Spectral-Based Illumination Estimation and Color Correction

Reiner Lenz,[1]* Peter Meer,[2] Markku Hauta-Kasari[3†]

[1] Image Processing Laboratory, Dept. Electrical Engineering, Linköping University, S-58183 Linköping, Sweden

[2] Dept. Electrical and Computer Engineering, Rutgers University, Piscataway, New Jersey 08855-0909

[3] Graduate School of Science and Engineering, Saitama University, 255 Shimo-okubo, Urawa, Saitama 338, Japan

Received 13 November 1997; revised 30 April 1998; accepted 8 September 1998

Abstract: We present a statistical technique to characterize the global color distribution in an image. The result can be used for color correction of a single image and for comparison of different images. It is assumed that the object colors are similar to those in a set of colors for which spectral reflectances are available (in our experiments we use spectral measurements of the Munsell and NCS color chips). The logarithm of the spectra can be approximated by finite linear combinations of a small number of basis vectors. We characterize the distributions of the expansion coefficients in an image by their modes (the most probable values). This description does not require the assumption of a special class of probability distributions and it is insensitive to outliers and other perturbations of the distributions. A change of illumination results in a global shift of the expansion coefficients and, thus, also their modes. The recovery of the illuminant is thus reduced to estimating these shift parameters. The calculated light distribution is only an estimate of the true spectral distribution of the illuminant. Direct inverse filtering for normalization may lead to undesirable results, since these processes are often ill-defined. Therefore, we apply regularization techniques in applications (such as automatic color correction) where visual appearance is important. We also demonstrate how to use this characterization of the global color distribution in an image as a tool in color-based search in image databases. © 1999 John Wiley & Sons, Inc. Col Res Appl, 24, 98–111, 1999

Key words: color correction; color constancy; principal components; robust statistics

INTRODUCTION

A color image is always the result of a complex interaction between three different components: the optical properties of the scene, the illumination sources, and the sensors. Estimating the influence of these three factors on the measured signals is one of the main goals of color image analysis. The influence of the sensors is usually known, and the remaining problem is to separate the effects of the scene properties and the influence of the illumination. The human visual system can approximately solve this problem, a phenomenon known as color constancy (see Ref. 1 for an introduction and references). The problem is also actively pursued in computer vision and image processing, but a satisfactory solution for real world conditions is yet to be found.[2–18]

In this article, we describe the overall color distribution in an image (or an image-patch) as follows. First the color of a single point in an image is defined by the logarithm of its spectrum. This log-spectrum is then approximated by the first few terms in a series expansion. The coefficients are computed for all points under consideration and for each

[†]Current address: Department of Information Technology, Lappeenranta University of Technology, P.O. Box 20, Fin-53851 Lappeenranta, Finland

* Correspondence to: R. Lenz, Department of Science and Engineering, Campus Norrköping, Norrköping University, SE-60174 Norrköping, Sweden

Contract grant sponsor: National Science Foundation; Contract grant number: IRI-9530546

Contract grant sponsor: Swedish Institute International Research Fellowship

Contract grant sponsor: Swedish Research Council for Engineering Sciences

Contract grant sponsor: Wihuri Foundation

Contract grant sponsor: Finnish Cultural Foundation/South Karelia Foundation

TABLE I. The relative mean approximation error.

Order N	Error $\hat{E}_N$	Error log coordinates E_N
1	28.7639	34.5235
2	13.5956	15.9981
3	6.6864	7.9114
4	4.7054	6.3125
5	3.4572	5.2342
6	2.9545	4.4786
7	2.4831	3.6910
8	2.1312	3.3019
9	1.6335	3.0091
10	1.2080	2.3878
11	0.9910	2.2866
12	0.7635	1.6729
13	0.6565	1.6152
14	0.5690	1.5403
15	0.4471	0.9986

coefficient the statistical distribution is characterized by its mode, i.e., its most probable value. The properties and the performance of the description will be demonstrated with some examples, where we apply it to color image normalization and color texture characterization.

In our image formation model, we assume that the spectral distribution $M(\lambda, x)$, measured at location x in the image can be written as $M(\lambda, x) = R(\lambda, x) \cdot L(\lambda)$, where $R(\lambda, x)$ is the reflectance function at that location and $L(\lambda)$ the spectrum of the light-source. Using the logarithm, a linear relation, $m(\lambda, x) = r(\lambda, x) + l(\lambda)$, is obtained. Note that the lighting is assumed to be uniform across the investigated part of the image. Here we do not model more complicated interactions between the scene, the illumination, and the sensors such as body-reflection or fluorescence.[19,20] Expressing the functions m, r and l in the same coordinate system spanned by functions $b_k(\lambda)$, we obtain the following series expansions:

$$m(\lambda, x) = \sum_k \mu_k(x) b_k(\lambda),$$

$$r(\lambda, x) = \sum_k \rho_k(x) b_k(\lambda), \quad l(\lambda) = \sum_k \alpha_k b_k(\lambda). \quad (1)$$

Thus, for a given k we obtain

$$\mu_k(x) = \rho_k(x) + \alpha_k, \quad (2)$$

i.e., the effect of the illuminant on the k-th expansion coefficient of the log-reflectance function is a location independent, constant shift. The coordinate system used in this article is based on an eigenvector expansion and will be described later.

In our applications, we will not use Eq. (2) pointwise, i.e., for individual positions x, but we will use the fact that the probability distributions of the coefficients $\mu_k(x)$ and $\rho_k(x)$ are related by the shift coefficient α_k. The value of α_k can, thus, be estimated by comparing the two probability distributions or parameters derived from them. In this article, we choose to describe probability distributions by their modes, i.e., their most probable value. The use of the mode is motivated by the need for a robust location estimator. We want to extract the "center" of the distribution for each coefficient. Using maximum likelihood would implicitly assume that the underlying distribution is normal, which is probably not the case for real data. The mode is insensitive to both skew and more, importantly, long tails (outliers) of the distribution. Note that the median (another robust location estimator) will be biased for a skewed or long tailed distribution. Modeling such distributions as a mixture of Gaussians makes the analysis cumbersome. The mode, on the other hand, can be computed by a simple algorithm, essentially a window sliding over the data. It can be shown that for unimodal distributions the result remains correct even when almost half the data is distributed along one of the tails. The mode is a maximum *a posteriori* (MAP) estimator, since it corresponds to the maximum of the p.d.f. of the given data.

Among the possible applications of Eq. (2) we mention the following:

- Color constancy: If the modes of the distributions of the coefficients μ_k and ρ_k are known, then the value of α_k and, thus, an estimation of the light source spectrum $L(\lambda)$ can be computed. The modes of the ρ_k distributions might be known from previous experience or, in the case of a dynamically changing illumination, they could be known from an estimation computed from a previous image.
- Image normalization: A desired global color impression of a color-corrected image can be obtained through a definition of the modes of ρ_k.
- Color description: In image database searches, it is often useful to find images with a given overall color impression. In this case, the modes of the ρ_k distributions are used as origin and the values of α_k characterize the color distribution in the image. We will later introduce the

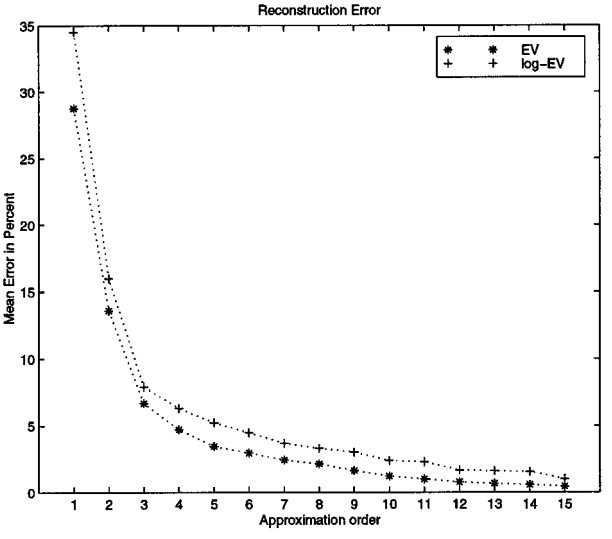

FIG. 1. Error function of the approximation order.

TABLE II. Modes from database and gray world assumption.

Eigenvector number	Mode database	Coefficient gray world
1	−9.1353	−9.7295
2	1.3931	1.0771
3	−0.5151	−0.2446
4	−0.0130	−0.1857
5	−0.1320	−0.0849
6	−0.1106	−0.0093
7	−0.2579	−0.0509
8	0.1048	0.0767
9	0.1099	0.0799
10	−0.2045	−0.1047
11	−0.0394	−0.0105
12	0.1370	0.0806
13	0.0134	0.0096
14	−0.0439	−0.0110
15	0.1156	0.0508

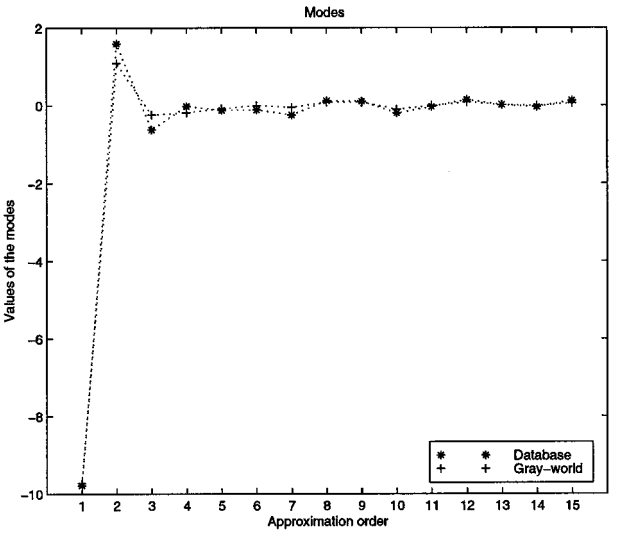

FIG. 2. Modes computed from the database and the "gray world" hypothesis.

concept of a relative illuminant and describe how to use it for color image retrieval.

The rest of the article is organized as follows: In Section 2, we describe how the basis functions $b_k(\lambda)$ are computed and we investigate the properties of the resulting coordinate system. In almost all applications, the spectral description of the color at a point in an image is not available. A conversion from the given color system to the spectral description is, therefore, necessary before the basic algorithm can be applied. In Section 3, we discuss this conversion for the case where the pixels are RGB-vectors. Section 4 describes the estimation of the illumination light spectrum, and Section 5 its application in the restoration or compensation of the color shift. Section 6 discusses some useful postprocessing techniques, and the last section illustrates the results of some experiments.

BASIC ALGORITHM

Since human color perception is based on three different receptor types in the retina, the traditional approach towards color is to use three-dimensional coordinate systems. Each of these systems was designed for a certain type of application, and a substantial part of color science deals with the study of their properties and the conversion procedures among them. Examples are color monitors (RGB), color printing (CMY), computer vision and computer graphics (IHS), and the colorimetry (CIE-systems) such as Lab, Luv, etc. For exact definitions and descriptions see Refs. 21 and 22. None of the systems is especially suited for investigation of color constancy algorithms, because a change in the spectral characteristics of the illumination source often leads to complicated transformations in the 3-D coordinate space.

Our approach is based on the assumption that the space of electro-magnetic spectra, which is relevant for human color vision, can be described by a set of the order of a few thousand representative colors.[23] The representative colors are chosen based on perceptual criteria, i.e., they incorporate the subjectiveness of human color perception. The most well-known color-appearance systems are the Munsell system[24] and the Natural Color System (NCS).[25]

High-resolution measurements of the spectra of color chips for these representative colors are now available.* For each of the 1269 chips of the Munsell System, their spectra was measured from 380−800 nm at 1-nm steps, while the 1513 samples from the NCS system were measured from 380−780 nm at 5-nm intervals. These measurements were

* The Munsell spectra are available from the Information Technology Dept., Lappeenranta University of Technology, Lappeenranta, Finland. The NCS spectra were obtained from the Scandinavian Color Institute in Stockholm courtesy of B. Kruse.

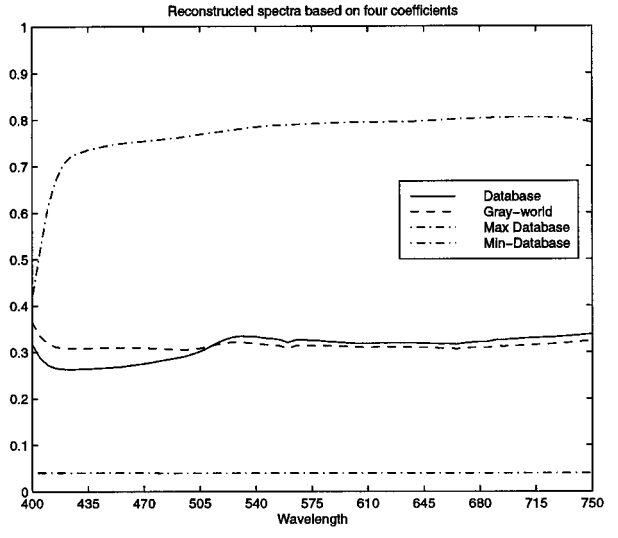

FIG. 3. Spectra computed from the database and the "gray world" hypothesis.

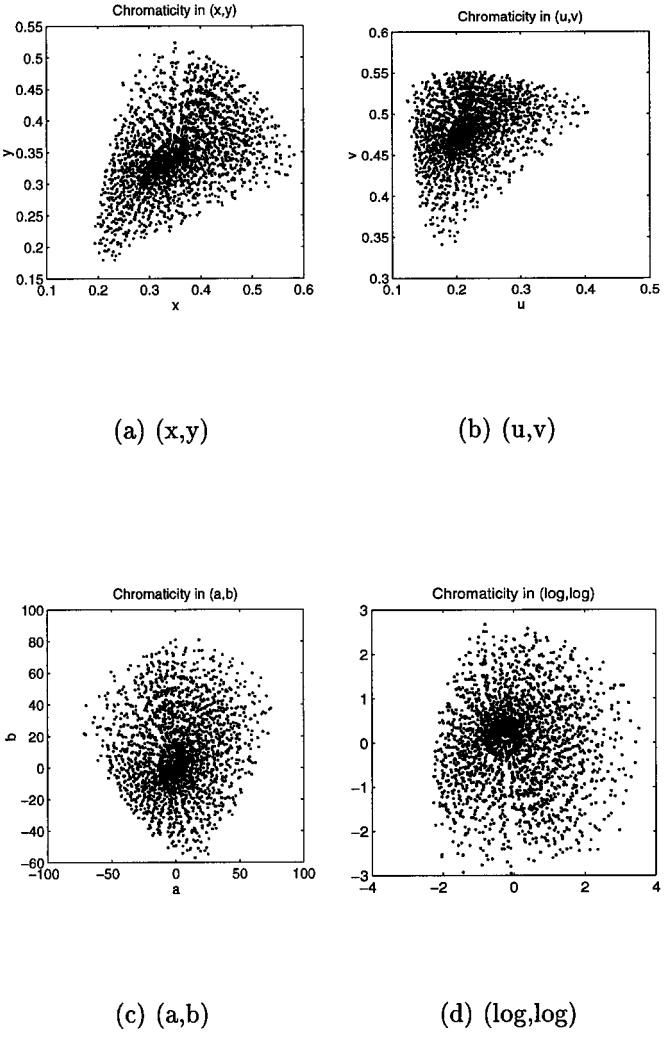

FIG. 4. Chromaticity coordinates of the database spectra.

combined in one set consisting of 2782 spectra (usually sampled in 5-nm steps from 400–750 nm). In the following we refer to this set of spectra as the spectral database.

These spectra are, of course, the spectra of the pigments used in the production of the color chips, they are not spectra of real world objects. The usage of these spectra is, however, necessary, because (to our knowledge) no equally comprehensible, representative set of natural spectra is available. Furthermore, we are only interested in some statistical properties of larger collections of colors, and the assumption that these properties are very similar for the color chips and the natural colors seems to be reasonable.

Most approaches to computational color constancy are based on the assumption that the reflectance spectra can be described by a low-dimensional model.[26–29] Usually, the coefficients in the eigenvector expansion of the spectra are used as variables. In this article, we do not approximate the spectra themselves by linear combinations of eigenvectors, but we use an expansion in logarithmic coordinates instead. We compute first the eigenvectors of the logarithmic spectra in the spectral database. Then we approximate the logarithm of a given spectrum by a linear combination of the first few eigenvectors. From this the original spectrum can be recovered by exponentiation. This expansion has always a higher minimum-mean-squared-error than the eigenvector expansion, but we found that the differences for approximation orders higher than two are small.

If $S(\lambda)$ is a positive function of the wavelength variable λ, then $s(\lambda) = \ln(S(\lambda))$ is the logarithm of this function. Lowercase letters always denote the logarithm of the capital symbols. By $\hat{b}_n(\lambda)$, we denote the eigenvectors computed from the original spectral database, and $b_n(\lambda)$ is the n-th eigenvector computed from the logarithm of the spectra in the same spectral database. The number of eigenvectors used in the approximation is N. The N-term approximation of the vector $S(\lambda)$ in the $\hat{b}_n(\lambda)$ system is given by $\hat{S}_N$, and the corresponding approximation in the $b_n(\lambda)$ system by S_N:

$$\hat{S}_N = \sum_{k=1}^{N} \alpha_k \hat{b}_k \text{ and } S_N = e^{\sum_{k=1}^{N} \beta_k b_k}. \tag{3}$$

The mean approximation errors are computed as

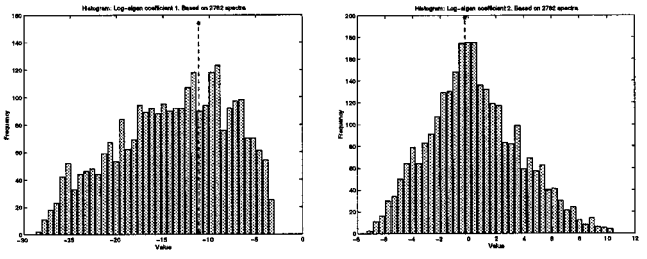

First Coefficient Second Coefficient

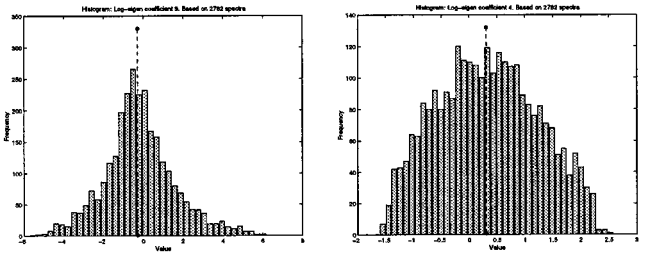

Third Coefficient Fourth Coefficient

FIG. 5. Histograms for log-database.

$$E_N = 100 \times \frac{\text{mean}\|S - S_N\|}{\text{mean}\|S\|}$$
$$\hat{E}_N = 100 \times \frac{\text{mean}\,\|S - \hat{S}_N\|}{\text{mean}\,\|S\|}, \tag{4}$$

where the mean is computed over all spectra in the database.

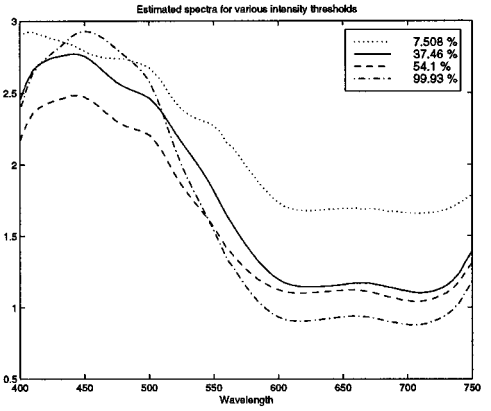

FIG. 6. Estimated spectrum as a function of the intensity threshold.

In Table I and Fig. 1, these errors are summarized. In our image formation model 1, we will always use the eigenvectors computed from the log-spectra.

In most applications, only a few measurements derived from the spectrum $M(\lambda, x)$ (such as the tristimulus values) will be available. In this case, the spectrum has to be estimated from the measurements. One way this can be done is described in the next section. Even when the whole spectrum is known, we see from Eq. (2) that only the distributions of the coefficients $\mu_k(x)$ are available, whereas the distributions of the reflection coefficients $\rho_k(x)$ and the constants α_k are unknown.

In such a case, we have to make some assumptions about the distributions of the coefficients r_k or α_k, or both. This is similar to the Bayesian framework used by Brainard and Freeman.[2] In our calculations, these assumptions will, however, enter only as the values of the modes of the distributions of the coefficients. We do not need to specify the complete distribution as in the Bayesian approach. Statisti-

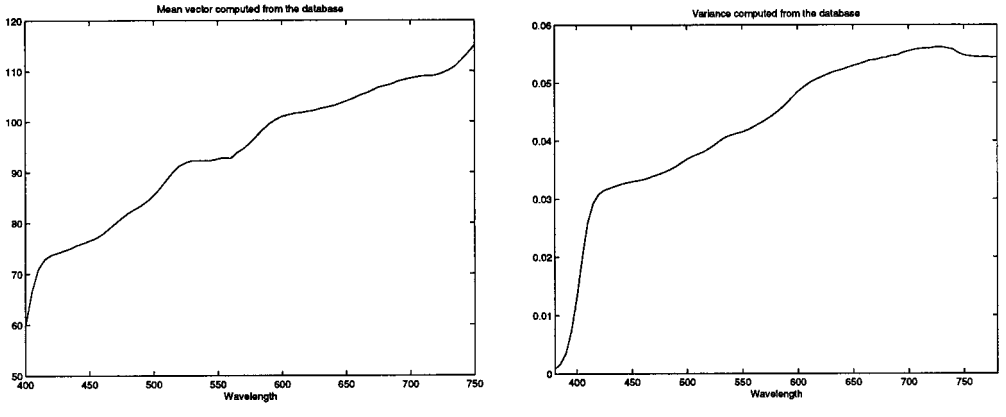

(a) (b)

FIG. 7. The mean- $A(\lambda)$ (a) and variance-function $B^2(\lambda)$ (b) computed from the database.

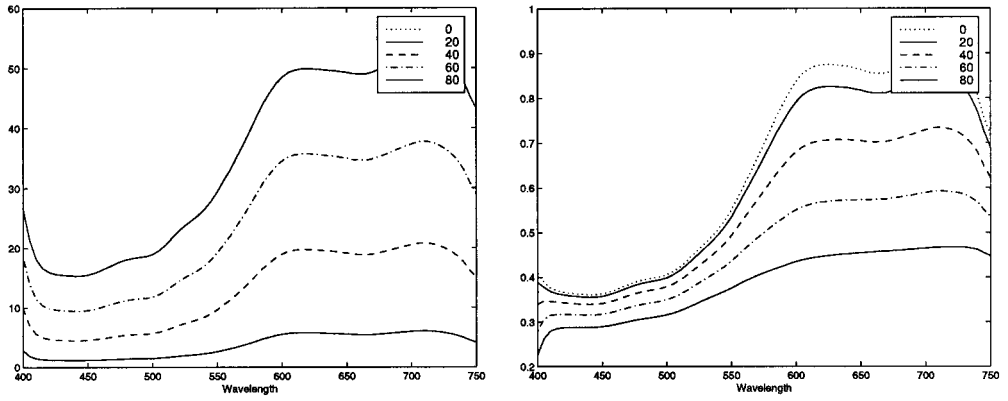

(a) (b)

FIG. 8. The multiplicative (a) and additive (b) restoration functions $a_\nu(\lambda)$ and $b_\nu(\lambda)$ for various regularization levels.

cal properties of the set of possible illumination spectra (like the daylight spectra[22]) can also be incorporated, but they will not be used in the following.

One hypothesis that we will often use is the assumption that the distribution of the reflection coefficients ρ_k has the same modes as the coefficients computed from the spectral database. This is certainly not the case in reality, since the spectra in the database appear with different probabilities in real scenes. Better estimates of the modes of the ρ_k distributions can be obtained by incorporating further knowledge about the image formation process. A simple way to compensate the different probabilities of the colors is by counting each color only once.

Another guess about the values of the modes of the distributions of the ρ_k coefficients can be obtained via the "gray world hypothesis" (see Ref. 2, p. 512). This assumption states that the mean over all reflectance functions is independent of λ:

$$\text{mean}_x\, r(x,\lambda) = \gamma. \tag{5}$$

In our experiments, we select λ in the range from 400–750 nm and choose the value of the constant γ such that the norm of the resulting constant vector has norm equal to the mean norm of the vectors in the spectral database. The logarithm of this function is approximated by the eigenvector expansion, and the coefficients in this expansion are listed in the second column of Table II. The first column of Table II contains the modes of the distributions of the expansion coefficients computed from the log-spectra in the database. Figure 2 shows these values in a diagram, and Fig. 3 shows the corresponding spectra computed from the modes of the first four coefficients (the upper and lower spectra are the spectra from the database, which have the highest and lowest norm, respectively). Apart from a small deviation in the shorter wavelength region, the two assumptions lead to nearly identical spectra.

From the modes of the distributions of the measurement coefficients μ_k and the modes computed from the spectral database (column one in Table II or the expansion coefficients in the second column of Table II), the values of the shift parameters α_k can be computed. From these α_k, a spectrum can be derived by exponentiation of the linear combination of the log-eigenspectra. This spectrum will be called a relative illuminant. In the case where the hypothesis

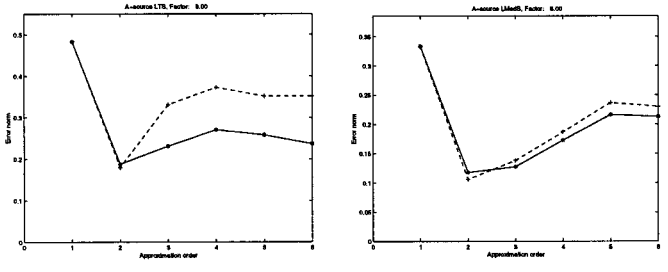

(a) LTS (b) LMedS

FIG. 9. Error distributions for two different images and A-source.

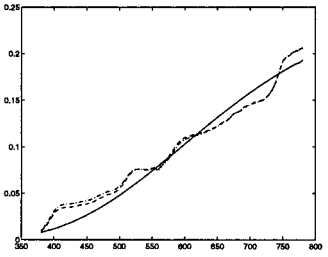

FIG. 10. A-source spectrum and its estimators.

about the statistical properties of the coefficients is true, the relative illuminant is an approximation of the true illuminant. For images in general, we can interpret the shift parameters α_k as a characterization of the global color distribution of the image. The parameter vector (α_k) has then the role of a coordinate vector, which can be used in color-image retrieval applications.

RGB TO SPECTRUM CONVERSION

True multispectral imaging is today only used in a few application areas such as remote sensing. The vast majority of color images is, however, stored in one of the three-dimensional color systems such as RGB, or CIE-related systems like Lab.[1,22,30]

A conversion from these descriptions to the spectral domain is needed before the algorithm can be applied. This is an ill-defined problem, since many different spectra will be mapped to the same coordinate vector in a three-dimensional coordinate system, an effect known as metamerism.

In the following, we assume that we are given a digital color image in RGB-format. Given the RGB vector at position x in an image, we have to estimate which spectral vector $m(\lambda, x)$ corresponds to this RGB-vector. The simplest estimation computes the RGB-coordinates for all elements in the spectral database and defines the spectrum $m(\lambda, x)$ as the nearest neighbor in RGB-space. This is unreliable due to the noneuclidean structure of the RGB-space. In the experiments described below, we first separate the intensity and the chromaticity properties of the RGB-vector and then we find the database spectrum with the best matching chromaticity values.

One motivation for this separation lies in the different ways that these values are obtained. The range of the RGB values (usually 0 . . . 255) is given by the hardware requirements of the digital image processing hardware, whereas the scaling of the spectral measurements in the database is determined by the physics of the measurement process.

Which spectrum is selected for a given RGB-vector depends on the coordinate system used in the chromaticity space. We experimented with the four different systems as follows: a given RGB-vector is first converted to *XYZ*-coordinates using the linear transformation specified by the CIE-1931 RGB- and *XYZ*-systems (Ref. 1, p. 139). For the spectra in the database, the *XYZ*-coordinates are computed using the CIE primary stimuli X, Y, and Z.[1] From the (X, Y, Z) vector, the chromaticity vector (x, y, z) is computed as

$$(x, y, z) = \frac{(X, Y, Z)}{(X + Y + Z)}. \tag{6}$$

In the following, denote the chromaticity vectors of the image RGB-vector and the database spectrum by $C_i = (X_i, Y_i, Z_i)$ and $C_s(X_s, Y_s, Z_s)$, respectively. The following methods to compute the distance between C_i and C_s were used:

(**x, y**): The distance between C_i and C_s is the l^1 norm of the (x, y)-part, i.e.:

$$\text{dist}_{xy}(C_i, C_s) = |x_i - x_s| + |y_i - y_s|. \tag{7}$$

(**u, v**): This metric is the l^1 norm in the coordinates: $u = 4 * X/(X + 15 * Y + 3 * Z)$ and $v = 9 * Y/(X + 15 * + 3 * Z)$ resulting in

$$\text{dist}_{uv}(C_i, C_s) = |u_i - u_s| + |v_i - v_s|. \tag{8}$$

(**a, b**): The chromaticity coordinates are defined as: $a = 500 * (X^{1/3} - Y^{1/3})$, $b = 200 * (Y^{1/3} - Z^{1/3})$ and

$$\text{dist}_{ab}(C_i, C_s) = |a_i - a_s| + |b_i - b_s|. \tag{9}$$

(This is essentially the La^*b^* system, where reference white has (X, Y, Z) coordinates (1, 1, 1) and large enough intensity values.)

(**log, log**): Since we would like to use a table-lookup based conversion, we want the converted chromaticity vectors of the database spectra to fill a rectangle as evenly as possible. Therefore, we introduced the following conversion of the (X, Y, Z) vectors: First we compute: $(\xi, \eta) = (\log(X) - \log(Y), \log(Y) - \log(Z))$, which is similar to the Lab conversion. Then we shift and scale them as $((\xi - E(\xi))/S(\xi), (\eta - E(\eta))/S(\eta))$ (where E and S denote the expectation and the standard deviation), and finally a 45° rotation is applied. The chromatic distance between C_i and C_s is now the l^1 norm of the difference in these new coordinates.

For each element in the spectral database, we computed its coordinates in each of the four coordinate systems described above. In Fig. 4, the distributions of the resulting position vectors are shown for the four coordinate systems.

ESTIMATING THE ILLUMINATION LIGHT SPECTRUM

Using the table lookup described in the previous section gives, for each selected RGB-vector in the image, the spectrum in the database with the most similar chromaticity properties. The norm of this spectrum vector is then multiplied with a normalizing factor, which compensates the intensity differences between the image RGB-vector and the database spectrum. Expanding the logarithm of this spec-

(a) (b)

FIG. 11. Postcard image: (a) indoor illumination; (b) blue light.

trum in eigenvectors of the set of log-database spectra results in the values $\mu_k(x)$ introduced in Eq. (1).

The distribution of the expansion coefficients can be very diverse. This can be seen in Fig. 5, which shows the distribution of the first four components ρ_k, $k = 1 \ldots 3$ computed from the spectral database.

To estimate the mode (the most probable value) for each coefficient, a robust mode estimator has to be employed. In a Bayesian framework, the mode is a MAP estimator, which minimizes the uniform error cost function [Ref. 31, p. 210]. In robust statistics several mode estimators were developed. We used the least trimmed squares (LTS) and the least median of squares (LMedS) estimators.[32] The same mode estimator is used to find the

(a) (b)

FIG. 12. Spectral normalization intensity threshold using (a) 37.46% vs. (b) 99.93% of the pixels.

(a) (b)

FIG. 13. Spectral normalization based on 99.93% of the pixels: (a) with and (b) without histogram equalization.

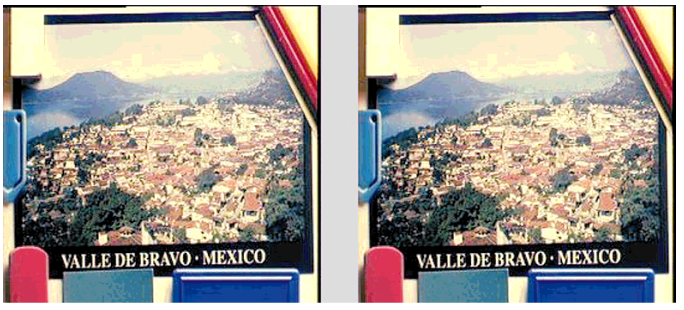

(a) (b)

FIG. 14. (a) Spectral normalization vs. (b) RGB normalization.

most probable value of the measurement coefficients $\mu_k(x)$ and the database coefficients $\rho_k(x)$. The value of the shift parameter α_k is given by the difference between the two modes.

In our experiments, we also found that, in the computation of the estimated relative illumination spectrum, one should take into account the intensity values at various positions. We include only the pixels with gray values above a given threshold in the estimation process. Figure 6 shows how the estimated spectrum depends on this threshold. (The numbers show the percentage of all pixels that were incorporated into the estimation.)

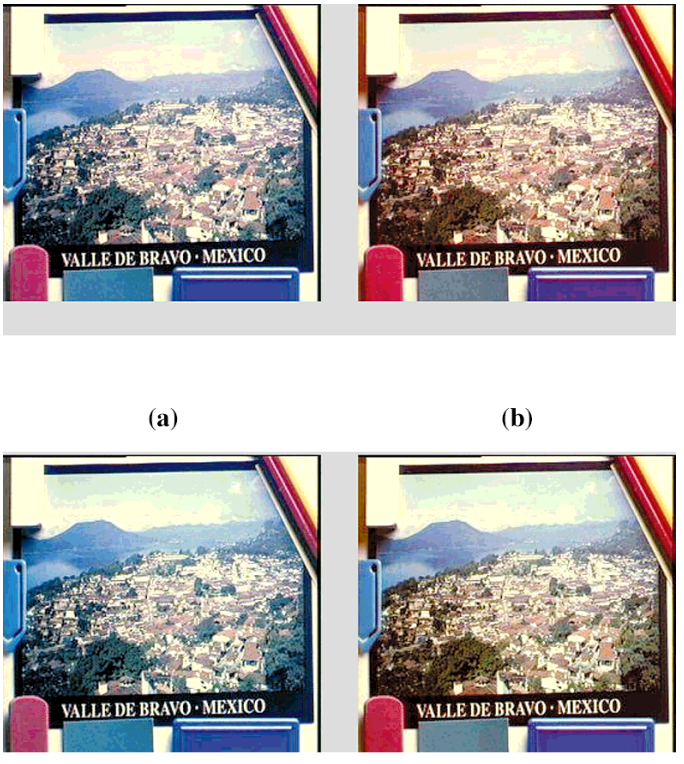

(a) (b)

(c) (d)

FIG. 15. Regularization: (a) threshold 2, regularization 0; (b) threshold 2, regularization 60; (c) threshold 5, regularization 0; (d) threshold 5, regularization 60.

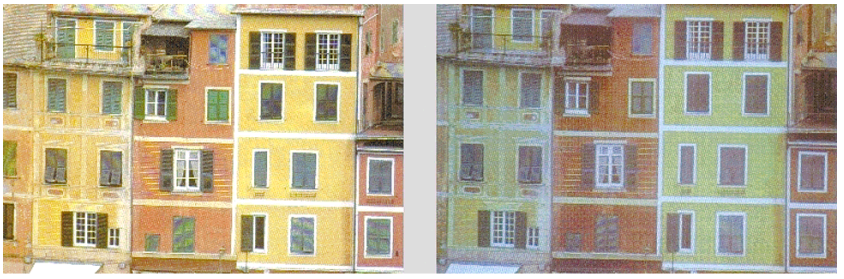

(a) **(b)**

FIG. 16. Image of the same scene captured with (a) normal lens and (b) telephoto lens.

CORRECTION

In the case where the image formation model $M = R \cdot L$ is correct and the estimated light source $\tilde{L}$ is equal to the true source L, the reflection function can be computed as $R = M/L$. In reality, the spectrum $\tilde{L}$ computed from the image is only an estimation of the true illumination characteristic. This is a typical ill-posed problem in which small estimation errors may lead to large errors in the final result. Regularization is a standard technique to avoid these effects. We applied it as follows.

In the original model: $M(\lambda, x) = R(\lambda, x) \cdot L(\lambda)$ the reflectance $R(\lambda, x)$ is regarded as a random variable with mean $A(\lambda)$ and variance $B^2(\lambda)$. The values of A and B^2 are estimated from the database spectra.

The differences between the real imaging process and its simplified model are collected in the random variable ϵ with mean zero and variance ν^2. This leads to $M(\lambda, x) = L(\lambda) \cdot R(\lambda, x) + \epsilon$.

The best linear estimator of the centered variable $R(\lambda, x) - A(\lambda)$ is [Ref. 33, Section 4.4.2]

$$\frac{L(\lambda)\,B^2(\lambda)(M(\lambda,x) - A(\lambda))}{L^2(\lambda)\,B^2(\lambda) + \nu^2}. \tag{10}$$

For the original variable R, this leads after some algebraic manipulations to the estimator

$$\begin{aligned} R(\lambda,x) &\approx \frac{L(\lambda)\,B^2(\lambda)\,M(\lambda,x)}{L^2(\lambda)\,B^2(\lambda) + \nu^2} \\ &\quad + \frac{([L^2(\lambda) - L(\lambda)]\,B^2(\lambda) + \nu^2)\,A(\lambda)}{L^2(\lambda)\,B^2(\lambda) + \nu^2} \\ &= M(\lambda,x) \cdot a_\nu(\lambda) + b_\nu(\lambda). \end{aligned} \tag{11}$$

This is the restoration formula used in the implementation. The values of A and B are estimated from the database, and the value of ν^2 is a free parameter that describes the confi-

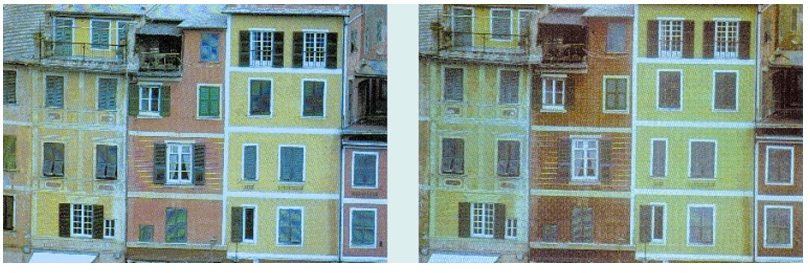

FIG. 17. Result of RGB normalization experiments.

FIG. 18. Result of spectral-based normalization experiments.

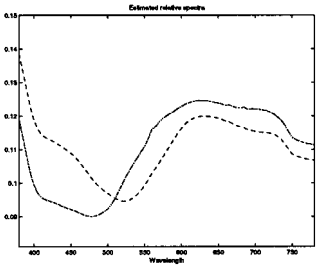

FIG. 19. Estimated illumination spectra from street images: (solid) normal lens; (dashed) telephoto lens.

dence in the estimate. (Note that ν^2 is global, while A and B depend on the wavelength.)

The distribution of the variance $B^2(\lambda)$ as computed from the spectral database is shown in Fig. 7. It shows that for short wavelengths λ the restored value of $R(\lambda, x)$ mainly depends on the average value $A(\lambda)$, whereas the measured value $M(\lambda)$ is dominating for the long wavelength region. Examples of the form of the functions $a_\nu(\lambda)$ and $b_\nu(\lambda)$ for various values of the regularization parameter ν are shown in Fig. 8.

In many correction problems, an input RGB-image has to be converted into an output RGB-image. For such applications it is often unnecessary to compute the nonlinear spectral-based correction pixel-by-pixel. A computationally more attractive approach is to replace the nonlinear estimation [like the one described in Eq. (11)] by the linear approximation

$$r_c = A \cdot r_o + b, \tag{12}$$

where r_c is the corrected output RGB-vector, r_o is the original RGB-vector, A is a 3×3 matrix, and b is a vector. For the zero vector b, this is one of the most often used methods for color correction (see Ref. 8 and Section 5.12 in Ref. 22).

One general method to find a suitable matrix A and a vector b for a given relative illuminant spectrum L is the following. In the first step, it is assumed that the estimated illuminant L is correct, and the spectra in the database are multiplied with the illuminant L to find a representative collection of reflected spectra. These spectra are then converted to RGB coordinates. The result is a matrix F containing the RGB-vectors of the spectra under the assumed illumination. In the second step, the desired inversion procedure (as in Ref. 8) is applied to these simulated illuminated object reflectance spectra given the corrected spectra. These corrected spectra are also converted to RGB resulting in a matrix G of RGB-vectors. Now the matrix A and the vector b are computed as solutions of the matrix equation $A \cdot F + b = G$. This equation can then be solved by familiar methods like the least-squares or total-least-squares. In our experiments, these two methods always produced comparable results, and the resulting RGB-images were visually more or less identical to the images obtained by the spectral-based correction images.

POSTPROCESSING

Using the spectral-based normalization method results in a restored spectrum $R(\lambda, x)$, which is an estimation of the reflectance properties of the object. It can, therefore, be used to compute the appearance of the object under any other illumination, simply by pointwise multiplication of the estimated spectrum and the spectrum of the light source.

For a conversion from the spectral domain to a three-dimensional coordinate system, like RGB, to be meaningful, the spectrum must lie in the gamut of the output device. This is not automatically the case, if we restore the spectrum with the procedure described above. In practice it is, therefore, often necessary to apply some postprocessing to the estimated spectra. In our current implementation, we first convert the spectra to RGB vectors and then we apply intensity-based postprocessing methods. Usually we truncate negative RGB values and we apply a gray-value based histogram-equalization to the corrected image. This is usually necessary, since the raw-restoration usually leads to a greatly reduced contrast in the image.

EXPERIMENTS

In the first series of experiments, we simulated the estimation of a known light source as follows. First, the RGB input image is converted to a multispectral image. Then pointwise multiplication with a known light source gives the simulated multispectral measurement image. The input to the algorithm is this RGB-converted multispectral image. From this RGB image, the illumination spectrum is estimated and compared with the true spectrum.

In the series of experiments we investigated the role of several factors:

1. 4 different input images
2. 1–6 eigenvectors to approximate the log-spectra
3. LTS and LMedS mode detectors
4. Four CIE light sources A, B, C, and D65
5. (x, y) and (log, log) chromaticity metric

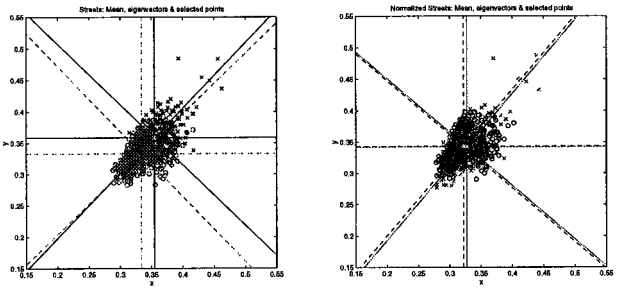

(a) (b)

FIG. 20. Chromaticity distributions of normal and telephoto images: (a) original images; (b) after normalization.

FIG. 21. Color-based search in image database.

The main conclusions from these experiments is that the LMedS and the LTS estimators nearly always gave identical results. The (log, log) metric was always slightly better than the (x, y) metric, only in one case [Fig. 9(a), solid line based on (log, log) and dashed line based on (x, y)] was this difference significantly higher. Using three eigenvectors to approximate the log-spectrum usually gave the lowest error. The error distribution for one of the images illuminated by the A-source (which has the strongest effect on the image colors) is shown in Fig. 9(b). The spectrum of the A-source and the estimated spectral distributions computed from the same image as in Fig. 9(b) is shown in Fig. 10. Here three eigenvectors and the LMedS and LTS-estimators were used.

The next experiment used more realistic conditions. A simple scene was captured under two different illuminations. In Fig. 11(a), normal indoor light conditions were used, whereas in Fig. 11(b) the light box 2412 of the Aristo Grid Lamp Products, Inc. provided a mostly blue illumination. In Figs. 12–15 some of the results with these images are summarized.

Figure 12 shows the result of the direct spectral normalization procedure [regularization parameter $\nu = 0$ in Eq. (11)] based on the spectra estimated from 37.46% and 99.93% of the input pixels (see Fig. 6 for the estimated spectra).

Figure 13 illustrates the effect of the gray-value based histogram equalization on the final image. Figure 14 compares the result of the normalization in the spectral domain with the normalization based on the RGB-to-RGB conversion method described in Eq. (12). Both normalization procedures are combined with the gray-value histogram-equalization. The resulting images are more or less identical.

The images in Fig. 15 finally show the effect of the regularization parameter ν in Eq. (11). In the upper row [images (a) and (b)], 37.46% of the pixels were used in

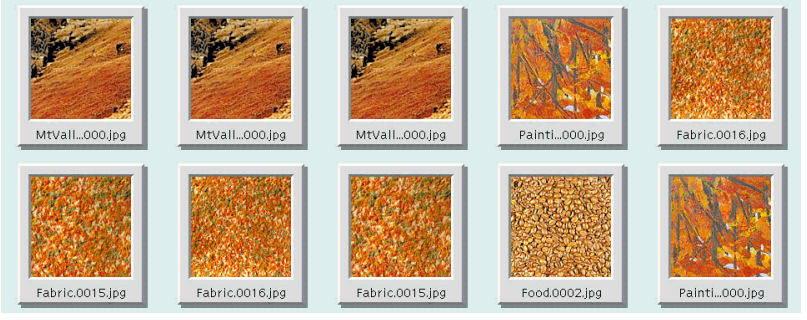

FIG. 22. Color-based search in image database.

FIG. 23. Color-based search in image database.

the estimation, whereas 66.39% entered in the computation of the images in the lower row [(c) and (d)]. In the computation of the images in the left column [(a) and (c)] no regularization is used, whereas for the right images [(b) and (d)] the regularization value 60 was used. (See Fig. 8 for the form of the multiplicative and additive restoration functions a_ν and b_ν.)

Beside the well-known dependence on the spectral composition of the illuminant (e.g., the recorded colors change significantly when the illumination changes from daylight to indoor incandescent light), an image can also appear significantly different when recorded in differently calibrated systems (photographic and/or electronic).

In Fig. 16, two color images of the same scene are shown, scanned in from Ref. 34, pp. 90–91. The image in Fig. 16(a) was taken from 50 m, while the image in Figure 16(b) was taken with a telephoto lens from 1000 m. In the latter, all the colors appear less saturated and the gray tones dominate. Note that the final result of the scanning process is not a photometrically correct description of the scanned image, since commercial scanners usually apply color transformations to produce visually more pleasing results.

We then investigated the normalization of the street images in Fig. 16. First, the illuminants relative to the complete Munsell/NCS system were computed. This information characterizes the chromatic properties of the images and can be used for indexing into a database. They are shown in Fig. 19. Note that the less saturated telephoto image leads to a flatter estimated spectrum.

For each of the two images, its estimated relative illuminant was then used to adjust the global color distribution of the image to the distribution of the Munsell/NCS system. Using both the spectral- and the RGB-normalization procedures we obtained two normalized images. The results of these experiments are shown in Fig. 17 for the RGB-based normalization and in Fig. 18 for the spectral based normalization.

The top row shows the original images as in Fig. 16. In the middle row are the results from the RGB-normalization, and in the bottom row the results for the spectral normalization.

A more quantitative description of the normalization effects can also be obtained. In Fig. 20, the chromaticity distributions (in CIE-xy space) of the two images (a) before and (b) after normalization are shown. The distributions are based on the chromaticity vectors of the same, 1000 randomly selected pixels. The points originating in the normal lens images (original and processed) are marked with a $+$, while the points from the telephoto images with a o. For all four images it was assumed that the images are viewed under D65 daylight and the chromaticity coordinates were computed accordingly.

The directions of the two eigenvectors of each chromaticity distribution were then computed, and are shown with solid lines for the normal lens images and dashed lines for the telephoto images. The centers of these eigenvector coordinate systems are located at the means of the chromaticity distributions. They are shown with the solid line cursor for the normal lens images, and with the dashed cursor for the telephoto images. The effect of the normalization is clearly revealed.

In the last experiment, we use the relative illuminant to locate images in a database. The database used consisted of the 473 color texture images in the VISTEX database.* For each of the images in the database we compute first the relative illuminant as described before. In the experiments below, we used 6 eigenvectors and we estimated the relative illuminant from 5000 randomly selected pixels in each image. Once these spectra are computed, they can be used to search in the database for images that have a certain global color distribution. In Figs. 21–23, we illustrate some of the results. In these experiments, we select first one image in the database as a prototype image and then we find those images in the database with the most similar color distribution. For this we normalize first all relative illuminant spectra to norm one to eliminate the influence of intensity variations, and then we use the scalar product of the two normalized relative illuminants as similarity measure for the global color distributions of these two images. Intensity properties of the images are ignored, because only normalized spectra enter the similarity computations. Figures 21–23 show the prototype image in the upper-left corner of the image. The other nine images are the most similar images found in the database (where the images are sorted from left to right, top to bottom with falling similarity values). Note that the database contains several images of the same texture at different resolutions and, therefore, the same texture may appear more than once.

Figures 21 and 22 show that the matching results are reasonable for homogeneous textures, whereas the results are less intuitive for images with different textures such as Fig. 23.

DISCUSSION AND CONCLUSIONS

We have shown that the log-eigenvector expansion of color spectra defines a coordinate system that allows an efficient solution to problems related to color constancy. Combining the log-spectral space with color-appearance systems (which are based on human color vision) allows a reliable estimation of the global color characteristics of a color image. The conversion procedures between this and other color systems however, need, further study.

The interpolation procedures used to convert RGB-vectors to spectra were sufficient in our application, because only statistical properties of a large number of spectra were needed. More difficult was the conversion from the spectral representation to quantized three-dimensional color representations like RGB. Here we found that a number of practical problems influenced the final results in a funda-

* http://www-white.media.mit.edu/vismod/imagery/VisionTexture/vistex.html.

mental way. Problems that had to be considered included: the handling of spectra that are outside the gamut of the chosen output device, and quantization methods that use the available number of colors efficiently and that at the same time preserve the structure in the spectral space.

ACKNOWLEDGMENTS

Markku Hauta–Kasari is on leave from the Department of Information Technology, Lappeenranta University of Technology, Lappeenranta, Finland.

The spectra in the Munsell database were measured at the Department Physics, Joensuu University, Finland and are available from the Department Information Technology, Lappeenranta University of Technology, Lappeenranta, Finland.

The NCS-database was obtained from the Scandinavian Color Institute in Stockholm and provided by B. Kruse.

1. Kaiser PK, Boynton RM. Human color vision. Washington, DC: OSA; 1996.
2. Brainard DH, Freeman WT. Bayesian color constancy. J Opt Soc Am A 1997;14:1393–1411.
3. D'Zmura M, Iverson G. Color constancy. i. basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. J Opt Soc Am A 1993;10:2148–2165.
4. D'Zmura M, Iverson G. Color constancy. ii. results for two-stage linear recovery of spectral descriptions for lights and surfaces. J Opt Soc Am A 1993;10:2166–2180.
5. D'Zmura M, Iverson G. Color constancy. iii: General linear recovery of spectral descriptions for lights and surfaces. J Opt Soc Am A 1994;11:2389–2400.
6. Finlayson GD. Color constancy in diagonal chromaticity space. In: Proc Int Conf Computer Vision. IEEE; 1995. p 218–223.
7. Finlayson GD, Drew MS, Funt BV. Color constancy: generalized diagonal transforms suffice. J Opt Soc Am A 1994;11:3011–3019.
8. Finlayson GD, Funt BV. Coefficient channels: derivation and relationship to other theoretical studies. Col Res Appl 1996;21:87–96.
9. Finlayson GD, Funt BV, Barnard K. Color constancy under varying illumination. In: Proc Int Conf Computer Vision. IEEE; 1995. p 720–725.
10. Forsyth DA. A novel algorithm for color constancy. Int J Computer Vision 1990;5:5–36.
11. Funt B, Cardei V, Barnard K. Neural network colour constancy and specularly reflecting surfaces. In: Proc AIC Colour 97. AIC; 1997.
12. Funt BV, Cardei V, Barnard K. Learning color constancy. In: IS&T fifth color imaging conference; 1996.
13. Funt BV, Finlaysson GD. Color constant color indexing. IEEE Trans Pattern Anal Mach 1995;17:522–529.
14. Finlayson GD. Color in perspective. IEEE Trans Pattern Anal Mach 1996;18:1034–1038.
15. Luong QT. Color in computer vision. In: Chen CH, et al., editors. Handbook of pattern recognition and computer vision. World Scientific; 1993. p 311–368.
16. Maloney LT, Wandell BA. Color constancy: a method for recovering surface spectral reflectance. J Opt Soc Am A 1986;3:29–33.
17. Sharma G, Trussell HJ. Digital color imaging. IEEE Trans Im Proc 1997;6:901–932.
18. Webster MA. Human color perception and its adaptation. Network: Comp Neur Sys 1996;7:587–634.
19. Klinker GJ, Shafer SA, Kanade T. The measurement of highlights in color images. Int J Computer Vision 1988;2:7–32.
20. Hecht E. Optics. Reading, MA: Addison–Wesley; 1987.
21. Ballard DB, Brown CM. Computer vision. Englewood Cliffs, NJ: Prentice Hall; 1982.
22. Wyszecki G, Stiles WS. Color science. 2nd Ed. London: Wiley & Sons; 1982.
23. Derefeldt G. Colour appearance systems. In: Gouras P, editor. The perception of colour. New York: CRC Press; 1991. p 218–261.
24. Munsell Book of Color, Matte Finish Collection. Baltimore: Munsell Color; 1976.
25. Hård A, Sivik L, Tonnquist G. NCS, natural color system—from concepts to research and applications. Col Res Appl 1996;21:180–205.
26. Brainard DH. Colorimetry. In: Bass M, editor. Handbook of optics. Vol. 1. New York: McGraw–Hill; 1995. p 26.1–26.54.
27. Jaaskelainen T, Parkkinen J, Toyooka S. Vector subspace model for color representation. J Opt Soc Am A 1990;7:725–730.
28. Maloney LT. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. J Opt Soc Am A 1986;3:1673–1683.
29. Parkkinen JPS, Hallikainen J, Jaaskelainen T. Characteristic spectra of Munsell colors. J Opt Soc Am A 1989;6:318–322.
30. The science of color. Washington DC: Opt Soc Am; 1973.
31. Mendel JM. Lessons in estimation theory for signal processing, communications, and control. Engelwood Cliffs, NJ: Prentice Hall; 1995.
32. Rousseeuw PJ, Leroy AM. Robust regression & outlier detection. New York: Wiley; 1987.
33. Louis AK. Inverse und schlecht gestellte Probleme. Stuttgart: Teubner Verlag; 1989.
34. De Grandis L. Theory and use of color. Englewood Cliffs, NJ: Prentice Hall; 1986.