Failure of Luminance-Redness Correlation for Illuminant Estimation

Florian Ciurea[*] and Brian Funt[**]
[*]Sony Electronics, Inc., San Jose, California, USA
[**]Simon Fraser University, Burnaby, British Columbia, Canada

Abstract

We investigate the hypothesis, recently published in *Nature,* that the human visual system may use some sort of luminance-redness correlation[2] together with the scene average for illuminant estimation. We found this idea interesting but not thoroughly tested. In particular, tests on real images were limited to scenes made up artificially from hyperspectral data,[4] spectral power distributions of various daylight illuminants, and the human cone sensitivity functions. The Ruderman database[4] of hyperspectral images is also quite peculiar because it consists of a small number of images of mostly foliage. Our experiments show that for scenes composed from a more diversified hyperspectral database combined with real illuminant spectra, the predicted correlation turns out to be very weak. For actual digital camera images, the luminance-redness correlation fails completely.

Introduction

Macleod and Golz recently proposed[1,2] that second-order statistics of image colors arising from the interplay between surfaces and illuminant in a scene could yield useful information about the illuminant color in real images. The simple mean color statistic used in the gray world assumption can not distinguish between a reddish room in white illumination versus a white room under reddish illumination. The luminance-redness correlation computed across all pixels in the scene could potentially be a distinguishing factor between the two scenes. Because of the distribution of the natural illuminants on the red-blue axis (often referred to as “warm” for reddish versus “cold” for bluish illuminants), a redness component of the illumination would account for most of its variation in chromaticity for natural light sources. In a series of experiments, we test the Macleod-Golz hypothesis as described in *Nature* postulating the existence of a luminance-redness correlation and find that the correlation, in fact, is very weak.

In probability theory, the correlation (or correlation coefficient) between two variables is defined as the ratio of their covariance by the product of their standard deviation. The correlation can vary from 1, in case of an increasing linear relationship, to -1 in case of a decreasing linear relationship. A zero correlation indicates that the two variables are independent.

In essence, the luminance-redness correlation is supposed to give insight as to the color of the illuminant, much the same as the gamut of observable colors under the current scene gives clues about the departure from the canonical gamut.[5] In fact, Mausfeld and Andres[3] propose that the mean and the covariance matrix, as first and second-order statistics respectively, may give valuable information about the shape and form of this gamut. Experiments by MacLeod and Golz[1,2] on images synthesized from a small hyperspectral database[4] of 12 images indicate that the luminance-redness correlation could be used, together with the mean of the sensor responses across the image, to disambiguate between reddish scenes under white light and white scenes under reddish light (Figure 1).

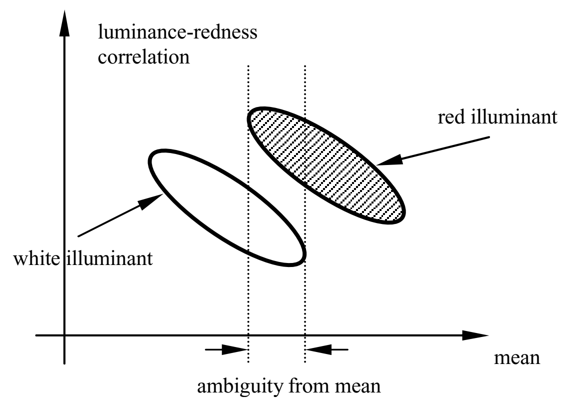

Figure 1. Luminance-redness correlation used in conjunction with the mean redness to determine the illuminant redness

In a two-dimensional space defined by the luminance-redness correlation and the mean image redness, the probability distribution of the white illuminant scenes can actually be distinct from the probability distribution of the red illuminant scenes. On the mean redness axis alone, there is a region where the two probability distributions overlap.

MacLeod and Golz use a logarithmic cone excitation space and define luminance as the sum of responses in the long and medium cones. The redness is defined as the luminance-normalized response in the long cone. Our experiments use the same definitions.

Is this method robust enough to be a good theory of human perception and does it work well to be used for automatic white-point estimation and color balancing in digital photography? We have devised several sets of experiments in which to measure the luminance-redness correlation on real images. Experiments 1-4 are done on images synthetically generated from real hyperspectral data and full spectral information of illuminants. For experiment 5 we use real images from the calibrated SFU dataset[6] and for experiment 6 we used the SFU large database.[7] In all the cases, we compute the luminance-redness correlation for each image of a scene under a certain illumination condition and then the correlation of the whole cluster of images under any given illumination.

Experiment 1

For this experiment, images were generated using the hyperspectral database collected by Ruderman et al.[4], and standard illuminants D40, D55, D85 and D200. This is a replication of the experiment on "real" data reported by MacLeod and Golz.[1,2]

Experiment 2

Here, images were constructed from the hyperspectral database[8] collected by Nascimento et al. and standard illuminants D40, D55, D85 and D200.

Experiment 3

For this experiment, we used the Ruderman hyperspectral database and the following four illuminants from the SFU calibrated database[6]: Philips Ultralume Tube, Sylvania Cool White Tube, Sylvania Warm White Tube and Solux 4700K with full Blue 3202 filter.

Experiment 4

In this case, the images were constructed from the hyperspectral database by Nascimento et al.[8] and the following four illuminants from the SFU calibrated database[6]: Philips Ultralume Tube, Sylvania Cool White Tube, Sylvania Warm White Tube and Solux 4700K with full Blue 3202 filter.

Experiment 5

We used real images of 8 scenes from the "Mondrian" set of the calibrated database of images from the SFU calibrated dataset[6] acquired with the Sony DXC-930. These indoor scenes were taken under controlled laboratory conditions in which the spectral power distribution of the illuminant is measured. The "Mondrian" set consists of images with minimal specular reflections. We used the same set of four illuminants as in the previous experiments using the SFU dataset.

Experiment 6

In this experiment we tested the luminance-redness correlation on the real images in the large database acquired with the Sony VX-2000 video camera.[7] These images represent mostly natural scenes, both indoor and outdoor. This database contains a description of illuminant chromaticity in each of the scenes, given in camera RGB coordinates. Unlike the first five experiments, here we do not have the same scene available under multiple illuminants. We simply make two large classes of illuminants labeled "red" and "neutral" based on the information about the illuminant chromaticity available in the database about each image.

Results on Real Images

The first experiment consists of reproducing the results on the Ruderman hyperspectral database under illuminants D40, D55, D85 and D200 using the human cone sensitivities.[9] We see a strong correlation between luminance-redness correlation and mean redness of the scene. The values for the correlations are respectively: -0.56, -0.65, -0.65, -0.55 and are in accordance to the values reported by MacLeod and Golz.[1,2] These values are significantly different from zero and so they indicate a strong correlation. Figure 2 shows the Ruderman database of hyperspectral images, and Figure 3 illustrates the results of experiment 1.

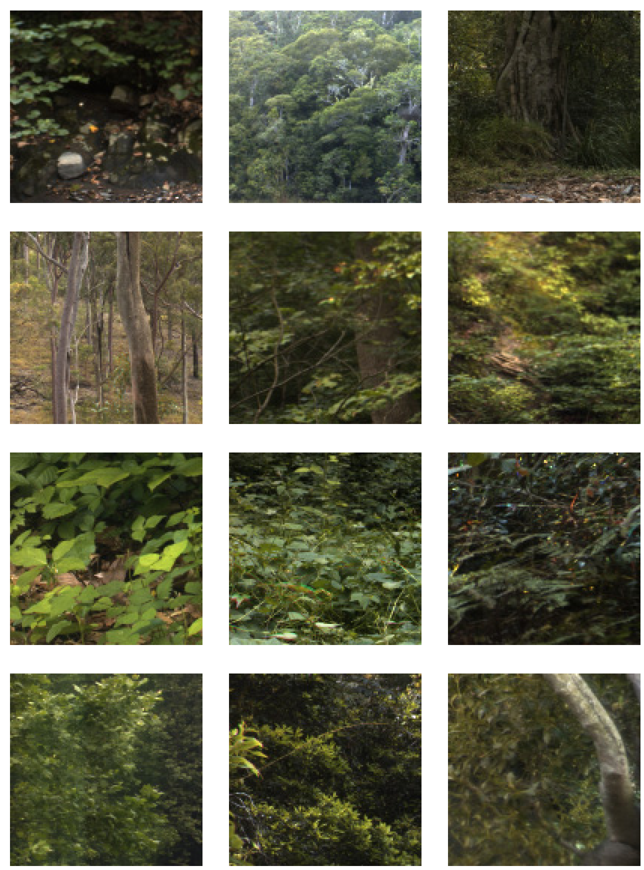

Figure 2. Hyperspectral database of 12 images by Ruderman et al

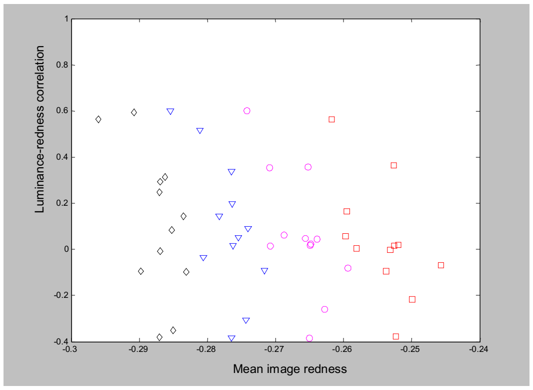

Figure 3. Experiment 1 in luminance-redness correlation. The 12 images are represented by black diamonds (illuminant D40), blue triangles (D55), pink circles (D85) and red squares (D200).

We believe that the dataset chosen for the original experiments by MacLeod and Golz was somewhat limited. It mainly contains images of scenes with foliage. In the second experiment, we used the 8 images from the hyperspectral database by Nascimento et al.,[8] which is more balanced.

Figure 4. Hyperspectral database of 8 images by Nascimento et al.

In this case, we found lower correlation scores for each cluster of images. The values are -0.40, -0.25, -0.09 and 0 respectively for the illuminants D40, D55, D85 and D200. The results are illustrated in Figure 5.

Experiments 3 and 4 are similar to experiments 1 and 2 respectively, with the exception that four real illuminants have been used instead of the ideal daylight sources proposed by Macleod and Golz. For experiment 3, the results are illustrated in Figure 6. The values for the correlations are -0.62, -0.47, -0.31, -0.20. While the pattern of correlation is still visible in experiment 3, this pattern is hardly visible in experiment 4 (see Figure 7). In this case, the correlation values are much lower: 0.28, 0.18, -0.01, 0.12 and in this case the luminance-redness correlation is virtually non-existent.

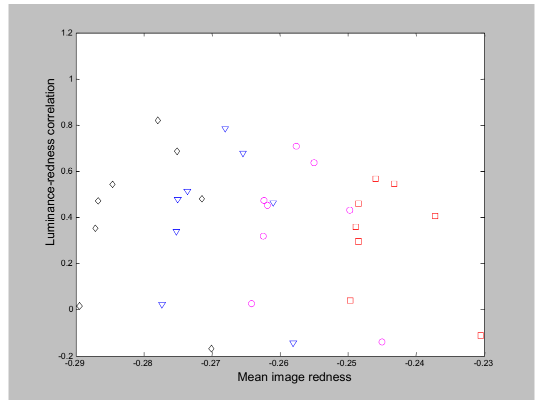

Figure 5. Experiment 2 in luminance-redness correlation. The 8 images are represented by black diamonds (illuminant D40), blue triangles (D55), pink circles (D85) and red squares (D200)

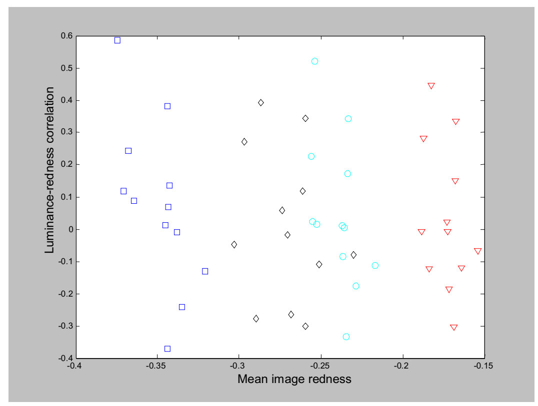

Figure 6. Experiment 3 in luminance-redness correlation. We used the same database of hyperspectral images as in Experiment 1, but with the set of four illuminants from the SFU database.

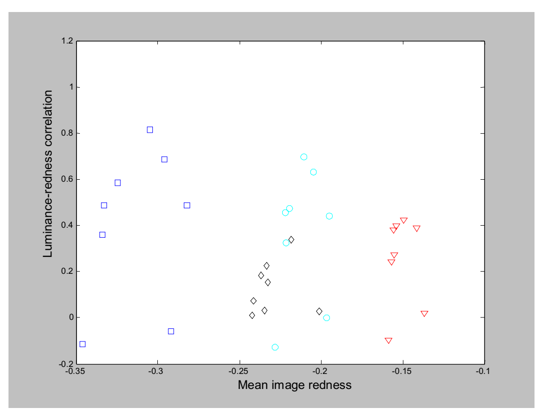

Figure 7. Experiment 4 in luminance-redness correlation. We used the hyperspectral database from Experiment 2, but with the SFU set of illuminants.

We used the hyperspectral database from Experiment 2, but with the SFU set of illuminants.

Experiments 5 and 6 are based on real images. Eight scenes from the SFU database[6] have been used in experiment 5, under each of the four chosen illuminants. In this case, the correlation values are even lower: 0.11, 0.19, -0.03, 0.21 (see Figure 8). For this experiment we also note that there is effectively no luminance-redness correlation. Experiment 6 uses the SFU large database of images.[7] The experiment differs from experiments 1-5 in that it is not based on having the same scene viewed under different illuminants. Here, we are simply interested in the probability distribution of scenes under “reddish” versus a “neutral” illuminant. A “reddish” illuminant is defined as having an r chromaticity component greater than 0.4. A “neutral” illuminant is defined as having an r chromaticity component less than 0.3.

It is worth mentioning that the major variability in the illuminant chromaticity is expressed by the red-blue component (r chromaticity). The probability distribution is shown in Figure 10. Here, we are looking for the pattern shown in Figure 11, which represents the ideal probability distribution in order for the luminance-redness correlation to be effective in illuminant estimation. Instead, the correlation values are –0.10 and 0.19 respectively. These results indicate that the luminance-redness correlation is again non-existent.

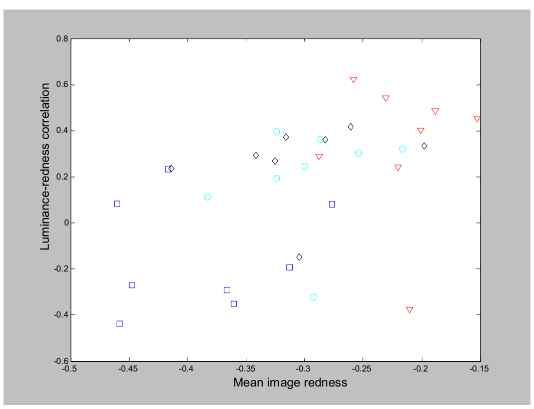

Figure 8. Experiment 5 in luminance-redness correlation. These are real images of the same scene under different illuminants.

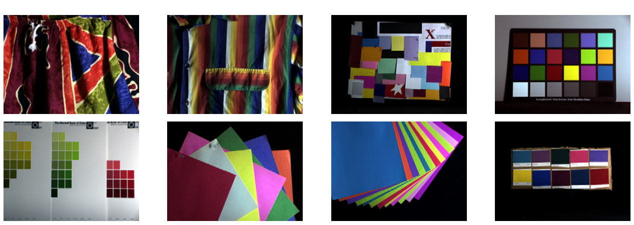

Figure 9. Experiment 5: the 8 scenes from the SFU database viewed under a canonical white illuminant.

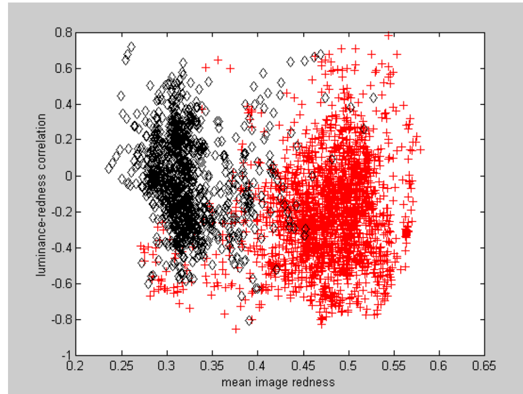

Figure 10. Experiment 6. The probability distribution of scenes under neutral (black crosses) and reddish illuminant (red stars)

Conclusion

We report a negative finding: Although it is possible to replicate the MacLeod-Golz experimental results showing luminance-redness correlation, their theory does not hold under more rigorous testing. The theory is attractive in that it relates the gray world method, which is based on the scene average alone, to the more advanced algorithms for illuminant estimation such as gamut mapping[5] and color by correlation.[10] However, the tests reported here show that the hypothesized luminance-redness correlation does not occur even on a second set of hyperspectral images. Furthermore, when testing with spectra of real light sources instead of ideal illuminant spectra, the correlation is even less pronounced. In the case of real digital camera images, the correlation does not exist at all.

References

1. D.I.A. MacLeod and J. Golz, *A computational analysis of colour constancy*, in *Colour Perception: Mind and the Physical World*, R. Mausfeld, D. Heyer, and H. Dieter, Editors. 2004, Oxford University Press.
2. J. Golz and D.I.A. MacLeod, *Influence of scene statistics on colour constancy.* Nature, 2002. **415**: p. 637-640.
3. R. Mausfeld and J. Andres, *Second-order statistics of colour codes modulate transformations that effectuate varying degrees of scene invariance and illumination invariance.* Perception, 2002. **31**: p. 209-224.
4. D.L. Ruderman, T.W. Cronin, and C.C. Chiao, *Statistics of cone responses to natural images: implications for visual coding.* Journal of the Optical Society of America, 1988. **15**: p. 2036-2045.
5. D.A. Forsyth, *A Novel Algorithm for Color Constancy.* International Journal of Computer Vision, 1990. **5**(1): p. 5-36.
6. K. Barnard, L. Martin, B. Funt, and A. Coath, *A Data Set for Color Research.* Color Research and Application, 2002. **27**(3): p. 148-152.

7. F. Ciurea and B. Funt. *A Large Database for Color Constancy Research*. in *Eleventh Color Imaging Conference*, 160-164, 2003. Scottsdale, AZ.
8. S.M. Nascimento, F. Ferreira, and D.H. Foster, *Statistics of spatial cone-excitation ratios in natural scenes.* Journal of the Optical Society of America A, 2002. **19**(1): p. 1484-1490.
9. A. Stockman, D.I.A. Macleod, and N.E. Johnson, *Spectral sensitivities of the human cones.* Journal of the Optical Society of America A, 1993. **10**: p. 2491-2521.
10. G.D. Finlayson, S. Hordley, and P.H. Hubel, *Color by correlation: a simple unifying framework for color constancy.* IEEE Transactions on Pattern Analysis and Machine Intelligence, 2001. **23**(11): p. 1209-1221.