Normalised Color Distances

André R. S. Marçal[1,2][a]
[1]*Departamento de Matemática, Faculdade de Ciências da Universidade do Porto, Rua do Campo Alegre s/n, 4169-007, Porto, Portugal*
[2]*INESC TEC, Porto, Portugal*

Keywords: Color Distance, Color Perception, Color Normalisation, Dunn Index.

Abstract: This paper presents normalised color distances based on widely used metrics in the RGB and L*a*b* color models, and an adjusted City Block distance for the HSV model. Three experiments were carried out, focusing on color perception, the identification of the actual range of the various normalised color distances, and their ability of compare and match images which have a predominant color perceived by a human observer. For this task, a spatially tolerant color distance is proposed. The image comparison experiment uses subsets of 6 images, out of 15 square tile images, with a total of 270 test cases considered. A modified Dunn index is proposed for the evaluation. L*a*b* based distances were found to be better adjusted to the human color perception. Color distances based on L*a*b* model were also more effective for image comparison, with spatially tolerant color distances having slightly better performance than using a direct image pixel pairing.

1 INTRODUCTION

Color plays a key role in the visual interpretation and understanding of images. The human perception of color is however highly subjective. Our visual system has a tendency to keep its perceptions invariant to illumination changes (Vanrell et al., 2011). Furthermore, the perception of color differences is distorted by categorical boundaries (to which we associate color names) (Hu et al., 2014). Image processing and computer vision systems also make use of the information provided by a color image, rather than using only a gray-scale (single band) version of the image. There are several mathematical/computational models (or spaces) for the representation of color, but the relation between these models and the human perception of color is not straightforward. Nevertheless, some of these colour spaces are optimized to correspond as best as possible to the human perception (Brychtova and Coltekin, 2015).

The process of measuring color differences must be designed to balance between the computed and the perceived difference (Vertan et al., 2003). A color distance provides a numeric representation of the similarity (or difference) between two colors. This distance can be used in relative terms or, in some contexts, as an absolute measurement where a normalised version is preferable. Despite the many color models and metrics available, there is not an establish color distance that exactly matches the human perception of color similarity / difference, as they rely on a simplified model of the world that allows these models to isolate the capabilities of the human visual system from the complexities introduced by real-world viewing (Sharma et al., 2005). Presently, the CIEDE2000 color-difference formula (ΔE_{00}) is regarded as the best, coinciding with subjective visual perception (Brychtova and Coltekin, 2015). However, as a number of discontinuities occur in the CIEDE2000 implementations (Sharma et al., 2005), the triangular inequality is not always satisfied and thus ΔE_{00} is in fact not a metric. Simpler color differences, such as the Euclidean distance, can thus be popular choices for color analysis and visualization tools (Szafir, 2018).

The objective of this paper is to present and evaluate normalised color distances. After this introduction, it has four additional sections: in section 2 the color models and distances used are presented; section 3 describes two small experiments focusing on color perception and the actual range and distribution of each color distance; section 4 presents an experiment to evaluate the ability of the color distances to compare images; and section 5 draws the conclusions.

[a] https://orcid.org/0000-0002-8501-0974

Marçal, A.
Normalised Color Distances.
DOI: 10.5220/0011965700003497
In *Proceedings of the 3rd International Conference on Image Processing and Vision Engineering (IMPROVE 2023)*, pages 134-141
ISBN: 978-989-758-642-2; ISSN: 2795-4943
Copyright © 2023 by SCITEPRESS – Science and Technology Publications, Lda. Under CC license (CC BY-NC-ND 4.0)

2 COLOR MODELS AND COLOR DISTANCES

A number of color models and metrics are presented next, without aiming in any way to be a comprehensive review of the topic. More details about color models and the conversion between models can be found in several digital image processing papers and textbooks, such as (Gonzalez and Woods, 2008).

2.1 Color Models

For digital image processing, additive models are generally preferred when working on standard displays (emitting light), whereas subtractive models are more suitable when the focus is printed media. Most color models have 3 independent variables, in line with the 3 degrees of freedom of the human visual system (Gonzalez and Woods, 2008).

The most widely used color model in digital image processing is probably the RGB color model (Bratkova et al., 2009), where the 3 variables R (Red), G (Green) and B (Blue) are the 3 primary colors of this additive model. For digital images, the intensity of each of these variables is quantised in a number of discrete values. For example, a RGB 24-bit color image has 8-bit for each color component, allowing for a total of 2^{24} number of different color to be represented ($\approx$ 16.7 millions).

One disadvantage of the RGB color model is that all 3 components are responsible for both color and intensity (brightness). There are however some models where the color component is decoupled from the intensity, such as the HSV model, where H (Hue) and S (Saturation) are related only to color, and V (Value) only to intensity (Gonzalez and Woods, 2008). Another example is the CIELAB color space, also referred to as L*a*b* (or $L\alpha\beta$) model (Reinhard and Pouli, 2011). The variable $L*$ is related to lightness, whereas the color is represented by $a*$ and $b*$, which are related to four unique colors of human vision: red, green, blue and yellow.

2.2 Color Distances

There are several possibilities for establishing distances in the various color models, including the Minkowsky distance or L-norm. The most basic versions are L1 (City Block or Manhattan distance) and L2 (Euclidean distance).

Often distances are used for comparisons in a context where only relative values are needed. For example, to select the the best match to a color, out of a set of reference colors. However, in some cases it might be relevant to have not only the best match (the color reference with the lowest distance), but also a numeric value that can provide a reliable measurement of the similarity between two colors. For this purpose, it is preferable to use normalized color distances, where the co-domain of the distance function is $[0,1]$. The most dissimilar colors possible will have a distance of 1, and the color distance will tend to 0 as the colors become more and more similar.

In the RGB model, a normalised distance is defined as $d_{RGB} : [0,1]^3 \rightarrow [0,1]$. The Euclidean distance (d^E_{RGB}) is computed by (1) and the City Block distance (d^{CB}_{RGB}) by (2), where R_i, G_i, B_i are the color intensities for element i (i=1,2).

$$d^E_{RGB} = \frac{1}{\sqrt{3}} \sqrt{(R_1 - R_2)^2 + (G_1 - G_2)^2 + (B_1 - B_2)^2} \tag{1}$$

$$d^{CB}_{RGB} = \frac{1}{3} \Big(|R_1 - R_2| + |G_1 - G_2| + |B_1 - B_2| \Big) \tag{2}$$

A normalised distance in the HSV model is defined as $d_{HSV} : [0,1]^3 \rightarrow [0,1]$. The domain is $[0,1]$ for all three model components, similarly to RGB. However, as Hue is an angular measurement, the difference in Hue $|H_1 - H_2|$ cannot be calculated directly (Montenegro et al., 2008), as that does not properly account for the artificial discontinuity in Hue imposed by the linear domain $[0,1]$. For example, two colors with $H_1 = 0.01$ and $H_2 = 0.99$ (and equal S and V) are almost indistinguishable reds, yet $|H_1 - H_2| = 0.98$. In fact, the Hue difference (Δ_{Hue}) in this case should be only 0.02 if the circular nature of the variable is considered, and thus Δ_{Hue} is computed by (3). An adjusted City Block distance (d^{ACB}_{HSV}) is proposed for the HSV model (4). The factor of 2 for Δ_{Hue} in (4) assures that the contribution from each parcel is the same, as the range of Δ_{Hue} is $[0,0.5]$ instead of $[0,1]$.

$$\Delta_{Hue} = min\Big\{ |H_1 - H_2|, 1 - |H_1 - H_2| \Big\} \tag{3}$$

$$d^{ACB}_{HSV} = \frac{1}{3} \Big(2 \times \Delta_{Hue} + |S_1 - S_2| + |V_1 - V_2| \Big) \tag{4}$$

In the L*a*b* model, the domain of the variables a* and b* is $[-1,1]$, thus a normalised distance is defined as $d_{Lab} : [0,1] \times [-1,1] \times [-1,1] \rightarrow [0,1]$. The Euclidean distance in the L*a*b* model (d^E_{LAB}) is computed by (5) and the City Block distance (d^{CB}_{LAB}) by (6). These equations are very similar to those for the RGB model (1, 2), except for the normalisation coefficients that are different. One alternative often

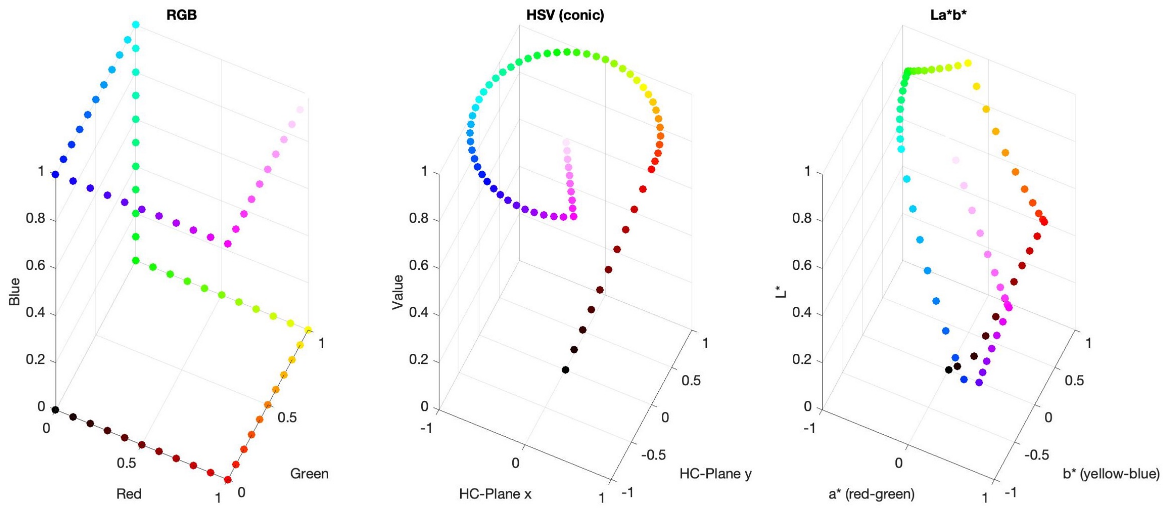

Figure 1: Color test sequence (uniformly spaced in RGB) representation in the 3D Cartesian space for the RGB model (left), HSV conic (centre) and L*a*b* model (right).

used in the L*a*b* color model is the Hybrid distance (Abasi et al., 2020). The normalised Hybrid distance in L*a*b* (d^H_{LAB}) is defined by (7). It is computed as an Euclidean distance in the a*b* plane combined with a City Block distance in what regards this plane and L*.

$$d^E_{Lab} = \frac{1}{3}\sqrt{(L^*_1 - L^*_2)^2 + (a^*_1 - a^*_2)^2 + (b^*_1 - b^*_2)^2} \quad (5)$$

$$d^{CB}_{Lab} = \frac{1}{5}\left(|L^*_1 - L^*_2| + |a^*_1 - a^*_2| + |b^*_1 - b^*_2|\right) \quad (6)$$

$$d^H_{Lab} = \frac{|L^*_1 - L^*_2| + \sqrt{(a^*_1 - a^*_2)^2 + (b^*_1 - b^*_2)^2}}{1 + 2\sqrt{2}} \quad (7)$$

It is worth mentioning that the actual range of the color distances is a subset of the co-domain. Another aspect that should be considered is that the domain amplitude of the variables responsible for the color information in the L*a*b* model is twice the amplitude of the domain of L^*. So the relative contribution of $a*$ and $b*$ to the color distance is potentially larger than L^*, thus making L*a*b* based distances more influenced by the color component than those distances based on the RGB and HSV models.

A normalised version of the CIEDE2000 color-difference (ΔE_{00}) is used for comparison purposes, as CIEDE2000 is considered closely aligned with the human color perception. A normalisation factor of 125 was used (for 24-bit color images), which for all tests carried out produced values within the $[0,1]$ range.

3 INITIAL EVALUATION

The normalised color distances proposed for RGB, HSV and L*a*b* color models were evaluated with the aim of identifying some of their potential advantages and limitations. Three experiments were designed and implemented using Matlab (MathWorks, 2021). The first experiment focus on the relation between distance and color perception, and the second one on identifying the actual range and distribution of each color distance. The third experiment, presented in a separate section, was designed to evaluate the ability of the color distances to compare images.

3.1 Color Perception

A sequence of uniformly spaced RGB colors was created. The first color is black (0,0,0), with increments of 0.1 used for only one of the R,G,B components at a time. The distance between any two consecutive colors in the sequence is thus constant for both Euclidean and City Block metrics. For the standard definitions, the distance is 0.1. However, due to the normalisation factors in (1) and (2), the values for normalised distances are different: $d^E_{RGB} = 0.0577$ and $d^{CB}_{RGB} = 0.0333$. The sequence of colors is presented in the 3D Cartesian space for the RGB model in figure 1 (left). This test sequence has a total of 71 colors along 7 edges of the RGB cube, including all primary (red, green, blue) and secondary colors (cyan, magenta, yellow), as well as black and white.

In figure 1 the color sequence is also presented in the HSV model (centre) and L*a*b* model (right).

SCITEPRESS SCIENCE AND TECHNOLOGY PUBLICATIONS

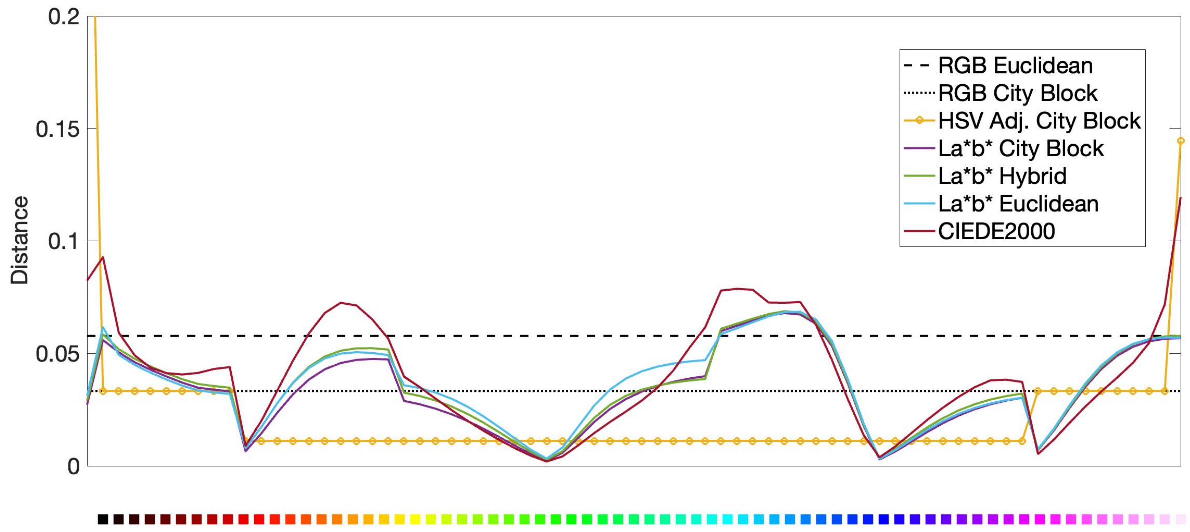

Figure 2: Color test sequence (bottom) and plot of the normalised distances for color pairs.

The HSV representation used for figure 1 is conic, although other geometries could be used, such as cylindrical or pyramidal (Gonzalez and Woods, 2008). It is worth noting that the Hue (H) and Saturation (S) components correspond to polar coordinates in the HS plane, thus the visual perception of distances from the figure is slightly misleading.

The color sequence presentation in the 3D Cartesian space for the L*a*b* model (figure 1, right) is particularly interesting. Unlike for the HSV model, the distances perceived in the L*a*b* 3D space directly relate to the L*a*b* Euclidean distance (d^E_{LAB}). It is clearly noticeable that some color pairs are very close together, such as those in the green region (top left), whereas some color pairs are much further away. For example, in the sequence of colors form blue to green the distance between consecutive colors is considerably larger than color pairs with two types of green. These differences in distances between consecutive color pairs seem to be better matched to our perception of color similarity than the constant values provided by the RGB based distances.

Figure 2 shows an alternative graphic presentation of the color sequence, together with a plot of the distances between consecutive pairs. In this figure it is possible to observe each color pair (in the bottom scale) and the corresponding normalised color distances. The black dotted and dashed lines are the distances computed in the RGB model, both having a constant value throughout.

The adjusted City Block distance for the HSV model (4) is presented in Figure 2 as an orange solid line, with large dots. The distance for the first color pair is very high due to a change of saturation (S) from 0 to 1. The remaining pairs have constant values along sub-sections of the sequence, with the lowest values in the centre of the sequence, where only the Hue component changes.

The 3 distances based on the L*a*b* model are presented in Figure 2 with solid lines without dots. There are some difference between them, but not significant. The distances are much lower for some color pairs than for others, which seems to be well aligned with the perceived color difference, considering the human interpretation perspective. The normalised version of CIEDE2000 (ΔE_{00}) is also presented in Figure 2. The general behaviour of ΔE_{00} is well aligned with the 3 L*a*b* distances, but with even larger values for highly dissimilar color pairs.

3.2 Color Distances Distribution

The second experiment consists of a brief statistical evaluation of the 6 normalised distances, using N RGB color pairs generated randomly. Initially one million color pairs were created ($N = 10^6$), with the color distances computed used to create the boxplots presented in Figure 3. The mean, standard deviation and 0.1 and 99.9 percentiles (Pctl) were computed for a larger set of random color pairs ($N = 10^8$), with the results presented in Table 1.

One first aspect that can be observed, is that the median (Figure 3) and mean (Table 1) are lower for the L*a*b* distances that for those based on the RGB and HSV models. The interquartile range is also lower for L*a*b* distances, particularly for d^{CB}_{LAB} and d^H_{LAB}. But perhaps the most relevant issue is the fact that the range of values actually used is far from covering the co-domain $[0,1]$. This can be observed by the whiskers in the boxplots, and by the differ-

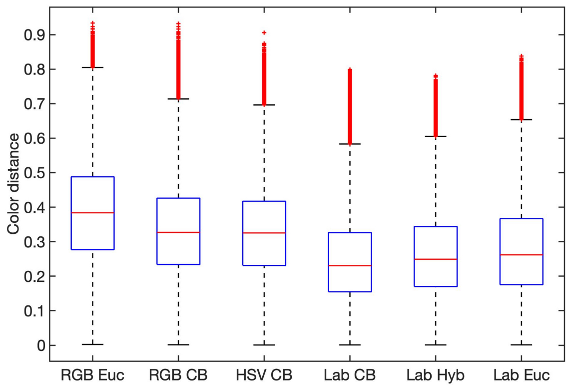

Figure 3: Boxplots for normalised distances of one million random color pairs.

ence between the 99.9 and 0.1 percentiles. The RGB Euclidean distance have the largest range, and the L*a*b* City Block and Hybrid the lowest.

Table 1: Statistic parameters for normalised distances of 10^8 random color pairs.

	Pctl 0.1	Pctl 99.9	Mean	St.Dev.
d^{E}_{RGB}	0.0369	0.7913	0.3835	0.1445
d^{CB}_{RGB}	0.0314	0.7804	0.3346	0.1366
d^{ACB}_{HSV}	0.0293	0.7322	0.3275	0.1304
d^{CB}_{Lab}	0.0163	0.7035	0.2511	0.1293
d^{H}_{Lab}	0.0182	0.6930	0.2650	0.1273
d^{E}_{Lab}	0.0184	0.7376	0.2795	0.1373

4 IMAGE COMPARISON

An experiment was designed to evaluate the ability of the normalised color distances to compare images and to identify the best image match. A total of 15 images of square tiles were used, all with 267 × 267 pixels and 24-bit RGB color. The images were grouped by visual interpretation into 5 classes, based on the predominant color present: yellow (A), red/pink (B), blue (C), brown (D) and green (E). Figure 4 shows the 15 images with the class assignment (3 for each class) and label. Despite the diversity of colors in some tiles (e.g. blue and white is present in tiles A1 and A2, both labeled as yellow), there is a predominant color on each tile that is likely the dominant aspect used for the class labeling made by human interpretation.

4.1 Image Similarity

The evaluation of the similarity between two images is based on the mean normalised color distance between all image pixel pairs. However, the direct pairing of pixels form two images might not be the most suitable approach, as there might be a small geometric misalignment between the two images. It is thus worth considering a small tolerance in what regards the positioning of the pixel pairing. To address this issue, a spatially tolerant color distance $D^{(v)}$ is proposed, considering a neighborhood v of a pixel. It is defined by (8), where each pixel (x,y) in image I_1 is compared with all pixel (x',y') in image I_2 that are within the neighborhood v of (x,y). The color distance $D^{(v)}$, for a pixel (x,y), is the smallest normalised color distance of all (x,y) - (x',y') pairs.

$$D^{(v)}_{12}(x,y) = min\Big\{d\{I_1(x,y), I_2(x',y')\}\Big\} \quad \text{with } (x',y') \in v(x,y) \qquad (8)$$

The neighborhoods used here were: direct pairing ($v=1$), where $x=x' \wedge y=y'$; 4-neighbors ($v=4$), where $|x-x'|+|y-y'| = \{0,1\}$; 8-neighbors ($v=8$), where $x-x' = \{-1,0,1\} \wedge y-y' = \{-1,0,1\}$. A total of 18 color distances were thus considered: 6 normalised color distances (1-7) × 3 neighborhoods ($v=1$, $v=4$ and $v=8$).

4.2 Experimental Procedure

Initially a set of 6 tiles is selected from the image set (Figure 4), with 3 classes used and 2 images from each class selected. The color distances between all pairs of images are then computed. For each image, the shortest distance is used to establish the best match. Ideally, each image would be matched with the other image in the set belonging to the same class, but that does not always happen. In fact, the best case scenario would be to have a very low distance for a pair of images of the same class, and large distances when two images belong to different classes.

4.3 Evaluation Criteria

An evaluation parameter inspired by the Dunn similarity index used for data clustering (Dunn, 1973) is proposed - the Modified Dunn Index. It is based on internal distances (for observations of the same class) and external distances (for observations of different classes). The Modified Dunn Index (*MDI*) is computed by (9), where $C(i)$ is the class of element i. *MDI* is the ratio of the minimum external and maximum internal distances, which one aims at maximising.

$$MDI_i = \frac{min\{d(i,j)\}}{max\{d(i,k)\}} \quad \text{with } C(j) \neq C(i)\,,\, C(k) = C(i) \qquad (9)$$

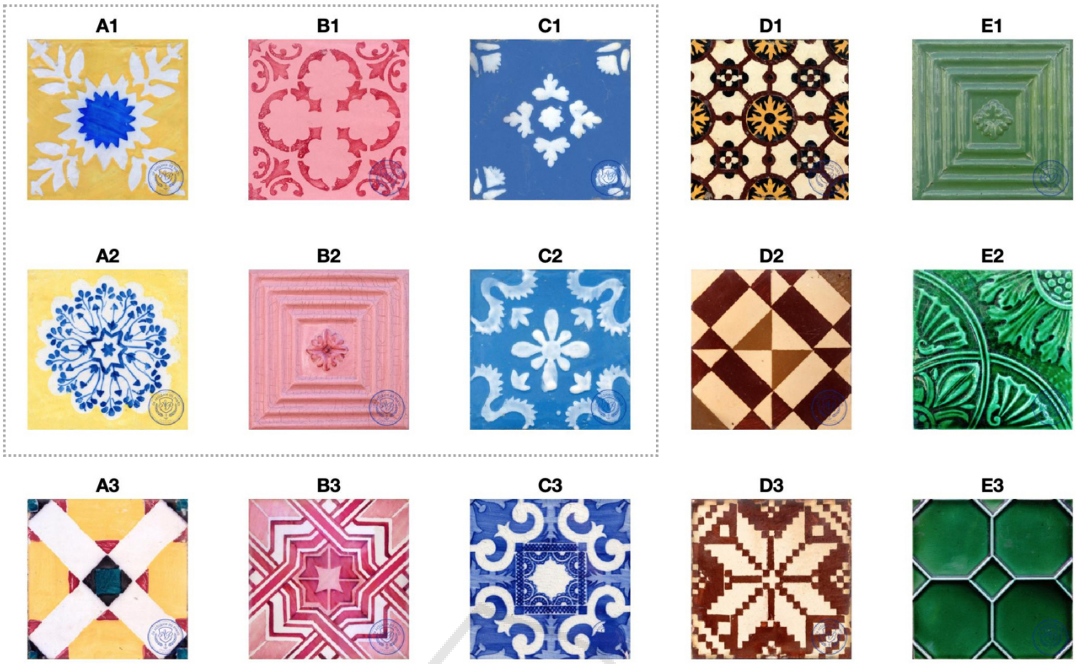

Figure 4: Test images (square tiles), grouped in 5 classes: yellow (A), red/pink (B), blue (C), brown (D) and green (E).

For each image in a test set, there are 4 external distances, and only one internal distance to consider. If $MDI < 1$ for an image, it means that it is mismatched, as there is an image from another class in the set with a shorter distance than the other image of its class.

4.4 Results for a Test Case

To illustrate the procedure, a test case (TC1) is presented in some detail next. The image tiles of TC1 are from classes yellow ($A1$, $A2$), red/pink ($B1$, $B2$) and blue ($C1$, $C2$), identified in the top left corner of Figure 4 by a gray dotted line.

Table 2: Color distances D^1 for the test case image pairs, with Normalised RGB Euclidean distance (d^E_{RGB}).

Tile	B1	C1	A2	B2	C2
A1	0.267	0.395	0.240	0.228	0.386
B1		0.383	0.302	0.130	0.394
C1			0.382	0.356	0.207
A2				0.263	0.359
B2					0.354

The Color distances D^1 for each image pair are presented in Table 2, using the normalised RGB Euclidean distance (and $v = 1$). For image $A1$, the internal distance is $D^1\{A1,A2\} = 0.240$, and the minimum external distance is $D^1\{A1,B2\} = 0.228$, resulting in $MDI = 0.951$ (< 1). Thus, somehow surprisingly, tile $A1$ is matched to tile $B2$ (with d^E_{RGB}, $v = 1$), possibly due to the fact that these tiles are slightly darker.

The MDI values for all color distances and tiles in TC1 are presented in Table 3, for a direct pairing of image pixels ($v = 1$). For most tiles, the MDI value is high (well above 1) for all color distances (e.g. tile B1). As it happens, the only MDI values below 1 (class mismatch) are for tile $A1$ using RGB based distances. The results clearly indicate that the HSV and L*a*b* based distances are more effective (higher MDI values), with the L*a*b* distances performing better than d^{ACB}_{HSV} for the harder cases (tiles $A1$ and $A2$).

Table 3: MDI values for all color distances and tiles in test case 1, using $D^{(1)}$ (direct pixel pairing).

	A1	B1	C1	A2	B2	C2
d^E_{RGB}	0.95	2.06	1.72	1.10	1.76	1.71
d^{CB}_{RGB}	0.92	1.86	1.43	1.08	1.59	1.43
d^{ACB}_{HSV}	1.04	3.19	2.29	1.13	3.06	2.23
d^{CB}_{Lab}	1.32	3.13	2.23	1.38	2.74	2.05
d^H_{Lab}	1.18	2.94	1.99	1.31	2.59	1.87
d^E_{Lab}	1.23	3.29	2.09	1.26	2.92	2.07

The procedure was repeated using the spatially tolerant color distance $D^{(v)}$ with neighborhoods (v) 4 and 8. The results are presented in Table 4 ($v = 4$) and Table 5 ($v = 8$). The results for $v = 4$ are better than using a direct pixel pairing ($v = 1$) for all cases (tiles and color distances). The results for $v = 8$ are equal or better than for $v = 4$.

Table 4: *MDI* values for all color distances and tiles in test case 1, using $D^{(4)}$ (neighborhood of 4).

	A1	B1	C1	A2	B2	C2
d^{E}_{RGB}	1.04	2.30	1.84	1.20	2.00	1.83
d^{CB}_{RGB}	1.03	2.12	1.53	1.19	1.85	1.57
d^{ACB}_{HSV}	1.18	3.49	2.30	1.19	3.44	2.38
d^{CB}_{Lab}	1.50	3.39	2.31	1.54	3.01	2.17
d^{H}_{Lab}	1.33	3.19	2.15	1.48	2.86	2.04
d^{E}_{Lab}	1.39	3.59	2.27	1.42	3.23	2.27

Table 5: *MDI* values for all color distances and tiles in test case 1, using $D^{(8)}$ (neighborhood of 8).

	A1	B1	C1	A2	B2	C2
d^{E}_{RGB}	1.09	2.36	1.85	1.25	2.08	1.88
d^{CB}_{RGB}	1.08	2.20	1.59	1.25	1.93	1.63
d^{ACB}_{HSV}	1.24	3.47	2.30	1.21	3.54	2.45
d^{CB}_{Lab}	1.57	3.44	2.31	1.60	3.09	2.23
d^{H}_{Lab}	1.39	3.24	2.23	1.55	2.93	2.11
d^{E}_{Lab}	1.46	3.65	2.36	1.49	3.31	2.36

It is worth noting that using a spatially tolerant image comparison, the distance between two images decreases as the neighborhood size increases. Thus, for any given image pair, $D^{(8)} \leq D^{(4)} \leq D^{(1)}$. However, for TC1 the MDI values of spatially tolerant color distances are higher (Tables 3 - 5). In this case, the best choice for pairing tiles according to their color classes is to use a neighborhood of 8 ($D^{(8)}$).

4.5 Global Results

A total of 270 test cases were created by grouping the 15 image tiles presented in Figure 4. This results from 10 possible class choices (combinations of 3 out of 5), with 27 tile selections possible for each class setting (combinations of 2 out of 3 for each class). The same procedure described for test case 1 was applied to all test cases. A total of 1620 observations were thus obtained (6 tiles $\times$ 270 test cases), for each color distance and neighborhood.

A summary of the *MDI* results is presented in Table 6, for a direct pairing of image pixels ($v = 1$). The table include the minimum, mean and median *MDI* values for each color distance, and the total number of failures (image mismatches, $MDI < 1$). The worst case for each distance (minimum *MDI* value) is better for L*a*b* based distances, although still with a value below 1. The number of failures is high for RGB based distances (about 20%) and rather low for L*a*b* based distances. The mean and median *MDI* values are also better (higher) for L*a*b* distances, with RGB based distances performing worst. The performance of ΔE_{00} is in line with d^{CB}_{Lab}, d^{CB}_{Lab} and d^{CB}_{Lab}, with slightly higher mean and median *MDI*, but also with a larger number of failures.

Table 6: Summary of *MDI* results for the complete experiment, with $v = 1$ (1620 observations).

	min.	mean	median	No. fails (%)
d^{E}_{RGB}	0.857	1.374	1.284	294 (18.1%)
d^{CB}_{RGB}	0.805	1.324	1.248	327 (20.2%)
d^{ACB}_{HSV}	0.705	1.542	1.348	174 (10.7%)
d^{CB}_{Lab}	0.942	1.602	1.510	36 (2.2%)
d^{H}_{Lab}	0.911	1.564	1.481	45 (2.8%)
d^{E}_{Lab}	0.920	1.606	1.516	72 (4.4%)
ΔE_{00}	0.922	1.663	1.613	86 (5.3%)

A summary of the *MDI* results using the spatially tolerant color distance $D^{(v)}$ is presented in Table 7 for a neighborhood of 4, and in Table 8 for a neighborhood of 8. The results are slightly better for $v = 8$, but the differences are small.

Table 7: Summary of *MDI* results for the complete experiment, with $v = 4$ (1620 observations).

	min.	mean	median	No. fails (%)
d^{E}_{RGB}	0.854	1.456	1.384	232 (14.3%)
d^{CB}_{RGB}	0.811	1.410	1.335	316 (19.5%)
d^{ACB}_{HSV}	0.690	1.631	1.439	183 (11.3%)
d^{CB}_{Lab}	0.957	1.716	1.624	36 (2.2%)
d^{H}_{Lab}	0.922	1.672	1.592	54 (3.3%)
d^{E}_{Lab}	0.942	1.723	1.635	68 (4.2%)
ΔE_{00}	0.917	1.802	1.718	101 (6.2%)

Table 8: Summary of *MDI* results for the complete experiment, with $v = 8$ (1620 observations).

	min.	mean	median	No. fails (%)
d^{E}_{RGB}	0.824	1.494	1.429	230 (14.2%)
d^{CB}_{RGB}	0.787	1.449	1.371	294 (18.1%)
d^{ACB}_{HSV}	0.685	1.666	1.471	183 (11.3%)
d^{CB}_{Lab}	0.967	1.765	1.700	36 (2.2%)
d^{H}_{Lab}	0.928	1.720	1.654	45 (2.8%)
d^{E}_{Lab}	0.949	1.773	1.676	77 (4.8%)
ΔE_{00}	0.892	1.861	1.763	105 (6.5%)

The mean and median *MDI* values are better for the spatially tolerant distances than using a direct pixel pairing ($v = 1$), for all color distances. The MDI of the worst case (minimum value) is little changed for all distances, with small positive and negative variations observed. The number of failures is improved for the RGB based distances, but only marginally changed for the other color distances. Overall, the best choice is the spatially tolerant $D^{(8)}$ L*a*b* City Block color distance (d^{CB}_{Lab}).

5 CONCLUSIONS

This paper presents 6 normalised color distances, based on widely used metrics in the RGB and L*a*b* color models. An adjusted City Block normalised color distance is proposed for the HSV model. The color distances were evaluated with 3 experiments. The first one, focusing on color perception, clearly indicated that the L*a*b* based distances are much better than those based on RGB and HSV for the evaluation of color similarity (and difference), being more closely aligned with the human color perception. The differences between the L*a*b* distances themselves were found to be negligible, in what regards the color perception evaluation performed.

The second experiment showed that although the normalised distances all have values between 0 and 1 potentially, in reality the range of values used is much smaller. This is particularly noticeable for the L*a*b* based distances. This fact is irrelevant if the distance is used in a relative context, but it can be an issue when the color distance is used as an absolute measurement. A possible solution could be to remap the distance, for example with a linear transformation mapping the 0.1 and 99.9 percentiles (Table 1) to $[0,1]$. This would result in some saturation (of 0.2% of the elements), which could be reduced by using more extreme percentile values (e.g. 0.01 and 99.99).

The third experiment was designed to evaluate the ability of the color distances to compare and identify the best image match. The test images selected have some diversity, but they also have a predominant color and are thus easily matched in color classes by a human observer. The goal was to verify the effectiveness of the normalised color distances to perform the same task. For this experiment, a spatially tolerant color distance $D^{(v)}$ was proposed, to account for a possible geometric misalignment between two images being compared, as well as a modified Dunn index for the evaluation of the results. The modified Dunn index was found to be an useful tool, allowing for large number of image (and color) comparisons to be summarised effectively. The spatially tolerant color distance $D^{(v)}$ was found to be slightly better than a comparison of images with a direct pixel by pixel pairing. The L*a*b* based distances proved to be much better than those based on the HSV and RGB color models for the comparison of images, with the L*a*b* City Block distance (d_{Lab}^{CB}) having the best performance.

REFERENCES

Abasi, S., M., A. T., and Fairchild, M. (2020). Distance metrics for very large color differences. *Color Research & Application*, 45:208–223.

Bratkova, M., Boulos, S., and Shirley, P. (2009). orgb: A practical opponent color space for computer graphics. *IEEE Computer Graphics and Applications*, 29(1):42–55.

Brychtova, A. and Coltekin, A. (2015). Discriminating classes of sequential and qualitative colour schemes. *International Journal of Cartography*, 1(1):62–78.

Dunn, J. (1973). A fuzzy relative of the isodata process and its use in detecting compact well-separated cluster. *Journal of Cybernetics*, 3(3):32–57.

Gonzalez, R. and Woods, R. (2008). *Digital Image Processing*. Pearson, 3rd edition.

Hu, Z., Hanley, J., Zhang, R., Liu, Q., and Roberson, D. (2014). A conflict-based model of color categorical perception: evidence from a priming study. *Psychonomic Bulletin and Review*, 21(5):1214–1223.

MathWorks (2021). Matlab, ver r2021a. MATLAB Central File Exchange. Retrieved September 7, 2022.

Montenegro, A., Calixto, E., Conci, A., and Clua., E. (2008). Introducing a new metric for automatic true color images granulometry. In *Proceedings of IWSSIP 2008 - 15th International Conference on Systems, Signals and Image Processing*.

Reinhard, E. and Pouli, T. (2011). Colour spaces for colour transfer. *Lecture Notes in Computer Science*, 6626:1–15.

Sharma, G., Wu, W., and Dalal, E. (2005). The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application*, 30(1):21–30.

Szafir, D. (2018). Modeling color difference for visualization design. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):392–401.

Vanrell, M., Murray, N., Benavente, R., Parraga, C. A., Otazu, X., and Baldrich, R. (2011). Perception based representations for computational colour. *Lecture Notes in Computer Science*, 6626:16–30.

Vertan, C., Ciuc, M., and Stoica, A. (2003). Correlating human color similarity judgements and colorimetric representations. *Proceedings of SPIE*, 5227:51–58.