On the Effectiveness of Color Similarity Measures in Background–Frame Differencing Applications

Michael Hild
Osaka Electro-Communication University, Dept. of Engineering Informatics
Neyagawa, Osaka, Japan
E-mail: hild@hilab.osakac.ac.jp

Abstract

The effectiveness of color similarity measures in background-frame differencing applications is examined in this paper. Twelve published color similarity measures are investigated in RGB sensor space and CIE *L*a*b** color space and it is shown that the absolute-value exponential method is most effective. In contrast to other methods, this method uniquely relates brightness differences, saturation differences, and hue differences in the two input colors to similarity values as a one-to-one and monotonous function and the full value ranges of brightness, saturation and hue differences are mapped to the full range of similarity values.

1. Introduction

Background-frame differencing is used in the fields of computer vision, pattern recognition and robotics as a first, low-cost processing step for extracting regions of moving objects from image sequences.[1,2,3] The camera must be kept stationary and a photo of the background is used as a reference image. The technique determines the color difference between pixels of the image frames of the sequence and corresponding pixels of the reference image, and if this difference exceeds a pre-set value at a pixel, it is considered to be a pixel of a moving object.

Background-frame differencing can be realized by computing the similarity between two colors: If color vectors are sufficiently similar, pixels are judged as belonging to the background. In this context, the question of which color similarity measure might be most effective for this task is central. This question has been investigated in Ref. 6 by testing twelve known color similarity measures on twelve real-world scenes (indoor as well as outdoor scenes under various lighting conditions and photographed with two different types of CCD cameras). As color spaces, both the RGB sensor spaces of (off-the-shelf) CCD cameras and their corresponding CIE *L*a*b** color spaces were used. The results of that study suggested that there exists one color similarity measure – the *absolute-value exponential method* – which outperforms all other measures and also is effective in absolute terms.

In this paper we examine why the absolute-value exponential method is better than competing color similarity measures. We show that among the twelve examined measures the absolute-value exponential method is the only measure which relates brightness differences, saturation differences, and hue differences in the two input color vectors in RGB–space to similarity values as a function that is one-to-one and monotonous. It is also shown that this measure maps the full value ranges of brightness, saturation and hue differences to the full range of similarity values, and that its sensitivity is adjustable within certain bounds. The measure's performance in *L*a*b** color space is shown to be less attractive.

2. Color Similarity Measures

The twelve color similarity measures all have been proposed in the literature.[4,5] All of them take two vectors as input and compute a real number in the range (0.0, 1.0), where value 1.0 indicates "identical" and "0.0" is synonymous with "not similar at all." The similarity measures used are listed below. They are represented as functions $S_k(\mathrm{x}_i, \mathrm{x}_j)$, $k = 1, 2, \ldots, 12$, where x_i, x_j are the two p-dimensional color vectors. The arguments $(\mathrm{x}_i, \mathrm{x}_j)$ are skipped in the formulae shown below.

Measure 1

$$S_1 = \frac{\mathbf{x}_i \mathbf{x}_j^t}{\left|\mathbf{x}_i\right|\left|\mathbf{x}_j\right|} = \cos(\theta) \tag{1}$$

where θ is the angle between vectors x_i and x_j.

Measure 2

$$S_2 = \left(\frac{\mathbf{x}_i \mathbf{x}_j^t}{\left|\mathbf{x}_i\right|\left|\mathbf{x}_j\right|} \right) \left(1 - \frac{\left|\left|\mathbf{x}_i\right| - \left|\mathbf{x}_j\right|\right|}{\max\left(\left|\mathbf{x}_i\right|, \left|\mathbf{x}_j\right|\right)} \right) \tag{2}$$

Measure 3

$$S_3 = \frac{|\mathbf{x}_i|\cos(\theta) + |\mathbf{x}_j|\cos(\theta)}{\left(|\mathbf{x}_i|^2 + |\mathbf{x}_j|^2 + 2|\mathbf{x}_i||\mathbf{x}_j|\cos(\theta)\right)^{\frac{1}{2}}} \quad (3)$$

Measure 4

$$S_4 = \frac{\cos(\theta)\left(|\mathbf{x}_i|^2 + |\mathbf{x}_j|^2 + 2|\mathbf{x}_i||\mathbf{x}_j|\cos(\theta)\right)^{\frac{1}{2}}}{|\mathbf{x}_i| + |\mathbf{x}_j|} \quad (4)$$

Measure 5

$$S_5 = 1 - \frac{\left(|\mathbf{x}_i|^2 + |\mathbf{x}_j|^2 - 2|\mathbf{x}_i||\mathbf{x}_j|\cos(\theta)\right)^{\frac{1}{2}}}{\left(|\mathbf{x}_i|^2 + |\mathbf{x}_j|^2 + 2|\mathbf{x}_i||\mathbf{x}_j|\cos(\theta)\right)^{\frac{1}{2}}} \quad (5)$$

Measure 6: Correlation Coefficient Method

$$S_6 = \frac{\sum_{k=1}^{p} \left|\mathbf{x}_{ik} - \bar{\mathbf{x}}_i\right|\left|\mathbf{x}_{jk} - \bar{\mathbf{x}}_j\right|}{\left(\sum_{k=1}^{p} \left(\mathbf{x}_{ik} - \bar{\mathbf{x}}_i\right)^2\right)^{\frac{1}{2}} \left(\sum_{k=1}^{p} \left(\mathbf{x}_{jk} - \bar{\mathbf{x}}_j\right)^2\right)^{\frac{1}{2}}} \quad (6)$$

where $\bar{\mathbf{x}}_i = \frac{1}{p}\sum_{k=1}^{p} \mathbf{x}_{ik}$.

Measure 7: Exponential Similarity Method

$$S_7 = \frac{1}{p}\sum_{k=1}^{p} \exp\left(\frac{-3}{4} \cdot \frac{(\mathbf{x}_{ik} - \mathbf{x}_{jk})^2}{\beta_k^2}\right) \quad (7)$$

where $\beta_k^2 > 0$ is a parameter that is determined experimentally.

Measure 8: Absolute-Value Exponent Method

$$S_8 = \exp\left(-\beta \sum_{k=1}^{p} \left|\mathbf{x}_{ik} - \mathbf{x}_{jk}\right|\right) \quad (8)$$

where $\beta > 0$.

Measure 9: Absolute-Value Reciprocal Method

$$S_9 = 1 - \beta \sum_{k=1}^{p} \left|\mathbf{x}_{ik} - \mathbf{x}_{jk}\right| \quad (9)$$

β is selected through experiments.

Measure 10: Maximum-Minimum Method

$$S_{10} = \frac{\sum_{k=1}^{p} \min(\mathbf{x}_{ik}, \mathbf{x}_{jk})}{\sum_{k=1}^{p} \max(\mathbf{x}_{ik}, \mathbf{x}_{jk})} \quad (10)$$

Measure 11: Arithmetic-Mean Minimum Method

$$S_{11} = \frac{\sum_{k=1}^{p} \min(\mathbf{x}_{ik}, \mathbf{x}_{jk})}{\frac{1}{2}\sum_{k=1}^{p} (\mathbf{x}_{ik} + \mathbf{x}_{jk})} \quad (11)$$

Measure 12: Geometric-Mean Minimum Method

$$S_{12} = \frac{\sum_{k=1}^{p} \min(\mathbf{x}_{ik}, \mathbf{x}_{jk})}{\sum_{k=1}^{p} (\mathbf{x}_{ik}\mathbf{x}_{jk})^{\frac{1}{2}}} \quad (12)$$

In Ref. 6, the above color similarity measures were examined with respect to their effectiveness for the background frame differencing task, using both the RGB sensor space and the CIE *L*a*b** color space for representing color. Based on experiments conducted on twelve different scenes (indoor as well as outdoor scenes under various lighting conditions and photographed with two different types of CCD cameras) it was shown that the *Absolute-value exponent method* performed better than all other measures for both color spaces. The performance ranking found in that study is summarized in Table 1.

Table 1. Ranking of Color Similarity Measures

Rank	Measure (RGB)	Measure (*L*a*b**)
1	S8	S8
2	S7	S7
3	S2	S10
4	S10	S9
5	S12	S11
6	S5	S2
7	S11	S5
8	S9	S12
9	S6	S6
10	S4	S4
11	S3	S3
12	S1	S1

The gap in performance between higher–ranked color similarity measures and lower–ranked ones was very wide, with the best similarity measures being able to classify more than 80% of pixels correctly and the worst ones only achieving about 10%.

3. Test Method

For investigating why the examined color similarity measures have such a wide spread of effectiveness, criteria describing in more detail what constitutes an effective color similarity measure are necessary. In the context of the background–frame differencing problem, such criteria can be proposed as follows:

1. The color similarity measures should be able to account for differences in the perceptual color features *hue*, *saturation* and *brightness* in a balanced way.
2. The functional relationship between a given color similarity measure and each of these perceptual color features should be *one–to–one*.

3. The functional relationship between a given color similarity measure and each of these perceptual color features should be *monotonous*.
4. A reasonable value range of each of the perceptual color features should map to the full value range of the color similarity measure.
5. It should be possible to adjust the *sensitivity* of a color similarity measure within reasonable bounds (i.e. control the rate of similarity value change per unit feature value difference).

In order to examine the color similarity measures, three sets of color vectors were prepared as test data: The first set was prepared to exhibit a wide variation of brightness, the second one was prepared to exhibit a wide variation of saturation, and the third set was optimized for wide variation of hue values.

For the first set, twenty color vectors $(R_k,G_k,B_k)^T = k \cdot (10,\ 10,\ 10)^T$, $k = 1, 2, ..., 20$, were computed. These colors are monochrome, and only brightness varies within the set. Brightness was computed as the mean value of the vector components. An equivalent set of color vectors was generated by transforming these RGB–vectors to *L*a*b**–vectors, where brightness is represented by the *L**–component.

For the second set, ten RGB–color vectors

$$(R_k,G_k,B_k)^T = 200 \cdot \left(\frac{10-k}{10},1,\frac{10-k}{10}\right)^T, k = 0,1,...,9$$

were generated and then scaled so that the brightness values stayed the same for all vectors. Saturation was computed as $S_{RGB} = 1 - 3 \cdot \min(r, g, b)$, where the *r, g, b* values correspond to *R,G,B* sensor values that are normalized by $(R + G + B)$. Consequently, this set of color vectors has constant brightness and hue (i.e. in this case green), and saturation varies between 1.0 and 0.035. A related set of color vectors was generated by transforming these RGB–vectors to *L*a*b**–vectors, and saturation was then computed as the chroma value

$$S_{L^*a^*b^*} = \left((a^*)^2 + (b^*)^2\right)^{\frac{1}{2}}$$

For the third set, RGB–vectors were chosen such that the entire hue range was evenly covered, while saturation remained constant. The vectors were then scaled for achieving the same brightness for all vectors. Hue H_{RGB} was computed using the formulae

$$H_{RGB} = \begin{cases} 2\pi - \theta, & B > G \\ \theta, & B \le G \end{cases},$$

$$\theta = \arccos\left(\frac{1}{2}\frac{(R-G)+(R-B)}{((R-G)^2+(R-B)(G-B))^{\frac{1}{2}}}\right),$$

where hue is represented as an angle in the range 0^o to 360^o. A related set of color vectors was generated by transforming these RGB–vectors to *L*a*b**–vectors, and hue $H_{L^*a^*b^*}$ was then computed as the angle $\theta = \arctan(b^*/a^*)$ such that $H_{L^*a^*b^*}$ has a value range of 0^o to 360^o.

The color similarity measures were tested by forming all possible combinations of two color vectors (x_i, x_j) contained in each of the data sets and using these color vector pairs as input for similarity computations. Their differences in brightness, saturation and hue were then computed and related to the color similarity of the vector pairs. Hue differences were computed such that they would not exceed 180^o, taking into account the periodicity of hue.

4. Results

Based on the test method outlined above, graphs of functional relationships "S_{xx}-Similarity versus Brightness difference", "S_{xx}-Similarity versus Saturation difference", and "S_{xx}-Similarity versus Hue difference" were computed for color vectors in RGB–space. Typical examples of these graphs are shown in Figure 2. Corresponding graphs for color vectors in *L*a*b**–space are shown in Figure 3. The characteristics of some typical results are outlined as follows:

Measure 2, Measure 5 (for RGB):

These measures can discriminate brightness differences, but the functional relationship between brightness differences and similarity values is not one-to-one (there exists a broad band of similarity values). The measures map the full saturation difference range to only about 50% of the similarity value range, and these functional relationships are not completely one-to-one, either. The full hue difference range is mapped to the full similarity value range, but the functional relationships are neither one-to-one nor monotonous, and small hue differences do not cause significant similarity value changes. The sensitivity of the measure cannot be adjusted.

Measure 6 (for RGB):

This measure cannot discriminate brightness differences nor saturation differences. The hue difference range is mapped to only about 30% of the similarity value range, and this functional relationship is strongly non-monotonous (i.e. from the middle of the hue difference range similarity values increase). This functional relationship also is slightly none-one-to-one, and the sensitivity of the measure cannot be adjusted. Another problem is that the similarity value is not defined when at least one of the input color vectors is monochrome (this causes division by zero).

Measure 7, Measure 8 (for RGB):

These measures can discriminate brightness differences and saturation differences, with the functional relationships for brightness differences being perfectly one-to-one and monotonous, and the functional relationships for saturation differences being almost one-to-one and monotonous for both measures. Both brightness and saturation value ranges are mapped to the entire similarity value range. The functional relationship between hue difference and similarity values is almost one-to-one and monotonous for Measure S_8, but for Measure S_7 this relationship is strongly non-monotonous. The sensitivity of both measures can be adjusted.

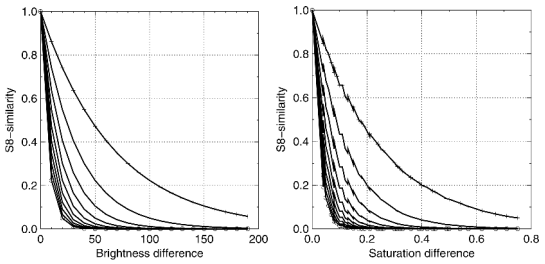

(a) Brightness (b) Saturation

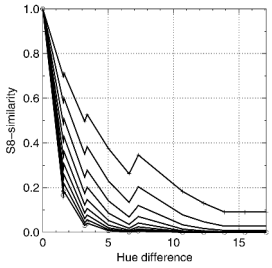

(c) Hue

Figure 1. S8 sensitivity depending on parameter β

When the color similarity measures are tested on color vectors in *L*a*b**–space, the properties of the color similarity measures are essentially the same as in RGB–space, except for the following:

1. The functional relationship between hue difference values and similarity values is not one-to-one for all similarity measures (including *S*8), and the similarity value spread is worse than it is in RGB–space.
2. The perfect characteristic of *S*7 in the case of brightness differences in RGB–space turns out to be much worse in the case of lightness in *L*a*b**–space.

Among the tested color similarity measures, measure S_8 gets closest to being ideal in the sense of satisfying the criteria proposed in Section 2. Next is measure S_7, followed by the groups of measures (S_2, S_5) and (S_9, S_{10}, S_{11}, S_{12}). The worst measures are S_1, S_3, S_4 and S_6. This ranking essentially coincides with the ranking obtained through experiments with real scenes (which are reported in Ref. 6 and are summarized in Table 1).

In order to examine the best color similarity measure *S*8 further, parameter β was varied in the range (0.005 ... 0.05). The resulting graphs (Fig. 1) show that the sensitivity of this measure can indeed be adjusted in a certain range, but the functional relationship between hue difference values and similarity values becomes more wiggly the smaller the value of β gets.

5. Conclusion

Based on the results presented above, it can be concluded that the absolute-value exponential method has more ideal characteristics than competing color similarity measures, because it is the only measure that has favorable properties with respect to the perceptual color features brightness, saturation and hue. In RGB–space, this measure relates brightness differences, saturation differences, and hue differences in the two input colors to similarity values as a function which is one-to-one and monotonous. It also maps the full value ranges of brightness, saturation and hue differences to the full range of similarity values, and its sensitivity is adjustable within certain bounds. In contrast, other measures do not have these characteristics. This explains why the absolute-value exponential method performs background-frame differencing most effectively in experiments involving indoor and outdoor scenes. The data also indicate that the measures' performance in *L*a*b** color space is worse than in RGB color space.

References

1. I. Haritaoglu, D. Harwood, and L. S. Davis, W4: Who? when? where? what? a real-time system for detecting and tracking people, Proc. The Third IEEE International Conference on Automatic Face and Gesture Recognition, pg. 222-227 (1998).
2. C. R. Wren, A. Azarbayejani, T. Darrell, and A. Pentland, IEEE Trans. on Pattern Analysis and Machine Intelligence, 19, 780-785 (1998).
3. M. Yamada, K. Ebihara, and Jun Ohya, A new robust real-time method for extracting human silhouettes from color images, Proc. The Third IEEE International Conference on Automatic Face and Gesture Recognition, pg. 528-533 (1998).
4. K. N. Plataniotis and A. N. Venetsanopoulos, Color Image Processing and Applications, Springer-Verlag, Berlin, 2000, pg. 70.
5. Y. Leung, Spatial Analysis and Planning under Imprecision, North-Holland Publishers, The Netherlands, 1988, pg. 116.
6. M. Hild and T. Emura, Which color similarity measure is most effective for background-frame differencing?, Proc. 9th Color Imaging Conference, pg. 168-173 (2001).

Biography

Michael Hild received his PhD from Osaka University, Japan in 1993. He has served as Lecturer at Osaka Electro-Communication University from 1994 to the present.

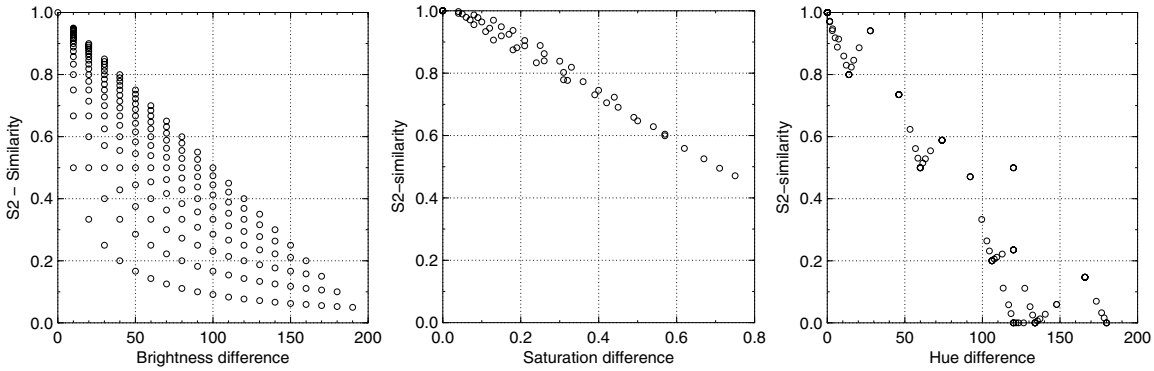

(a) S2 vs. Brightness (b) S2 vs. Saturation (c) S2 vs. Hue

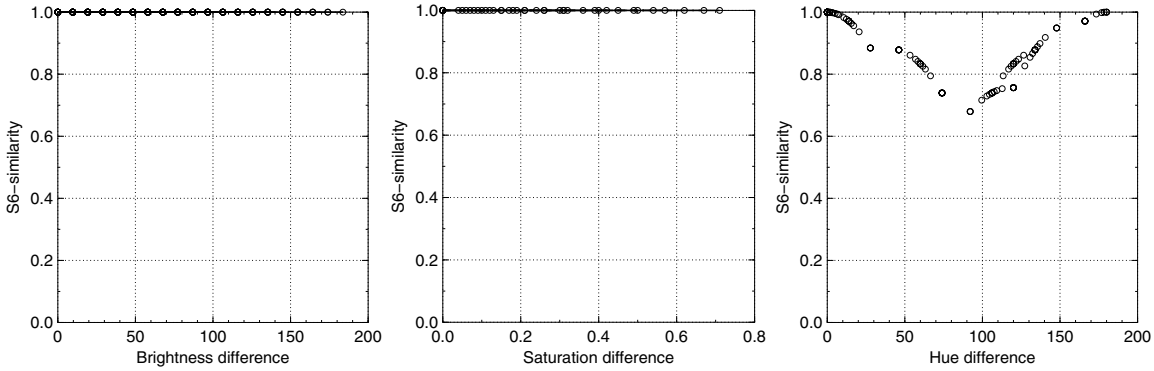

(d) S6 vs. Brightness (e) S6 vs. Saturation (f) S6 vs. Hue

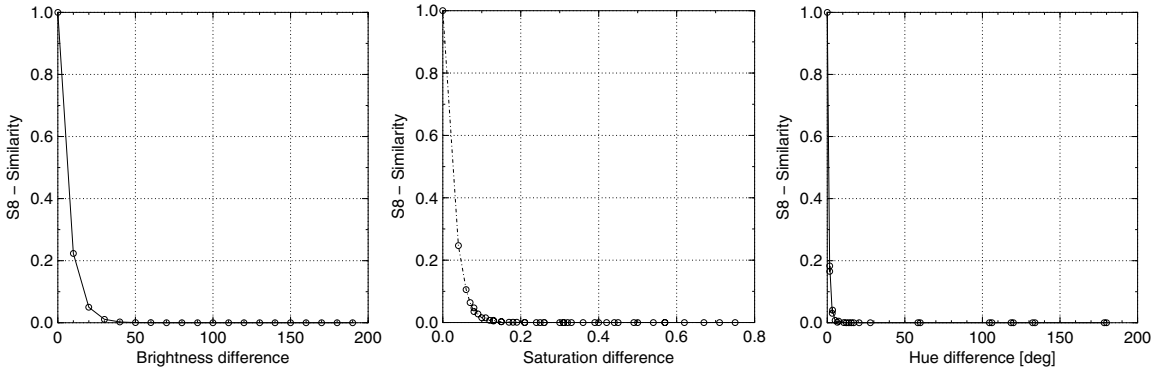

(g) S8 vs. Brightness (h) S8 vs. Saturation (i) S8 vs. Hue

Figure 2: *S2-similarity, S6-similarity, S8-similarity versus brightness, saturation and hue in RGB-space.*

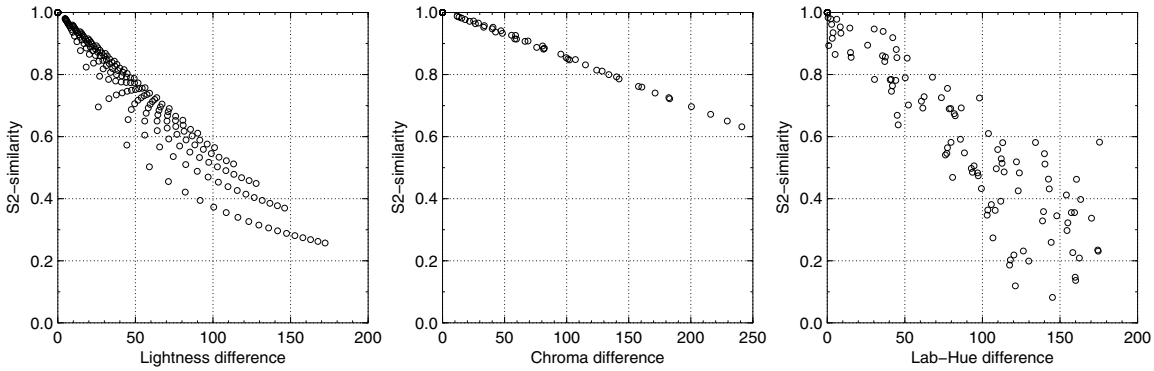

(a) S2 vs. Brightness (b) S2 vs. Saturation (c) S2 vs. Hue

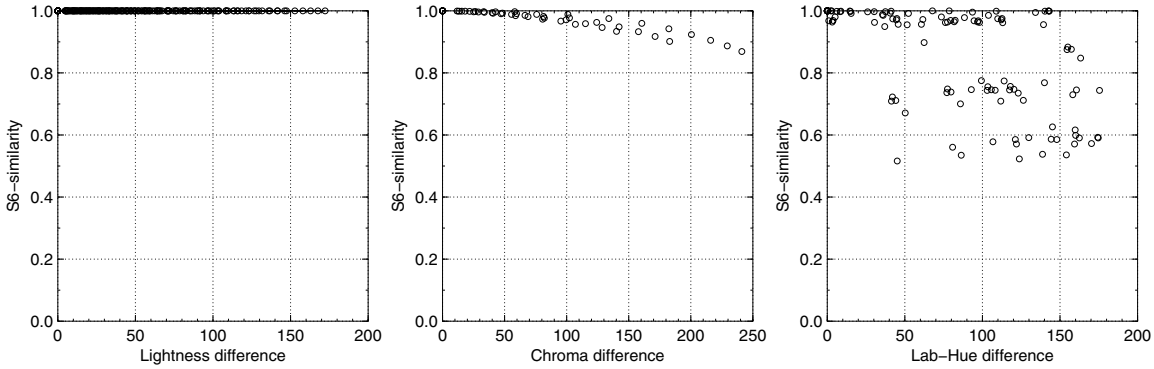

(d) S6 vs. Brightness (e) S6 vs. Saturation (f) S6 vs. Hue

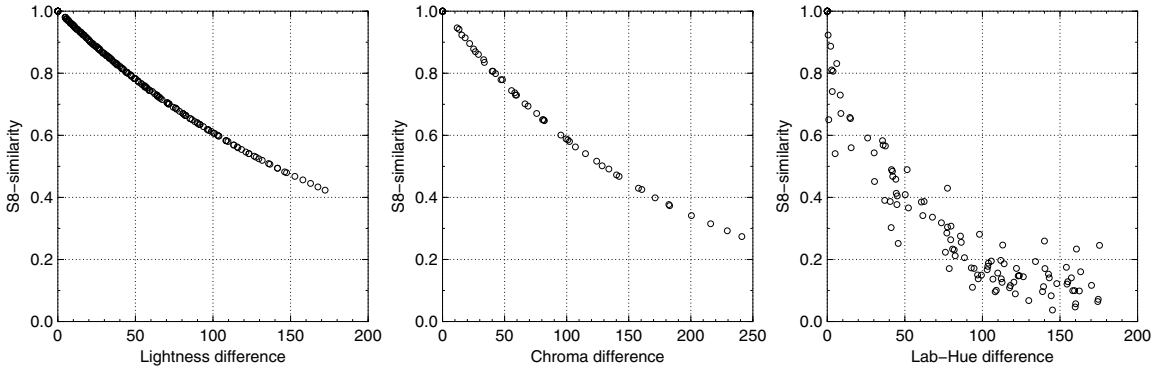

(g) S8 vs. Brightness (h) S8 vs. Saturation (i) S8 vs. Hue

Figure 3: *S2-similarity, S6-similarity, S8-similarity versus lightness, chroma and hue in $L^*a^*b^*$-space.*