planes. Therefore, we applied distance mapping, where each point gets a distance measure to the set of categorized points by humans.

The speed of distance transforms is determined by the precision needed. Images are a priori an approximation of reality due to various forms of noise. It might be that introducing additional approximations in the image processing chain to make it faster, has no effect on the final quality of the application. The best way to test that is to compare with no additional approximations in the chain. Therefore, we preferred an as accurate as possible distance transform, preferably exact.

Distance transforms can be applied on all kinds of data. In this paper, we discuss the 11 color categories, which are generally applicable. However, the categories that are needed depend on the specific application; e.g., a catalog of paintings or a stained cell tissue database. There might be the need to adapt the color categories quickly to specifications based on a certain domain or task. Moreover, the perception of individual users differs and systems are needed that use user profiles, which would be in our case: a user specific (i.e., personalized) color space segmentation. The latter is of great importance since users are in interaction with the systems, which use image analysis techniques, and judge their results. Therefore, we wanted a fast color space segmentation regarding computer and human resources.

The distance transform to be applied both needs to be fast enough and preferably exact. For this purpose, we developed the Fast Exact Euclidean Distance (FEED) transform, which is introduced in the next section.

4.1. Fast Exact Euclidean Distance (FEED)

In contrast with the state-of-the-art approaches such as Shih and Wu’s two scan algorithm (EDT-2) (Shih and Wu, 2004), we have implemented the Euclidean Distance (ED) transform starting directly from its definition. Or rather its inverse: each object pixel q , in the set of object pixels (O), feeds its ED to all non-object pixels p . The naive algorithm then becomes:

```

initialize  $D(p) = \text{if}(p \in O) \text{ then } 0, \text{ else } \infty$ 
foreach  $q \in O$ 
  foreach  $p$ 
    update :  $D(p) = \min(D(p), ED(q, p))$ 

```

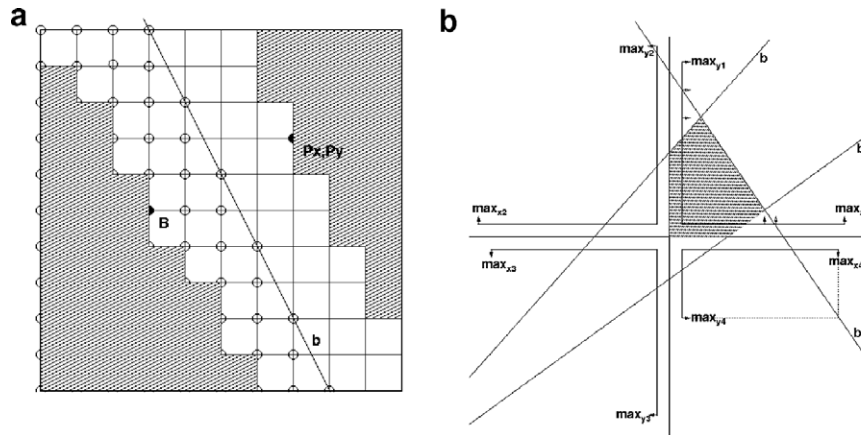


Fig. 4. (a) Principle of limiting the number of background pixels to update. Only the pixels on and to the left of the bisection line b have to be updated. B is the border pixel under consideration, p is an object pixel. (b) Bookkeeping of the sizes (the “max”) of each quadrant. Updating process: On each scan line the bisection lines b determine the range of pixels to update.

This algorithm is extremely time consuming, but is speeded up by:

- Restricting the number of object pixels q that have to be considered, to the border pixels of the objects.
- Pre-computation of $ED(q, p)$.
- Restricting the number of background pixels p that have to be updated for each border pixel using bisection lines (see Fig. 4.)

The method used for searching for other object pixels, bookkeeping of the bisection lines, and determining which background pixels to update is carefully designed (see also Fig. 4). This, to ensure that it takes much less time than the time gained by not updating the other background pixels.

This resulted in an exact but computationally less expensive algorithm for the ED transform: the Fast Exact Euclidean Distance (FEED) transformation. It is recently introduced by Schouten and van den Broek (2004). For both algorithmic and implementation details we refer to this paper.

4.2. Benchmarking FEED

In the latest experiments, we have compared FEED with EDT-2 Shih and Wu (2004) and with the city-block (or Chamfer 1,1) distance, as baseline. In Table 2, the timing

Table 2

Average timing results (in seconds) for three sets of images on the city-block transform, Shih and Wu’s (2004) two scan method (EDT-2), and for FEED (Schouten and van den Broek, 2004)

Images	Timing (errors) per distance transform		
	City-block	EDT-2	FEED
Standard	8.75 s (2.39%)	38.91 s (0.16%)	17.14 s
Rotated	8.77 s (4.66%)	38.86 s (0.21%)	18.02 s
Larger obj.	8.64 s (4.14%)	37.94 s (0.51%)	19.94 s

Between brackets the errors (in %) of the city-block (or Chamfer 1,1) and EDT-2 transform. Note that no errors of FEED are mentioned since FEED provides truly exact EDs.

results are provided for the city-block measure, for EDT-2, and for FEED. As was expected, with a rough estimation of the ED, the city-block distance outperformed the other two algorithms by far. More surprising is that FEED is more than twice as fast as EDT-2 (see Table 2). The time complexity of the city-block and EDT-2 methods is $\mathcal{O}(n)$ with n the number of pixels in the image. FEED behaves over a large range of tested images as having a time complexity of $\mathcal{O}(n)$; however, this could not be proved.

The aim of FEED is to utilize exact ED transforms. Hence, next to the timing results, the percentage of errors made in obtaining the ED is of interest to us. The city-block transform resulted for all images in an error-level of less than 5%; see Table 2. Shih and Wu claimed that their EDT-2 provided exact EDs. However, in 1% of the cases errors occur in their algorithm, as reported in Table 2. So, FEED is the only algorithm that provided the truly exact ED for all instances.

4.3. FEED for multi class data

Now, let us consider the case that multiple labeled classes of data points (e.g., color markers assigned to color categories) are present and, subsequently, FEED is applied for data space segmentation. In such a case, the class of the input pixel that provides the minimum distance is placed in a second output matrix. To achieve this, the update step of FEED is changed to:

$$\text{update : if } (ED(q, p) < D(p)) \\ \text{then } (D(p) = ED(q, p); C(p) = C(q))$$

where C is a class matrix, in our case the colors assigned to one of the 11 color categories.

Similar changes were applied to the city-block and EDT-2 methods; thus, three distance transforms and classification methods are obtained. These methods are then applied on a set of hue-intensity images as used in the remaining of this paper. The results are shown in Table 3.

Table 3

Average timing results in milliseconds on an ADM Athlon XP 2100 machine for a set of hue-intensity images of size 815×941 for the city-block transform, Shih and Wu's (2004) two scan method (EDT-2), and for FEED (Schouten and van den Broek, 2004)

Method	Timing and classification errors		
	DT (ms)	Classification (ms)	% Wrong
City-block	23.3	28.5	1.9
EDT-2	64.4	73.2	0.12
FEED	32.7	41.5	0

The last column gives (in %) the number of wrongly assigned pixels. FEED provides truly exact EDs.

The addition of the classification increases the time for all three methods. As was expected FEED is faster than EDT-2 and slower than the city-block approximation of the ED. But, only FEED provides no misclassified pixels.

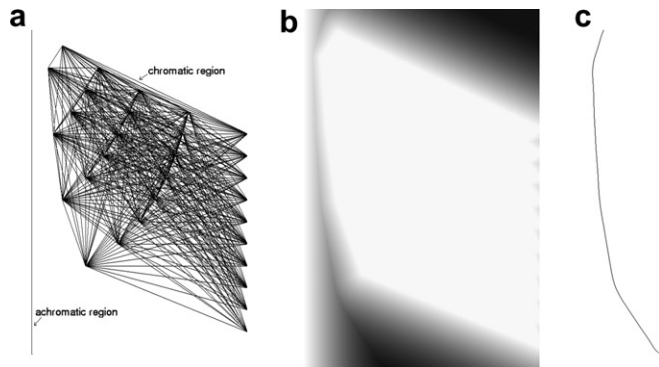


Fig. 5. The processing scheme of the separation of the chromatic from the achromatic color categories, in the saturation-intensity plane, using human color categorization data (see Section 2): (a) The fully connected graphs of the categorized CLUT markers. (b) The weighted distance map, created using Fast Exact Euclidean Distance (FEED) transformations (Schouten and van den Broek, 2004). (c) The resulting chromatic-achromatic border. Note that saturation is presented on the horizontal axis and intensity on the vertical axis.

The minimum distance value then indicates the amount of probability (or weight) that the pixel belongs to the class. This can be visualized by different color ranges, for each class. By extracting the pixels that express minimum probability, a Voronoi diagram can be generated. In our application this diagram defines the borders between the color categories (see Figs. 5c and 6c).

5. The segmentation process

Sections 3 and 4 described the means to do color space segmentation with scarce data. In this section, we describe the actual segmentation process of the color space with as input the color markers from the experiment described in Section 2. Result of this process is a fully segmented color space and, subsequently, a complete Color LookUp Table (CLUT) for the 11 color categories.

The first phase in preprocessing is the conversion of the RGB color markers (see Section 2) to HSI color markers. The conversions as given by Gevers and Smeulders (1999) were adopted:

$$H(R, G, B) = \arctan\left(\frac{\sqrt{3}(G - B)}{(R - G) + (R - B)}\right) \quad (1)$$

$$S(R, G, B) = 1 - \frac{3 \cdot \min(R, G, B)}{R + G + B} \quad (2)$$

$$I(R, G, B) = R + G + B \quad (3)$$

Please note that the original conversion was adapted by the use of a factor 3 instead of 1 in Eq. (2). This changes the range of the saturation (S) from $[\frac{2}{3}, 1]$ to $[0, 1]$. The next phase in preprocessing is the generation of the 2D planes. The 11 color categories of the HSI CLUT were divided into two groups: the achromatic categories (i.e., black, gray, and white) and the chromatic categories (i.e., blue, yellow, green, purple, pink, red, brown, and orange). In this way,

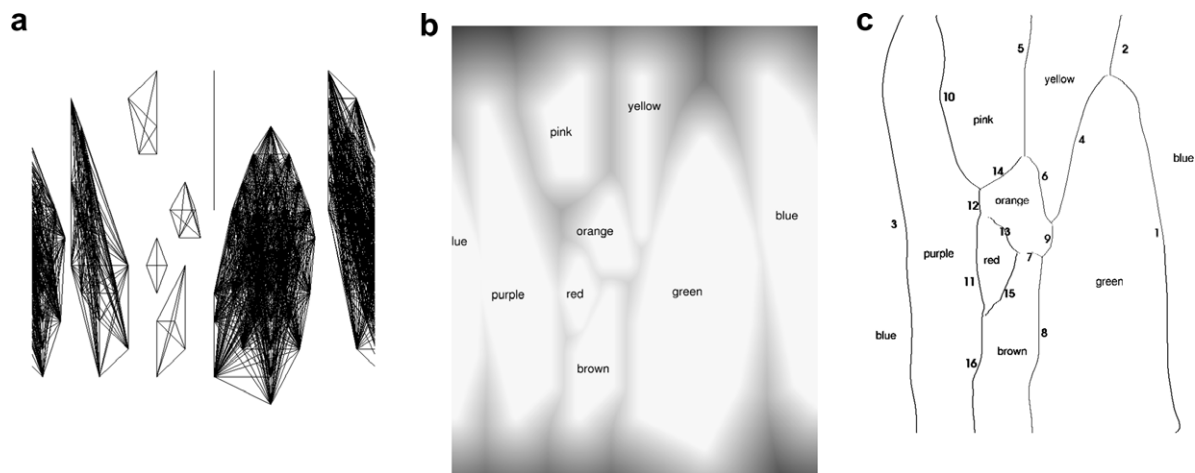


Fig. 6. The processing scheme of the separation of the chromatic color categories in the hue-intensity plane, using human color categorization data (see Section 2) (note that the hue-axis is circular): (a) The fully connected graphs of the categorized CLUT markers. (b) The labeled weighted distance map created using Fast Exact Euclidean Distance (FEED) transformations (Schouten and van den Broek, 2004). (c) The resulting borders between the chromatic color categories. Note that hue is presented on the horizontal axis and intensity on the vertical axis.

categories, as is shown in Table 4 with Ns and Cs. Hence, the color space segmentation mimics human color categorization and correctly classifies both more prototypical colors (i.e., the non-fuzzy color markers) and the colors over which no consensus is among people (i.e., the fuzzy color markers).

7. Discussion

We have explained our approach toward color analysis, which exploits human perception instead of mere image processing techniques. The importance of the 11 color categories (or focal colors) is discussed and sustained by a question and answer and by two experiments. The resulting experimental data (i.e., color markers) is used as input for a coarse color space segmentation process. The HSI color space is segmented, using two 2D projections of the HSI color space on which the recently developed Fast Exact Euclidean Distance (FEED) transform for multi class data is applied. The segmented HSI color space is transformed to a Color LookUp Table (CLUT) for the 11 color categories. We will now discuss its flexibility, a result of its modular implementation, followed by other issues of concern.

The advantage of the color space segmentation method as proposed is threefold: (i) It is based on human color categorization (ii) segmentation can be done on scarce data, and (iii) it is easily adapted to other, application, task, and/or user dependent color categories. Using this segmented color space (or better the CLUT) as a quantization scheme for image retrieval, two more advantages (Jou et al., 2004) can be mentioned: (i) it yields perceptually intuitive results for humans and (ii) it has a low computational complexity.

Color categories, as used in the current research, have also proved to be very useful in research toward color constancy (Hansen et al., 2007). Among other things, Hansen et al. (2007) propose to “infer color constancy from the boundaries between color categories”. So, an executable model of human color categorization could aid research toward another intriguing features of human color perception: color constancy.

The flexibility of the model is illustrated by the possibility to do a smooth personalization of the generic color space segmentation through a quick calibration. This is easily done by utilizing personal judgments of fuzzy color markers that can be used to redefine the boundaries between the color categories. Through the standard processing steps, this finally results in a tailored CLUT. More generally, three other modifications of the processing scheme can be applied easily: (i) another set of colors can be incorporated, which is user (e.g., in the case of color blindness), task, or application specific, (ii) instead of the HSI color space, each arbitrary color space can be used (Derefeldt et al., 2004; Hansen et al., 2007; Schettini et al., 2001), which can be achieved simply by incorporating another conversion scheme in the process; and (iii) the FEED transform can be replaced by an arbitrary alter-

native transform; e.g., Fouard et al. (2007) and Shih and Wu (2004). In general, the indirect segmentation of the 3D color space (by using 2D representations) can be considered as an disadvantage of the current processing scheme. Nevertheless it had two benefits for the segmentation process: (i) a limited number of color markers is sufficient and (ii) its computational costs are much lower. In our opinion this strategy, given the scarce categorization data that was available, leads to the best possible result. In some other setting, a 3D implementation of FEED could possibly limit the number of errors in the end result. In general, the choice whether or not to adopt a direct 3D distance transform is founded on the trade-off between the available color markers, precision, and speed.

The human color categorization model as described in the current paper has been benchmarked with various color quantization schemes and distance measures (van den Broek et al., 2005). Moreover, since 2005, the model of human color categorization serves as the foundation of the online color-based image retrieval system <http://www.M4ART.org> (van den Broek et al., 2006). M4ART contains over 30,000 photos, mainly of pieces of art, provided by the Rijksmuseum (National museum of the Netherlands) and ArtStart.nl, which coordinates the Dutch art rental centra.

In general, it is our belief that the combination of human perception and statistical image processing will improve the performance of data-mining systems that rely on color analysis. Moreover, such a combination enables us to eventually bridge the semantic gap present in automated color analysis. From this perspective, a unique color space segmentation is proposed, which is generic, computationally inexpensive, and easy to tune or to personalize.

Acknowledgements

This research was supported by the Netherlands Organization for Scientific Research (NWO) under project number 634.000.001. We thank Leon van den Broek, Frans Gremmen, Thijs Kok, Makiko Sadakata, Louis Vuurpijl (Radboud University Nijmegen), Merijn van Erp (Planon B.V.), Maarten Hendriks (MEDOX.nl), and Eva van Rikxoort (Image Sciences Institute) for their support.

References

- Cox, I.J., Miller, M.L., Omohundro, S.M., Yianilos, P.N. PicHunter: bayesian relevance feedback for image retrieval. In: Kropatsch, W.G., Aloimonos, Y., Bajcsy, R. (Eds.), Proceedings of International Conference on Pattern Recognition, pp. 361–369. Vienna, Austria, August 1996.
- Derefeldt, G., Swartling, T., Berggrund, U., Bodrogi, P., 2004. Cognitive color. *Color Research & Application* 29 (1), 7–19.
- Flickner, M., Sawhney, H., Niblack, W., Ashley, J., Huang, Q., Dom, B., Gorkani, M., Hafner, J., Lee, D., Petkovic, D., Steele, D., Yanker, P., 1995. Query by Image and Video Content: The QBIC system. *IEEE Computer* 28 (9), 23–32.

- Fouard, C., Strand, R., Borgefors, G., 2007. Weighted distance transforms generalized to modules and their computation on point lattices. *Pattern Recognition* 40 (9), 2453–2474.
- Gevers, Th., Smeulders, A.W.M., 1999. Color based object recognition. *Pattern Recognition* 32 (3), 453–464.
- Hansen, T., Walker, S., Gegenfurtner, K.R., 2007. Effects of spatial and temporal context on color categories and color constancy. *Journal of Vision* 7(4) (2), 1–15.
- Jou, F.-D., Fan, K.-C., Chang, Y.-L., 2004. Efficient matching of large-size histograms. *Pattern Recognition Letters* 25 (3), 277–286.
- Kay, P., Regier, T., 2006. Language, thought and color: Recent developments. *Trends in Cognitive Sciences* 10 (2), 51–54.
- Lin, T., Zhang, H.J. Automatic video scene extraction by shot grouping. In: Sanfeliu, A., Villanueva, J.J. (Eds.), In: *Proceedings of the 15th IEEE International Conference on Pattern Recognition*, vol. 4, pp. 39–42, Barcelona, Spain, 2000.
- Petrushin, V.A., Khan, L., 2007. *Multimedia Data Mining and Knowledge Discovery*. Springer-Verlag, Berlin-Heidelberg, ISBN-13: 978-1-84628-436-6.
- Schettini, R., Ciocca, G., Zuffi, S., 2001. *A Survey of Methods for Colour Image Indexing and Retrieval in Image Databases*. J. Wiley.
- Schouten, Th. E. van den Broek, E.L. Fast Exact Euclidean Distance (FEED) Transformation, In: J. Kittler, M. Petrou, and M. Nixon, (Eds.), *Proceedings of the 17th IEEE International Conference on Pattern Recognition (ICPR 2004)*, vol. 3, pp. 594–597, Cambridge, United Kingdom, 2004.
- Shih, F.Y., Wu, Y.-T., 2004. Fast Euclidean distance transformation in two scans using a 3×3 neighborhood. *Computer Vision and Image Understanding* 93 (2), 195–205.
- van den Broek, E.L., Kisters, P.M.F., Vuurpijl, L.G., 2005. Content-based image retrieval benchmarking: Utilizing color categories and color distributions. *Journal of Imaging Science and Technology* 49 (3), 293–301.
- van den Broek, E.L. Kok, T. Schouten, Th. E., Hoenkamp, E. Multimedia for Art ReTrieval (M4ART) In: *Proceedings of SPIE (Multimedia Content Analysis, Management, and Retrieval)*, 6073, 2006, 60730Z.
- W3 Schools. HTML Colors, URL: http://www.w3schools.com/html/html_colors.asp (accessed 19.07.07).