

Color by Correlation

Graham D. Finlayson
Colour Research Group
University of Derby
Derby DE22 3BL
United Kingdom

Paul M. Hubel
HP Laboratories
Hewlett-Packard Company
Palo Alto, California 94304
USA

Steven Hordley
Colour Research Group
University of Derby
Derby DE22 3BL
United Kingdom

Abstract

Under a large variety of scene illuminants, a human observer sees the same range of colors; a white piece of paper remains resolutely white independent of the color of light under which it is seen. In contrast, color imaging systems (e.g. digital cameras) are less color constant in that they will often infer the color of the scene illuminant incorrectly. Unless the color constancy problem is solved, color appearance models cannot be used to guide image processing, and such processing is necessary for accurate (and acceptable) color reproduction.

In this paper we present a new theory of color constancy, *Color by Correlation,* which solves for the white-point in images by exploiting the correlation that exists between image colors and scene illuminants. For example, because the reddest red camera measurement can only occur under the reddest red light we say that the reddest camera measurement correlates strongly with the reddest light. Importantly all camera measurements correlate to a greater or lesser degree with different colors of light. By examining the correlation between all image colors and all lights we show that it is possible to make a very accurate estimate of the color of the scene illuminant.

Color by Correlation not only performs significantly better than other methods but is a simple, elegant solution to a problem that has eluded scientists working on color for over a century[1].

Introduction

In our work on testing color appearance models we have found that several of the models perform well when asked to compensate for a range of illuminants[2]. The main factor prohibiting the use of such models in digital photography (and probably most other applications) is the requirement that the color of the scene illumination must be known. In most situations we simply do not have this information.

In processing the digital camera image we must either measure the color of the scene illumination or estimate its color from the image data. Of course in working on digital imaging systems it is not practical to have an illumination sensor and expect users to calibrate to a white reference. If biological imaging systems achieve color constancy without an illumination color sensor, then it should be possible for us to achieve color constancy from just the image data (otherwise we would have evolved with spectrophotometers and white reference tiles mounted on our foreheads! - see figure 1).

The invention described here[3] is an improvement on an earlier technique[4] to determine the color of illumination in a scene.

Figure 1: What we would look like if our color vision had evolved using color appearance models.

Previous Methods

Many solutions have been proposed for the white-point estimation problem. Land[5], Buchsbaum[6], and Gershon[7], and others proposed that the average color of a scene is gray and so the white-point chromaticity corresponds to the average image chromaticity (we refer to this method as Gray World). Land[8] proposed that the maximum pixel responses, calculated in the red, green, and blue color channels individually, can also be used as a white-point estimate (we refer tho this method as Max.RGB). Maloney[9] & Wandell[10], Dzmura[11,12] & Iverson, Funt[13] & Drew and others have formulated the white-point estimation problem as an equation-solving exercise. In contrast, Tominaga[14]& Wandell, Funt[15], Tsukada[16], Drew[17] and others have shown that in principle the white-point might be found by exploiting the physics of the world, for example by finding specularity or interreflection regions in images. All these methods are similar, however, in the respect that they afford poor estimation performance[18,19,20]. These methods fail because they make assumptions about images which do not in general hold: the average of a scene is not always gray, and specularities may or may not appear in images (and when specularities do appear they are not easily found). Each of the methods is easily discredited.

That these methods fail has inspired other authors to search for color constancy algorithms which are based only on weak (that is reasonable) scene assumptions. Forsyth[21] developed a theory of estimation based soley on the observation that the range of colors measured by a camera (or the eye) depends on the color of the light (the reddest red color cannot occur under the bluest light). This idea was refined by Finlayson[4] (the *Color in Perspective* method) who observed that illuminant color is itself quite restricted.

Because the Color in Perspective method is the closest precursor to the correlation method presented here it is worth reviewing the details of how it works. In a preprocessing stage, Color in Perspective calculates models of plausible surface colors and plausible illuminant colors. These correspond to bounded regions of chromaticity space. A chromaticity image, of many surfaces viewed under a single scene illuminant, must be simultaneously consistent with both these constraint sets. That is, solving for color constancy amounts to a constraint satisfaction task; the output of Color in Perspective is the set of possible estimates of the white-point in an image. The mathematics of how Color in Perspective solves the constraint task is somewhat laborious (it involves calculating and intersecting many convex constraint sets). In addition, the method is highly sensitive to spurious inconsistencies. For example the presence of an aperture color in an image can force the solution set to be empty. The correlation method presented in this paper can be used to calculate the Color in Perspective constraint set. However, the new method is very much simpler (faster!) and is also more robust (is not sensitive to spurious outliers).

Adopting only the weak assumptions made in the Forsyth and Finlayson methods makes it impossible to return a unique estimate of the white point. Rather, a range of possible answers is returned, any one of which might be possible. Of course a single estimate must still be chosen from this set, and a variety of estimators have in fact been proposed for this task. Forsyth suggests that after white-balancing (discounting any color biases due to illumination), the image colors should be as colorful as possible. Finlayson and Hordley[22] propose the mean as a more robust estimate, and D'Zmura and Iverson[23] (and Brainard and Freeman[24]) suggest a maximum likelihood estimation. The latter estimator is particularly relevant to this work since our proposed solution can, as a special case, also support the maximum likelihood case. However, unlike the D'Zmura and Iverson method, our solution is computationally simple. Our method is so simple that maximum likelihood estimation could be provided at video frame rate.

The key observation that we exploit in our method is that the number of colors, and the range of white-points that a camera can sense, is finite. That is the white-point estimation is an intrinsically discrete problem. Funt[25] et al. recently proposed white-point estimation as a discrete neural computation problem. Here, image chromaticities are fed into a `trained' neural network which then returns a white-point estimate as output. Unfortunately, this method works as a `black-box' and so one cannot say too much about the estimates that are made, such as the estimate confidence. Moreover, physically impossible estimates can also be made.

Color by Correlation

The new method proposed here uses a "correlation matrix memory" or "associative matrix memory" which will not only achieve equivalent results to the Color in Perspective method, but will improve the method by adding Bayesian statistics to the process.

In this new method, a correlation matrix memory is built to correlate the data from any image (for example, a RGB image from a digital camera) to the set of possible scene illuminants. The vertical dimension of the matrix memory (the columns) is a rearrangement of the two-dimensional chromaticity space into a list of binary (1-0) points. For a particular color (formed under a particular illuminant), a point (chromaticity coordinate) is set to 1 if and only if that color can occur under that illuminant. For example, the reddest red chromaticity can only occur under the reddest illumination. The horizontal dimension corresponds to a list which corresponds to all plausible illuminants as seen by the device. To compute the data for the matrix, a set of reference surface colors are used (these could be a color chart or a set of standard surfaces). For each column (which corresponds to an illuminant), the chromaticities of the reference set are computed: the illuminant is multiplied by the reference surface reflectances, and the chromaticities as seen by the imaging device are calculated (plotted in figure 2a). The reference gamut for this illuminant is simply the polygon found when we take the convex hull of the points (figure 2b). Then in the column corresponding to that illuminant, the chromaticities of the points within this reference gamut are turned on (set to 1), and the others are turned off (set to 0). (In figure 2c the first column of the matrix corresponds to the shaded polygon plotted in figure 2b, which contains the chromaticities of the reference gamut under illuminant 1). This procedure is repeated for each column corresponding to all the illuminants. In practice, the number of illuminants can be limited to the precision of the desired results. (For example we may want to choose an illuminant from a group of 10 sources.)

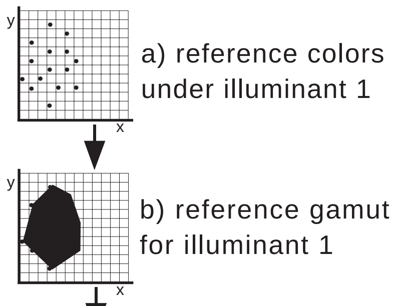

	ill1	ill2	ill3	ill4	ill5	ill6	ill7	ill8
x_1y_1	0	1	0	0	0	0	1	0
x_1y_2	0	1	0	0	0	1	0	0
x_1y_3	0	1	0	1	0	1	0	0
:	1	1	1	0	0	0	0	0
:	0	0	1	0	1	0	1	0
:	0	1	1	0	1	1	1	0
x_1y_n	0	0	0	1	0	1	1	0
x_2y_1	0	1	0	1	0	1	0	0
x_2y_2	1	1	0	1	0	0	0	1
:	1	1	1	1	0	1	0	1
:	1	1	1	1	0	1	1	0
:	1	0	1	1	1	1	1	1
:	0	0	1	0	1	0	1	0
:	0	1	1	0	1	0	1	0
:	1	1	0	0	0	1	1	0
x_ny_n	1	1	0	1	0	0	0	1

c)

Figure 2: Building a correlation matrix memory: a) first a set of reference surfaces illuminated by a particular source are plotted in chromaticity space, then b) by taking the convex hull of these chromaticities, we obtain the reference gamut for this illuminant, and finally c) a rearrangement of the chromaticities are listed in the matrix where 1 denotes the presence of the chromaticity in the reference gamut and 0 denotes its absence.

The above steps depend on the spectral sensitivities of the detector, the reference surfaces, and the illuminants, all of which are known for the reference image. This procedure is performed as part of the design or calibration of the imaging device. No assumption made of device linearity nor of spectral basis.

Estimating the white point

When the camera produces an image, the RGB data is converted to chromaticities and a vector is created corresponding to the values existing in the scene (left most binary vector in figure 3). This vector similar to the columns in the matrix memory, but contains 1's in the positions of chromaticities that appear in the image and 0's for chromaticities that do not appear in the image.

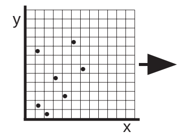

image vector		ill1	ill2	ill3	ill4	ill5	ill6	ill7	ill8
0	x_1y_1	0	1	0	0	0	0	1	0
0	x_1y_2	0	1	0	0	0	1	0	0
1	x_1y_3	0	1	0	1	0	1	0	0
1	:	1	1	1	0	0	0	0	0
0	:	0	0	1	0	1	0	1	0
1	:	0	1	1	0	1	1	1	0
0	x_1y_n	0	0	0	1	0	1	1	0
0	x_2y_1	0	1	0	1	0	1	0	0
1	x_2y_2	1	1	0	1	0	0	0	1
0	:	1	1	1	1	0	1	0	1
1	:	1	1	1	1	0	1	1	0
0	:	1	0	1	1	1	1	1	1
0	:	0	0	1	0	1	0	1	0
1	:	0	1	1	0	1	0	1	0
1	:	1	1	0	0	0	1	1	0
0	x_ny_n	1	1	0	1	0	0	0	1

Figure 3: Image data is then plotted in the same chromaticity space and a list of these chromaticities is listed in vector form.

Next, we multiply this vector with each column in the correlation matrix giving a new matrix (figure 4). In this new matrix every row that represents a chromaticity that did not exist in the image contains all zeros, and the rows representing chromaticities that were in the image have data values of either 0 or 1 (a 1 indicates that a particular image chromaticity is consistent with a particular illuminant, and a 0 indicates inconsistency). Another way of thinking of this is that the rows are turned off or allowed to be left on (as they were in the correlation matrix) depending on the existence of that color in the image.

	ill1	ill2	ill3	ill4	ill5	ill6	ill7	ill8
x_1y_1	0	0	0	0	0	0	0	0
x_1y_2	0	0	0	0	0	0	0	0
x_1y_3	0	1	0	1	0	1	0	0
:	1	1	1	0	0	0	0	0
:	0	0	0	0	0	0	0	0
:	0	1	1	0	1	1	1	0
x_1y_n	0	0	0	0	0	0	0	0
x_2y_1	0	0	0	0	0	0	0	0
x_2y_2	1	1	1	1	0	0	0	1
:	0	0	0	0	0	0	0	0
:	1	1	1	1	0	1	1	0
:	0	0	0	0	0	0	0	0
:	0	0	0	0	0	0	0	0
:	0	1	1	0	1	0	1	0
:	1	1	0	0	0	1	1	0
x_ny_n	0	0	0	0	0	0	0	0
sum	4	7	5	3	2	4	4	1

Figure 4: Multiplication of the image vector by the rows in the correlation matrix turns on the rows that exist in the image and turns off the rows corresponding to non-image chromaticities.

Each column is then summed, and the resulting values form a sum-vector that represents the number of image chromaticities which are consistent with a given illuminant.

In our example, the image vector has 7 components equal to 1 and the sum-vector is:

4 7 5 3 2 4 4 1

It is apparent that 4 of the input colors are consistent with the first illuminant, 7 with the second, 5 with the third, and so on. The 2nd illuminant is consistent with all 7 input colors and so is the correct answer.

Ideally, if there are 7 input colors we should threshold the sum-vector at a value 7, so applying this threshold results in the binary vector:

0 1 0 0 0 0 0 0.

Notice only the 2nd component is 1 and this indicates that all input colors are consistent with the 2nd illuminant. In practice, however, a threshold which is less than 7 might be chosen (since aperture colors will not in general be consistent with other image colors). Thus, in this case, we find the illuminant that is *most* consistent with the input colors.

The operation of the algorithm depicted in figures 3 and 4 can also be described as follows:

If M is the correlation matrix, v is the vector chromaticity representation of the image, and v^t is the transpose of v, then the sum-vector in figure 4 is simply:

v^tM

Adding Probability

Interestingly, when the matrix M has elements set to 1 then this can be shown to be a very specific incarnation of the Bayesian statistical model. Specifically, all colors under a particular illuminant are assumed to be equally likely, as are the illuminants themselves. This is also equivalent to the previous Color in Perspective method (the equivalence is true for the case where the threshold value is equal to the total number of image chromaticities).

Given experimental data, we can update the correlation memory to exploit the full power of Bayesian statistics. Remember that the element at row i and column j of the correlation matrix memory is set to 1 if the ith image chromaticity is consistent with the jth illuminant. Let us suppose that we know the probability that chromaticity i occurs under illuminant j: p(i|j) (this is easily determined experimentally or from a theoretically determined distribution - figure 5 is one such distribution). Then Bayes' rule allows us to calculate the probability of illuminant j given the fact that chromaticity i appears in an image: p(j|i). Assuming independence of the surface reflectances that could appear in a scene, the probability of illuminant j given image chromaticities i and a second image chromaticity, k, is proportional to: p(j|i)p(j|k). Denoting log probabilities as p' then p(j|i)p(j|k) becomes p'(j|i)+p'(j|k).

If we initialize the position i,j in the correlation matrix to the value p'(j|i), then the correlation matrix memory approach can be used to find the most probable estimate of white. In this framework the maximum value in sum-vector corresponds to the most likely illuminant.

Experiments

We evaluated our new algorithm by testing it on a large number of synthetic images generated by taking a random subset of surface reflectances from the Munsell set. Though we have also tested our algorithm on real digital camera images, (and observed equally good performance) these synthetic images have the advantage of allowing us to test our algorithm quickly on hundreds of different images and to compare our technique easily to other approaches, such as Color in Perspective and Gray World.

It should be emphasized here that these tests favor methods such as Gray World – a random sampling of the Munsell surfaces do average to gray if enough surfaces are considered. Therefore we expect the Gray World approach to converge to the correct answer. If we obtain better estimation of the illumination than the Gray World approach in this framework, we would expect considerably greater superiority if we were to test more realistic situations where we know the Gray World approach would fail.

To form an image we require three components: surface reflectances, illumination, and sensors. For our experiments we randomly selected surface reflectances from a set of 462 Munsell[26] chips (the results below show how the algorithm performed when given images of between 5 and 25 surface reflectances). We selected an illuminant for each image from a set of common illuminants. These included Judd's[27] daylights together with a variety of tungsten and fluorescent lights. Finally, for our sensors we used three reasonably narrow band digital camera sensors.

Given a surface reflectance $S(\lambda)$, an illuminant spectral power distribution $E(\lambda)$, and a set of sensor spectral sensitivities $R_k(\lambda)$, a sensor response P_k is given by:

$$P_k = \int\{S(\lambda)E(\lambda)R_k(\lambda)\}\, d\lambda.$$

The set of sensor responses generated in this way form the synthetic image which is the input to our algorithm.

Before running our algorithm, we precompute the probability distributions described in the previous section, for each of our possible illuminants. For a given image, the algorithm calculates chromaticities, and uses these values together with the probability distributions, to generate a likelihood for each illuminant. Once we have computed a probability for each illuminant we can choose our estimate of the illuminant in a number of ways: we can choose the maximum likelihood, the mean likelihood, or a local area mean[24]. In the case of the Color in Perspective method - where all surfaces are deemed equally probable - we want to use a mean selection[22].

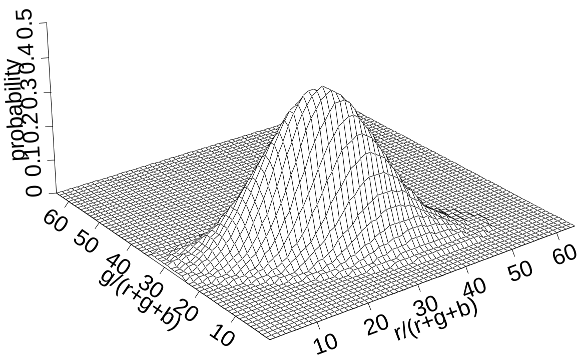

Figure 5: A probability distribution of the chromaticity i occurring under illuminant j: p(i|j).

Results

Figure 6 plots the CIELa*b* delta E errors between the estimated illuminant and the known illuminant (the correct answer), versus the number of surfaces in the test image. Each data point used to generate the curves is an mean of 100 estimates. The solid line summarizes the Gray World algorithm performance, dotted line for the Max.RGB (retinex type) method, short dashed line for Color in Perspective and long dashes for the new algorithm being proposed. (Similar denotations are used in figures 7 and 8.)

When we plot the median CIELa*b* error (figure 7) the Gray World and Max. RGB methods give similar results as before, but the Color in Perspective methods decreases, and the Color by Correlation method drops to zero after only ten surfaces - **which means that over half of the estimates calculated are perfect**. In figure 8 we have plotted the percentage of times an estimate is obtained with delta E less than 5 units. The Color by Correlation gives much better performance than the other techniques, even when only a few surfaces are contained in the image.

Clearly the correlation matrix technique performs considerably better than the other methods. Certain illuminants perform extremely well (tungsten), and some illuminants are more difficult to distinguish due to their similarity to neighboring illuminants (D65 and D75). In all cases, the correlation matrix gives a smaller error.

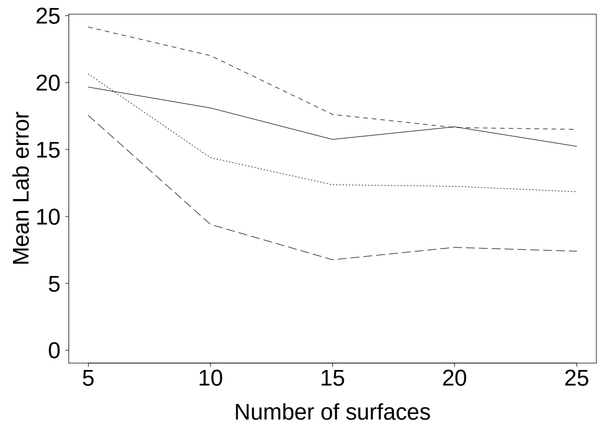

Figure 6: Plot of mean CIELab delta E difference between the estimated illumination color and the correct illumination color for synthetic images created using random collections of Munsell surfaces under 32 different illuminants. The estimation techniques include: Gray World estimation (solid line), Max. RGB estimation (dotted line), Color in Perspective correlation Matrix (short dashed line), and correlation matrix with probability (long dashed line).

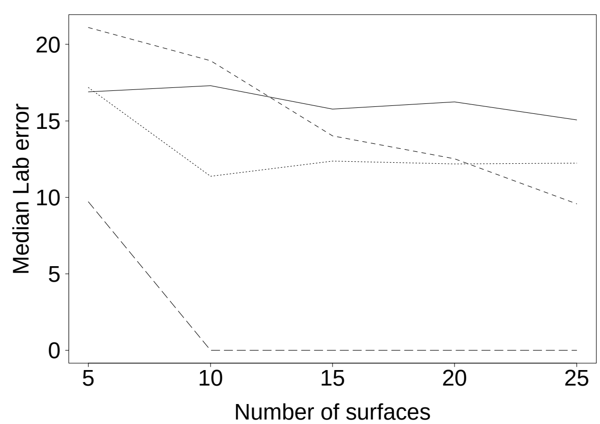

Figure 7: Plot of median CIELab delta E difference from the same data shown in figure 6. The Color by Correlation method (long dashed line) gave zero error (perfect illuminant estimation) in more than half of the estimates calculated even with onlu 10 surfaces in the scene.

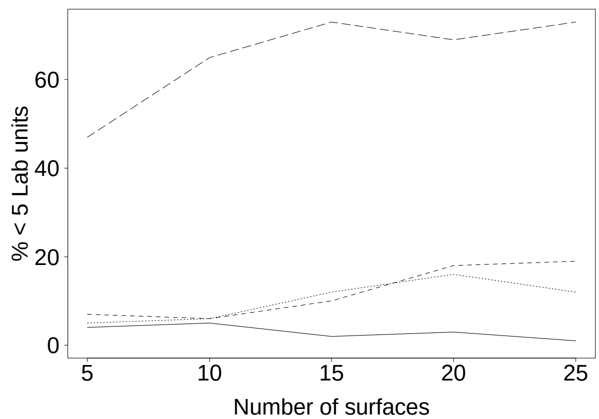

Figure 8. This set of curves shows the percentage of times the given methods give an illuminant estimate within 5 CIELab deltaE units of the correct answer as a function of the number of surfaces in the image.

Conclusions

Color by Correlation is a simple and elegant solution to the color constancy problem, that gives consistently better results than other methods. Since the method is based on constraints of possible illuminants in addition to probability statistics, it exploits the advantages of several previous techniques without the disadvantages (Gray World failure, for example).

References

1. H. von Helmholtz Handbuch der Physiologischen Optik (2nm ed.), Voss, Hamburg, (1896).
2. P.M. Hubel, and G.D. Finlayson, “Sharp Transforms for Color Appearance”, Accepted for publication and presentation at SPIE Electronic Imaging Conference: Device Independent Color, San Jose, CA, (1998).
3. G.D. Finlayson, G.D, and P.M. Hubel, “White point determination using correlation matrix memory”, U.S. Patent Pending, submitted October 1997.
4. G.D. Finlayson , "Color in Perspective”, IEEE transactions on Pattern Analysis and Machine Intelligence, p. 1034-1038, (1996).
5. E.H. Land, “Recent advances in retinex theory”, Vision Research, 26, p. 7--21, (1986).
6. G. Buchsbaum, "A spatial processor model for object color perception”, Journal of the Franklin Institute 310, p. 1-26 (1980).
7. R. Gershon and A.D. Jepson and J.K. Tsotsos, "From [R,G,B] to Surface Reflectance: Computing Color Constant Descriptors in Images”, Perception, p. 755-758 (1988).
8. E. H. Land, “The Retinex theory of Color Vision," Scientific American, p. 108-129, (1977).
9. L.T. Maloney and B.A. Wandell, “Color constancy: a method for recovering surface spectral reflectance”, J. Opt. Soc. Am. A, p. 29-33, (1986).
10. `B.A. Wandell, "The synthesis and analysis of color images”, IEEE Trans. Patt. Anal. and Mach. Intell. PAMI-9, p. 2-13, (1987).
11. M. D'Zmura, "Color Constancy: surface color from changing illumination”, J. Opt. Soc. Am. A (9), p. 490--493 (1992).
12. M. D'Zmura and G. Iverson, "Color constancy. I. basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces”, J. Opt. Soc. Am. A (10), p. 2148--2165 (1993).
13. B.V. Funt and M.S. Drew, “Color constancy computation in Near--Mondrian scenes using a finite dimensional linear model", IEEE Computer Vision and Pattern Recognition Proceedings, p. 544-549, (1988).
14. S. Tominaga and B.A. Wandell, "Standard surface-reflectance model and illuminant estimation”, J. Opt. Soc. Am. A, (6), p. 576-584, (1989).
15. B.V. Funt, M.S. Drew and J. Ho, “Color constancy from mutual reflection”, IJCV, (6), p. 5-24 (1991).
16. M. Tsukada and Y., "An Approach to Color Constancy Using Multiple Images”, IEEE Comp. Soc, (1990).
17. M.S. Drew, "Optimization approach to dichromatic Images”, Journal of Mathematical Imaging, (1993).
18. D.A. Brainard and B.A. Wandell, "Analysis of the Retinex theory of Color Vision”, J. Opt. Soc. Am. A (36), p. 1651-1661, (1986).
19. G.D. Finlayson, "Color constancy in diagonal chromaticity space”, IEEE Computer Society: Proccedings of the fifth International Conference on Computer, p. 218-223, (1995).
20. K. Barnard, "Computational color constancy: taking theory into practice”, MSc thesis, Simon Fraser University, School of Computing (1995).
21. D. Forsyth, “A novel algorithm for color constancy", IJVC (5), p. 5-36, (1990).
22. G. D. Finlayson and S. Hordley, “Selection for Gamut Mapping for Color Constancy”, Proceedings of the British Machine Vision Conference, to appear (1997).
23. M. D'Zmura, G. Iverson and B. Singer, "Probabilistic Color constancy”, in Geometric representations of perceptual phenomena: papers in honor of Tarow Indow's 70th birthday, ed. R.D. Luce et al, Lawrence Erlbaum Associates, (1994).
24. D.H. Brainard and W.T. Freeman, "Bayesian method for recovering surface and illuminant properties from photosensor responses”, in Human Vision, Visual Processing, and Digital Display V, p. 364 (1994).
25. B.V. Funt, V. Cardei and K. Barnard, "Learning Color Constancy”, 4th IS&T and SID Color Imaging (1996).
26. Newhall, Nickerson and D.B. Judd "Final Report of the {OSA} Subcommittee on the Spacing of the Munsell Colors”, J. Opt. Soc. Am. (33) p.385-418, (1943).
27. D.B. Judd, D.L. MacAdam and , and G. Wyszecki", Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature", J. Opt. Soc. Am."(54) p. 1031-1040 (1964)