Journal of Electronic Imaging 12(1), 161–172 (January 2003).

Retinex preprocessing of uncalibrated images for color-based image retrieval

Gianluigi Ciocca
Tecnologie Informatiche Multimediali, ITC
Consiglio Nazionale delle Ricerche, Milano, Italy
E-mail: ciocca@itim.mi.cnr.it

Daniele Marini
Università di Milano
Dip. Scienze dell'Informazione, Italy
E-mail: daniele.marini@unimi.it

Alessandro Rizzi
Università di Milano
Dip. Tecnologia dell'Informazione, Italy
E-mail: rizzi@dti.unimi.it

Raimondo Schettini
DISCO, Universita' degli Studi di Milano Bicocca, Italy
E-mail: schettini@disco.unimib.it

Silvia Zuffi
Tecnologie Informatiche Multimediali, ITC
Consiglio Nazionale delle Ricerche, Milano, Italy
Istituto di Biostrutture e Bioimmagini
Consiglio Nazionale delle Ricerche, Napoli, Italy
E-mail: zuffi@itim.mi.cnr.it

Abstract. *In image databases, variations in imaging conditions and preprocessing may result in similar originals that exhibit a low measure of similarity when color information is used in standard image retrieval methods. We examine the performance of various color-based retrieval strategies to see whether, and to what degree, the effectiveness of retrieval improves with Retinex-based preprocessing, regardless of the strategy adopted. The results of experiments performed on four different databases are reported and discussed.*
 [DOI: 10.1117/1.1526844]

1 Introduction

Recently the problem of meaningful retrieval from image databases has been addressed by many authors exploiting a great variety of approaches.[1,2] Most of the systems index only the syntactic content of the images in terms of visual features (color, shape, texture, size, distance, relative position, etc.). Visual features can be automatically extracted from raw data, without considering semantic information, and used to represent the entire image, a subregion of it, or a single object. Color-based features, in particular, have been extensively studied and experimented on in the last few years since they can be efficiently computed, and are particularly robust with regard to noise; image degradation; and variations in image size, resolution, and orientation. However, color-based features are unstable image indices when the acquisition conditions and imaging devices are not known *a priori*, or are not carefully controlled. Digital images depend on the physical content of the scene depicted, the illumination of the scene, and the characteristics of the imaging device. In many practical situations, we have very little or no information about these three factors. Other factors that may strongly affect color-based image retrieval, and are often underestimated: the images collected in a database, where they are usually represented only in terms of device-dependent color coordinates, may have been derived from many different sources, and may also have passed through a number of processing stages between entry and indexing; this may have introduced color shifts and nonlinear transformations, making similar, or even identical originals appear quite different and exhibit low similarity measures when standard image retrieval methods are used.[3]

We examine here the performance of various color-based retrieval strategies when coupled with retinex-based

Paper JEI 01017 received Mar. 27, 2001; revised manuscript received Apr. 4, 2002 and Jun. 24, 2002 accepted for publication June 28, 2002.
1017-9909/2003/$15.00

prefiltering to see whether, and to what degree, this preprocessing improves the effectiveness of retrieval, regardless of the strategy adopted. The experiments were performed on four databases, one of 310 paintings, another of 387 ceramic objects, the Simon Fraser University database of 55 images, and the University of East Anglia (UEA) uncalibrated color image database of 392 images.

2 Related Work

Querying a database with an image acquired under uncontrolled illumination, or with an imaging system different from that used for other archive items, may produce unsatisfactory results if the problems of the illuminant and the device dependence of the color were not addressed when designing the retrieval strategy. These problems were dealt with using mainly two different approaches. Some authors have introduced new illuminant-invariant color features for image indexing,[4–8] while others have followed a color constancy approach by preprocessing the query and database images to normalize the colors to a standard neutral illuminant. Image normalization can be accomplished in various ways. One method considers the color phenomenon primarily from a physical point of view, and the corresponding algorithms estimate the scene illuminant on the basis of the acquired image, e.g., using statistical techniques,[9] or constraints on the spectral matching between the image and some possible illuminants.[10] Another approach considers the color a result of the interaction between the physical stimulus and the adaptation mechanisms of human vision.[11,12] In this case, the equalization property of the algorithms derives from the spatial distribution of the chromatic contents of the image.

A completely different normalization strategy for image retrieval was recently presented by Finlayson *et al.*[13] In their approach, prefiltering is still used, but the objective is not to recover the appearance of image contents under some canonical conditions. Their normalization mechanism renders the image independent of lighting geometry and illumination color, but the final image produced is very different in appearance from the original.

Funt and Barnard[14] compared various preprocessing algorithms in the context of color-based object recognition using histogram intersection, and their conclusion is that prefiltering for color constancy improves results in object recognition tasks.

However, Finlayson and Schaefer[15] have recently reported that none of the existing color invariant and color normalization techniques, tested on a database of 28 designs acquired under 14 typical conditions, perform well enough to support color-based object recognition.

Our aim here is to investigate the performance of the effective image normalization algorithm retinex in a content-based image retrieval system when coupled with some of the retrieval strategies most commonly used in target-search tasks. We have focused on retinex as it's known to enhance images of unknown origin and unknowable color content.[16,17] Retinex maintains color appearance so that the necessary parameters for color image indexing preserve a meaning that would otherwise be lost. Moreover, retinex redistributes the dynamics of the image, expanding the dynamics of the color content, and maximizing the system's capacity for discrimination of the color based feature indexes.

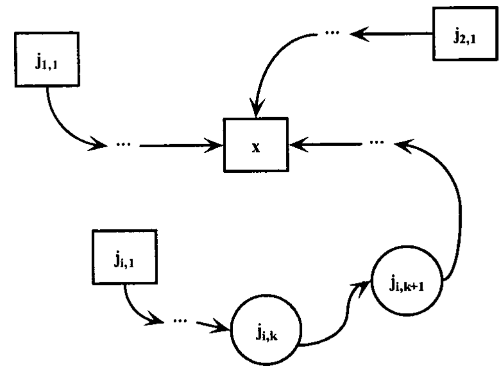

Fig. 1 Different paths to pixel I^x.

Sections 3 and 4 of this paper present retinex and the retrieval algorithms that have been implemented, while Sec. 4 describes the experiments performed, and provides a critical review of our results.

3 Retinex Filtering

Retinex was developed by Land and McCann[12] as a model of color perception in human vision. Many software and hardware variations of the original method have been presented in the last few years.[18–21] These algorithms differ mainly in the strategy used to exploit the spatial distribution of the colors in the scene, and require, in general, the heuristic setting of a large number of parameters (threshold, number of paths, maximum path distance, etc.). However, experimentation has shown that when these parameters are properly tuned, the various versions produce comparable results.[22] In our experiments we have used the latest version of McCann's algorithm and a retinex algorithm, designed by some of the authors of this paper, that employs Brownian paths and look up tables. The Brownian algorithm has been amply experimented in recent years,[23] allowing a fine adjustment of its parameters, while McCann's has recently been critically analyzed,[16] and a very effective implementation, which requires the setting of only one parameter, is now available.

3.1 *Basic Retinex Algorithm*

The original retinex algorithm,[12] applied to a digital image, can be described as follows. Let I be the image to be processed; for each pixel I^x at position x, a number N of paths starting from different points I^{j_i} with $i=1,...,N$ and ending at pixel I^x are generated (see Fig. 1). For each chromatic channel c, let us indicate the ith path as a sequence of n pixels $\{I_c^{j_{i1}}, I_c^{j_{i2}}, \ldots, I_c^{j_{i,n}}\}$ with $I_c^{j_{i,n}}=I_c^x$ and $I_c^{j_{i1}}=I_c^{j_i}$. Along each path, a sequence of ratio products called the chain value (CV) is computed according to the following rules:

$$\mathrm{CV}_c^{i,1}=1$$

for $k=2$ to n:

$$\mathrm{CV}_c^{i,k}=\mathrm{CV}_c^{i,k-1}\,\delta_c^{i,k} \quad \text{where}$$

$$\delta_c^{ik}=\begin{cases}\dfrac{I_c^{i,k}}{I_c^{j_{i,k-1}}} & \text{if } |I_c^{j_{i,k}}-I_c^{j_{i,k-1}}|>\text{threshold}\\ 1 & \text{otherwise}\end{cases}. \tag{1}$$

If all the pixels along the path are darker than the first one, the CV is less than 1. If a lighter pixel is found, the CV may exceed 1, and the following reset mechanism is then applied:

$$\text{If } \mathrm{CV}_c^{i,k}>1 \text{ then } \mathrm{CV}_c^{i,k}=1. \tag{2}$$

In this way the CV is recomputed starting from that lighter pixel.

The output pixel value $\overline{I_c^x}$ for channel c of the original pixel I^x is the mean of all the CVs computed along all the N paths that end at I^x:

$$\overline{I_c^x}=\frac{1}{N}\sum_{i=1}^{N}\mathrm{CV}_c^{i,x}, \tag{3}$$

These values are calculated independently on the three RGB channels considered an approximation of the *l*, *m*, and *s* retinal wavebands.[12]

3.2 *Brownian Look-Up Table Retinex—Retinex 1*

The various implementations of the retinex algorithm differ in the characteristics of paths used to explore the image. For our version we have chosen random Brownian paths generated with a midpoint displacement technique.[17,24,25]

The threshold value in Eq. (1), which enables us to cope with nonuniform illumination, may vary up to 10% in color depth without appreciable differences in the output images;[17] we have set this parameter at 5% in all the experiments reported here.

The implementation has been expedited with the use of look-up tables (LUTs) and a subsampled version of the original image.[22] The subsampling level depends on the size and the frequency content of the images. No automatic procedure for setting this parameter has been developed to date; we have set it, on the basis of our experience, at one eighth of the image sides. To avoid high-frequency aliasing we have first applied a 7×7 low-pass Gaussian filter. Three LUT mapping functions between the subsample and its filtered image have then been created, one for each color band. These LUT functions have been applied on the original full size image to obtain the filtered image (see Fig. 2).

The local effects of retinex may change the value of originally equal pixels. The three mapping functions are built by associating each value of the input subsampled image with the average of all the corresponding values in the output subsampled image. Lost values of the original image, not present in the subsample, are recovered by linear interpolation.

3.3 *McCann Multilevel Retinex—Retinex 2*

In the McCann multilevel retinex[16] there are no paths exploring the image: nonadjacent pixels are compared by applying the multilevel propagation mechanism of the pixel ratios.

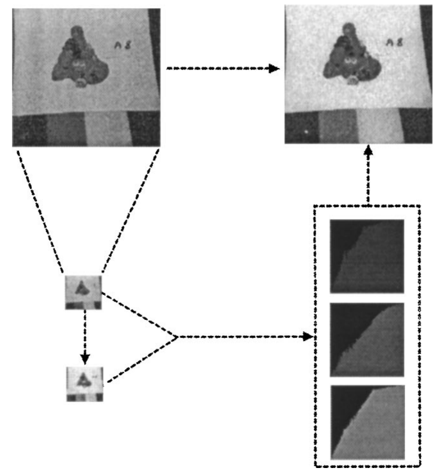

Fig. 2 LUT retinex algorithm.

To begin, a pyramid of subsampled images is computed from a logarithmic scaled version of the original image, halving the edge size at each step and averaging the pixel values, until an image of only a few pixels is produced (maximum 5×5). In this way, a series of images ranging from the original resolution Input_0 to the lower resolution $\mathrm{Input}_{\mathrm{low}}$ is created. To avoid zero argument, the logarithmic scaling is obtained in the following way:

$$\mathrm{Input}_0(x,y)=\log\left[\frac{I(x,y)+0.001}{1.001}\right].$$

The algorithm starts from the lowest level of resolution. At each level, three matrixes with dimensions corresponding to the resolution of that level, and called the new product (NP), the inner product (IP), and the old product (OP) are created. To begin with, at the lowest resolution level, the OP is initialized with the maximum value found in $\mathrm{Input}_{\mathrm{low}}$.

Each pixel in the NP is obtained by computing clockwise the ratio between the original pixel in the same position and its neighboring pixels. The reset operation is then applied on each ratio, which is averaged with the relative OP value. At the end of this operation, the NP becomes the OP, and the process is repeated for each level a prefixed number of times, according to a parameter value nIteration, which may differ from level to level, but has been set here at four iterations for each level.

$low=$ **SamplingOf**$(Input_0)$
$OP\leftarrow$ **MaxValue**$(Input_{low})$
for $j=low$ **downto** 0
  for $i=1$ **to** $nIterations_j$
    With Neighbors Do

$$IP=\frac{OP+Input_j}{Input_j}$$

$$NP=\frac{OP+IP}{2}$$

$$OP = \mathbf{ResizeDouble}(NP)$$

When all the operations of a level have been completed, the NP is enlarged to twice the edge size, by pixel replication, and becomes the OP of the next level.

At the end of the last iteration on the last level the NP will be a gray-level image constituting one of the three chromatic channels of the output image. These operations must be performed once for each chromatic channel.

Finally, the image is reversed from the logarithmic scaling;

$$\text{output} = \exp(\text{NP}),$$

and rescaled to the output device range.

4 Color-Based Image Retrieval

There is no single ''best'' representation of the color content of an image, but only multiple representations that characterize the content from different perspectives.[26] Among all the available color features we have chosen hue, saturation, and value (HSV) color moments, the color histogram, the color coherence vector, and the spatial chromatic histogram interactions as representative of the many variations of color-based retrieval algorithms. HSV color moments matching does not require any preprocessing (except color space transformation); the other methods require *a priori* color quantization.

4.1 *Color Features Quantization*

The effective and efficient computation of the color features has required a drastic reduction in the number of colors used to represent the contents of our 24-bit images. Formally, we let C be a color space, and $P = \{c_1, c_2, \ldots, c_i, \ldots, c_n | c_i \in C, n \ll \|C\|\}$ be a subset of C called a quantization space. A function Q, which maps each color in C to an element in P, and is called a quantizer is defined as:

$$Q: C \rightarrow P. \tag{4}$$

Let I be a 24-bits $n \times m$ image. Colors in I, defined on color space C, are reduced by applying the quantizer Q using the palette of c chosen colors (P).

We have used the HSV color space here, as color features in the HSV color space are known to yield better results in image retrieval than color features in other spaces.[27] The subset P was selected by nonuniformly sampling the HSV color space at 64 intervals grouping colors with the same appearance. To reduce quantization noise, and preserve chromaticity contents, a vector majority filter was applied to the quantized image, using a working window W of 3×3 pixels. Letting $\tilde{I}$ be the quantized image, each pixel $\tilde{I}(x,y)$ was assigned a value according to the following rule:

$$\tilde{I}'(x,y) = c_k \quad \text{with} \quad c_k = \underset{c_k}{\operatorname{argmax}}\{\|A_{c_i}\|\} \text{and}$$

$$A_{c_i} = \{\tilde{I}(x,y) | \tilde{I}(x,y) = c_i, \tilde{I}(x,y) \in W\}. \tag{5}$$

The discriminability of color histograms depends largely on the color quantization method used, as well as on the size of the final palette.[26] An advantage of retinex-based preprocessing is that it preserves the color appearance of the image contents. We can reasonably assume that a quantization method tuned for content-based color image retrieval can be applied as is to an image database preprocessed with a retinex algorithm.

4.2 *HSV Color Moments (HSVM)*

The color distribution of an image can be considered a probability distribution. Because any probability distribution is uniquely characterized by its central moments, color distribution can be characterized in the same manner. In Ref. 28, the first three moments (mean, variance, and skewness) of each color channel of the HSV color space were used to evaluate image similarity. The feature entries for the i'th color channel were:

$$E_i(\tilde{I}) = \frac{1}{N} \sum_{x,y} \tilde{I}_i(x,y)$$

that is, the average color channel values,

$$\sigma_i(\tilde{I}) = \left\{ \frac{1}{N} \sum_{x,y} [\tilde{I}_i(x,y) - E_i(\tilde{I})^2] \right\}^{1/2}$$

that is, the standard deviation, (6)

$$s(\tilde{I}) = \left\{ \frac{1}{N} \sum_{x,y} [\tilde{I}_i(x,y) - E_i(\tilde{I})]^3 \right\}^{1/3}$$

that is, the third root of the skewness.

The evaluation function proposed was a user-specified weighted L_1 distance. Each feature entry was weighted by a value selected by the user in view of the specific application:

$$d_{\text{mom}}(\tilde{I}_1, \tilde{I}_2) = \sum_i [w_{i1}|E_i(\tilde{I}_1) - E_i(\tilde{I}_2)| + w_{i2}|\sigma_i(\tilde{I}_1) - \sigma_i(\tilde{I}_2)| + w_{i3}|s_i(\tilde{I}_1) - s_i(\tilde{I}_2)|], \tag{7}$$

where w_{i1}, w_{i2}, $w_{i3} \geqslant 0$, and $i = H$, S, and V. In our experiments, all the weights were set at 1.

4.3 *Color Histogram (HIST)*

Color histograms are frequently used to compare images because they are simple to compute and tend to be robust regarding small changes in camera viewpoint. Retrieval using color histograms to identify objects in image databases has been investigated in Refs. 29 and 30. An image histogram refers to the probability mass function of image intensities. The histogram has been extended for color images to capture the joint probabilities of the intensities of the three color channels. More formally, each entry of the histogram is defined by

$$h_{c_i}(\tilde{I}) = N \cdot \mathrm{Prob}[\tilde{I}(x,y) = c_i]$$

$$x = 1, \ldots, m, y = 1, \ldots, n \quad \text{and} \quad N = m \times n. \tag{8}$$

Computationally, the color histogram is formed by counting the number of pixels of each color. There are several distance formulas for measuring the similarity of color histograms. Techniques for comparing probability distributions are not appropriate for color histograms because the similarity is determined by visual perception, rather than the closeness of the probability distributions. One of the most commonly used measures for color histogram comparison is the histogram intersection.[29] The intersection formula is given by

$$d[h(I_1), h(I_2)] = \frac{\Sigma \min[h_{c_i}(I_1), h_{c_i}(I_2)]}{\min[\|h(I_1)\|, \|h(I_2)\|]}, \tag{9}$$

where $\|h(I_1)\|$ and $\|h(I_2)\|$ give the area of each histogram.

4.4 *Color Coherence Vector (CCV)*

Color coherence vector histograms proposed in Ref. 31 are a refinement of color histograms, which classifies the pixels of each image: a pixel is said to be coherent if it is part of a large, similarly colored region; otherwise, it is labeled as noncoherent. To extract the color regions, the image is quantized in n colors, and an eight-neighbor connected component algorithm is then applied. Pixels in regions of a size exceeding a predefined threshold (typically, 0.5 to 1% of the whole image) are considered coherent pixels; those in smaller regions are noncoherent. For each color c_i the number of coherent pixels α_{c_i} and the number of noncoherent pixels β_{c_i} are then computed. Each entry in the CCV is thus a pair $(\alpha_{c_i}, \beta_{c_i})$, called a coherence pair. The whole coherence vector is defined as

$$\mathrm{CCV}(\tilde{I}) = \langle (\alpha_{c_1}, \beta_{c_1}), \ldots, (\alpha_{c_i}, \beta_{c_i}), \ldots, (\alpha_{c_n}, \beta_{c_n}) \rangle. \tag{10}$$

Clearly the sum $\alpha_{c_i} + \beta_{c_i}$ is the number of pixels of color c_i present in the image; the set of the sums for $i = 1, \ldots, n$ represents the color histogram. The L_1 distance can then be used to compare two CCVs:

$$\Delta \mathrm{CCV}(\tilde{I}, \tilde{I}') = \sum_{i=1}^{n} (|\alpha_{c_i} - \alpha'_{c_i}| + |\beta_{c_i} - \beta'_{c_i}|). \tag{11}$$

4.5 *Spatial Chromatic Histogram (SCH)*

Spatial chromatic histograms[32] are extended histograms that preserve, together with information about the color content of the image, the spatial distribution of each color within the image. Each entry in a SCH is composed of three values: $h_{c_i}(\tilde{I})$, the ratio of pixels in $\tilde{I}$ of color c_i; $\mathbf{b}_{c_i}(\tilde{I}) = (\overline{x}_{c_i}, \overline{y}_{c_i})$, the baricenter (in relative coordinates) of the spatial distribution of color c_i; and $\sigma_{c_i}(\tilde{I})$, the standard deviation of the distribution of color c_i. The elements of the SCH for an image are then:

$$S_{c_i}(\tilde{I}) = [h_{c_i}(\tilde{I}), \mathbf{b}_{c_i}(\tilde{I}), \sigma_{c_i}(\tilde{I})], \quad \text{where} \quad i = 1, \ldots, m. \tag{12}$$

Letting A_{c_k} be the set of pixels in the image having the same color c_k, $A_{c_k} = \{\tilde{I}(x,y) | \tilde{I}(x,y) = c_k\}$, and the three elements of the SCH can then be computed as follows:

$$\begin{aligned} h_{c_k}(\tilde{I}) &= \frac{|A_{c_k}|}{n \times m}, \quad \text{where } n \text{ and } m \\ &\quad \text{are the width and height of the image,} \\ \overline{x}_{c_k}(\tilde{I}) &= \frac{1}{n} \frac{1}{|A_{c_k}|} \sum_{\tilde{I}(x,y) \in A_{c_k}} x, \\ \overline{y}_{c_k}(\tilde{I}) &= \frac{1}{m} \frac{1}{|A_{c_k}|} \sum_{\tilde{I}(x,y) \in A_{c_k}} y, \\ \sigma_{c_k}(\tilde{I}) &= \left\{ \frac{1}{|A_{c_k}|} \sum_{p \in A_{c_k}} d[p, b_{c_k}(\tilde{I})]^2 \right\}^{1/2}, \end{aligned} \tag{13}$$

where $d(\ )$ is the Euclidean distance between two pixels.

The similarity function proposed by the authors has been designed to separate color information from spatial information:

$$\begin{aligned} f(\tilde{I}_1, \tilde{I}_2) = \sum_{c_i} \Bigg(& \min[h_{c_i}(\tilde{I}_1), h_{c_i}(\tilde{I}_2)] \\ & \times \left\{ \frac{\sqrt{2} - d[\mathbf{b}_{c_i}(\tilde{I}_1) \mathbf{b}_{c_i}(\tilde{I}_2)]}{\sqrt{2}} \right. \\ & \left. + \frac{\min[\sigma_{c_i}(\tilde{I}_1), \sigma_{c_i}(\tilde{I}_2)]}{\max[\sigma_{c_i}(\tilde{I}_1), \sigma_{c_i}(\tilde{I}_2)]} \right\} \Bigg). \end{aligned} \tag{14}$$

5 Experiments

We examine here the performance of the described color-based retrieval strategies when coupled with retinex-based prefiltering to see whether, and to what degree, retinex improves the effectiveness of retrieval, regardless of the strategy adopted. All the experiments have been carried out using the Quicklook[4] Image Search Engine.[33] The experiments were performed on a database of 310 paintings, and one of 387 ceramic objects, on both of which we have simulated a change in imaging conditions, using different illuminants and different device color spaces; two public databases, the Simon Fraser University (SFU) database of 11 objects acquired under 5 different illuminants, and the UEA uncalibrated color image database of 28 designs acquired under 14 different conditions.[34,35]

Given the different nature of the databases, we performed two series of experiments. First, on the painting and ceramics databases, we performed ''target search'' experiments: fifteen query images were randomly selected from each database, we randomly changed their device color

Fig. 3 (a) Images used to query the database of ceramics, (b) images used to query the database of paintings, (c) and (d) query images under changed illumination and device color space, (e) and (f) query images processed with the retinex 1 algorithm, and (g) and (h) query images processed with the retinex 2 algorithm.

space and illumination, and then attempted to retrieve original image from the database. The original query images, the queries after the change in imaging conditions, and the results of retinex preprocessing of the modified queries are shown in Fig. 3.

Second, our objective for the UEA and SFU databases-was: given the image of an object (or a pattern), the retrieval of all the images of the same object (or pattern). To define the query sets, we randomly selected an image from each homogeneous group, for a total of 11 images from the

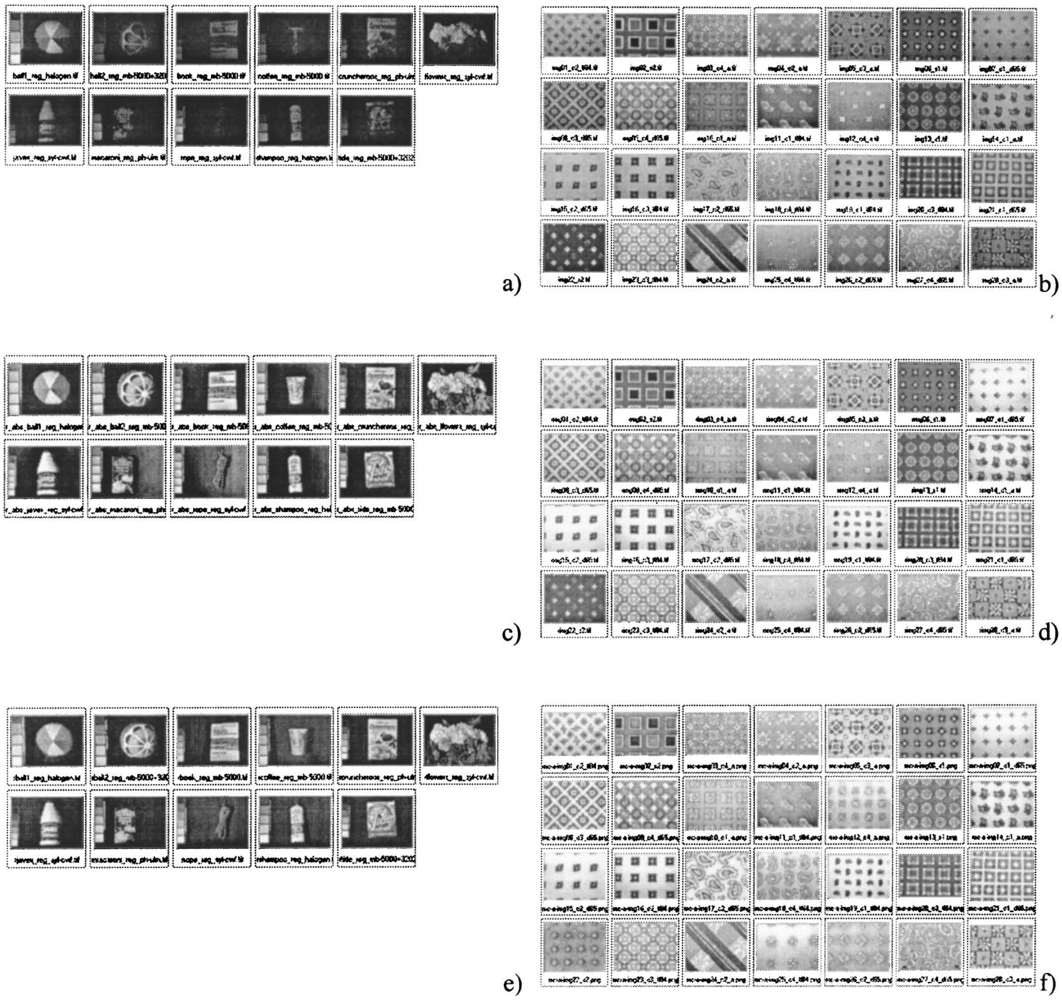

Fig. 4 (a) Images used to query the SFU database, (b) images used to query the UEA database, (c) and (d) the query images after application of the retinex 1 algorithm, (e) and (f) the query images after application of the retinex 2 algorithm.

SFU database and of 28 images from the UEA database. The original query images, and the results of retinex preprocessing of them are shown in Fig. 4. The search was performed three times, using the four retrieval strategies described in Sec. 4, first without any color processing, and then applying each retinex algorithm to both the query and the database images (see Fig. 5). The results are compared in Sec. 5.3.

5.1 *Databases of Paintings and Ceramics*

The management of archives of cultural artifacts is, without a doubt, an application in which color information is im-

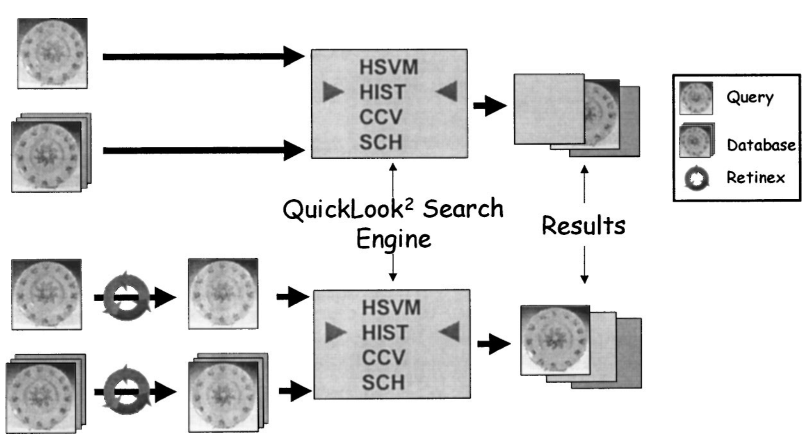

Fig. 5 Scheme of the experiment.

portant, and cannot be replaced, or emulated by other pictorial features. This was the reason for performing the first experiments on the databases of 310 painting and of 387 ceramic objects. The database of paintings contains images downloaded from the web sites of some museums and image providers. It was particularly interesting as most of the scenes depicted have unusual color casts and no reference whites. We wanted to see if this simulation of uncontrolled imaging conditions would ''confuse'' the retinex algorithms, leading them to produce inconsistent results. The ceramics object database is composed of all the color images available on the web site *Ceramica d'arte in Italia*[36] (Ceramic Art in Italy). These images, acquired from several art books and museum catalogues, were interactively modified by skilled operators to match the appearance of the real objects or, more precisely, ''to match the mental representation of those objects,'' since in many cases, the operators have never seen them. Experts in the field consider the quality of these images to be good on the whole. However, most of these images have completely black or colored backgrounds which, again, could ''confuse'' the retinex algorithms. To resemble real situations in content-based retrieval, in our experiments no background-foreground segmentation was applied. Note that retinex prefiltering was used here only for normalization purposes.

5.1.1 *Simulating changes in imaging conditions*

Our objective in these experiments was to test the robustness of retinex preprocessing in situations where the database is composed of images that have not been converted into a common color space, but preserve color coordinates deriving from various visualization systems. Image conversions among device color spaces are usually accomplished by color transform engines based on the standardized specification of the devices according to International Color Consortium (ICC) guidelines. ICC profiles for display devices contain information about the phosphors and white point chromaticities, and gamma values for each color channel. To perform the color space transform of the query image we have used the ''Profile to Profile'' function of Adobe Photoshop, and converted images from the sRGB profile[37] to a generic monitor profile, either the Generic EBU 1.5 Gamma Monitor, or the Generic EBU 1.8 Gamma Monitor. Output profiles having gamma values substantially different from the sRGB value of 2.2 were chosen to test the retinex algorithm with a nonlinear transformation.

If we consider two images representing the same scene under two different illuminants, a simple model for describing differences in corresponding pixels is a linear transformation (a 3×3 matrix of coefficients), regardless of the pixel's location. An even simpler model is one considering a diagonal matrix. The diagonal model has been proposed by Von Kries as a model for human adaptation.[38] For it to hold exactly, sensor filters must be assumed to have narrow-band properties. Although this hypothesis has not been completely verified in practice, it is frequently assumed in the definition of illuminant-invariant image descriptors.[4,5,12]

To simulate a change in illuminant conditions we used the Von Kries chromatic adaptation model. As the database images were coded in standard RGB color space, sRGB values were first converted to Commission Internationale de I'Eclairage (CIE) *XYZ* coordinates, and the Von Kries transform was then applied to map tristimulus values from the sRGB reference white to the target illuminant. The target illuminant was chosen from a set of eight different standard illuminants[39] (A, B, C, D50, D55, D75, D93, F2). The resulting *XYZ* values were then converted into sRGB values.

This procedure for simulating changes in illumination results in a uniform alteration of the color cast of the image. It could be objected that an algorithm simpler than retinex could be used in the presence of a uniform cast. But in real cases, the illumination may not always be uniform, due to the presence of shadows, and the effectiveness of retinex in discounting the illumination in real cases has already been verified.[23]

5.2 *SFU and UEA Object Databases*

The SFU database is composed of 55 images of 11 objects, acquired under five different illuminants using a Sony

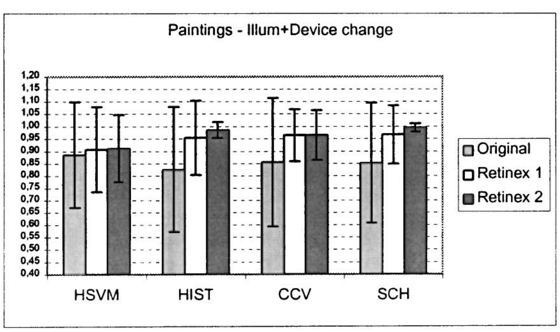

		HSVM	HIST	CCV	SCH
Original	Mean	0,885	0,826	0,854	0,851
	Std. Dev.	0,214	0,254	0,261	0,244
Retinex 1	Mean	0,911	0,954	0,964	0,967
	Std. Dev.	0,136	0,150	0,105	0,118
Retinex 2	Mean	0,908	0,986	0,965	0,995
	Std. Dev.	0,172	0,032	0,100	0,017

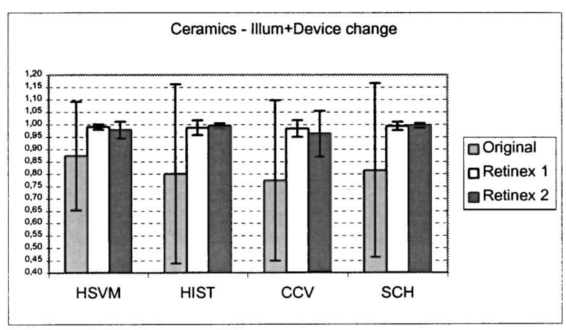

		HSVM	HIST	CCV	SCH
Original	Mean	0,873	0,800	0,773	0,815
	Std. Dev.	0,219	0,363	0,324	0,351
Retinex 1	Mean	0,991	0,988	0,984	0,994
	Std. Dev.	0,011	0,030	0,033	0,016
Retinex 2	Mean	0,978	0,995	0,963	0,997
	Std. Dev.	0,034	0,010	0,092	0,009

Fig. 6 Summary of results of the target search experiments on the database of paintings and ceramics in terms of STS.

Table 1 Percentage improvement in experiment scores compared with retreival without retinex prefiltering.

	Database			
	Paintings	Ceramics	UEA	SFU
Retinex 1	10,7763	21,5690	83,7343	35,7970
Retinex 2	11,1838	20,8283	118,4458	108,6645

DXC-930 three-chip CCD camera used with the gamma correction off, and the color temperature set at 3200 K (tungsten illuminant). More details regarding the acquisition system can be found on the Web.[34] The illuminants were the Macbeth Judge II illuminant A, a Sylvania Cool White Fluorescent, a Philips Ultralume Fluorescent, the Macbeth Judge II 5000 Fluorescent, and the Macbeth Judge II 5000 Fluorescent together with a Roscolux 3202 full blue filter, which produced an illuminant similar in color temperature to a very deep blue sky. Figure 4(a) shows the query images used for the experiment; Fig. 4(c), the results of the application of the retinex 1 algorithm; and Fig. 4(e), those of retinex 2.

The UEA uncalibrated color image database is a new, completely uncalibrated, database of 392 design images comprising 28 different designs acquired under three light sources using four digital cameras (ranging from a high-end studio camera to a low-end personal camera) and two randomly chosen commercial scanners. The images were acquired under a yellowish tungsten light, a whitish fluorescent light, and a bluish daylight[15,40] producing for each design 14 different images of often widely varying color appearance. Figure 4(b) shows the query images used for the experiment; Fig. 4(d), the results of the application of the retinex 1 algorithm; and (f), those of the retinex 2.

5.3 *Results*

5.3.1 *Experiments on the databases of ceramics and paintings*

The performance of each of the retrieval strategies on the ceramic and painting databases has been quantified in terms of a success of target search (STS) index defined as follows:

$$\mathrm{STS}=\left(1-\frac{\mathrm{Rank}-1}{N-1}\right), \tag{15}$$

where Rank is the retrieval position of the target image, and ranges from 1 to N, the number of images in the database. The STS ranges from 0, the worst result, to 1, indicating a perfect retrieval.

Figure 6 summarizes the results of the application of retinex 1 and of retinex 2. In each diagram, the STS values are reported for all the features used in retrieval, first with no color processing (''original''), and then after retinex processing.

The search on retinex preprocessed images constantly outperformed the search on nonfiltered images. The average improvement in performance was 16.2% for the retinex 1 algorithm and 16% for the retinex 2 algorithm (see Table 1 below). No significant differences in the performance of the two retinex algorithms were observed in the experiments on these databases.

5.3.2 *Experiments on SFU and UEA databases*

The results of searching the SFU and UEA databases are presented in Fig. 7. The index here is not the STS score, but the number of images of the same query object ranked among the first 5 retrieved for the SFU database, or among the first 14 for the UEA database, 5 and 14 being the number of different imaging conditions applied, respectively.

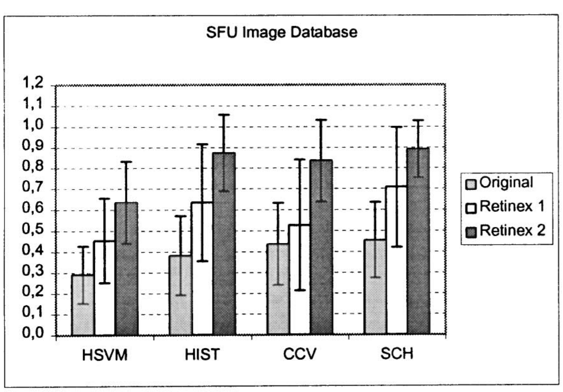

		HSVM	HIST	CCV	SCH
Original	Mean	0,291	0,382	0,436	0,455
	Std. Dev.	0,138	0,189	0,196	0,181
Retinex 1	Mean	0,455	0,636	0,527	0,709
	Std. Dev.	0,202	0,280	0,313	0,288
Retinex 2	Mean	0,636	0,873	0,836	0,891
	Std. Dev.	0,196	0,185	0,196	0,138

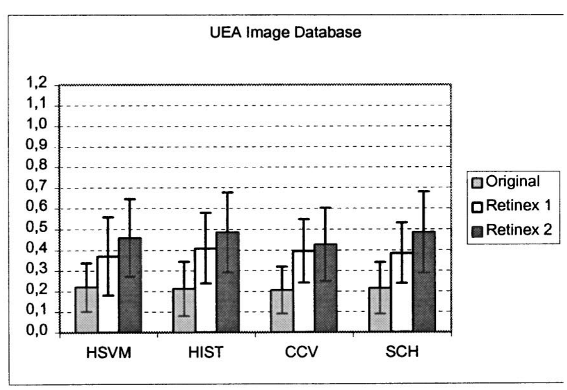

		HSVM	HIST	CCV	SCH
Original	Mean	0,219	0,212	0,204	0,214
	Std. Dev.	0,118	0,133	0,115	0,126
Retinex 1	Mean	0,370	0,408	0,395	0,385
	Std. Dev.	0,189	0,170	0,154	0,146
Retinex 2	Mean	0,459	0,485	0,426	0,485
	Std. Dev.	0,187	0,192	0,177	0,196

Fig. 7 Summary of results of the experiments performed on the SFU and UEA databases.

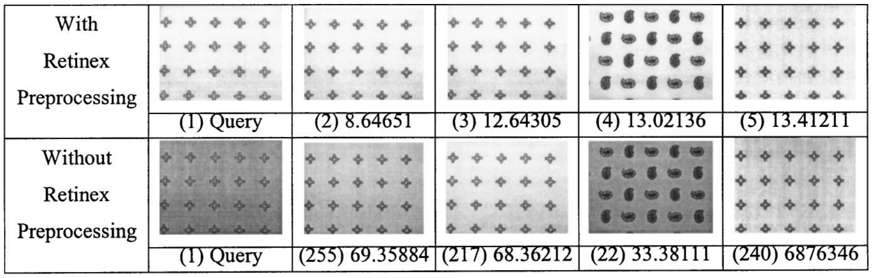

Fig. 8 Retinex, discounting the illuminant, renders the relevant images similar to the query but, if most of the images in the database are similar in color content, some non relevant images may also be found in top positions.

The score ranges from a minimum of 1 (i.e., only the query image itself has been found) to a maximum of 5 or 14 (i.e., all the images of the same object are ranked in the top positions). The numerical results were normalized to compare the performances on the two databases:

$$\text{score}=\frac{\text{objects retrieved in top } N}{N}, \quad N=5 \text{ or } 14. \tag{16}$$

Figure 7 shows that the use of retinex prefiltering produces an evident improvement in query effectiveness under all conditions and with all retrieval strategies. It also registers a difference in the performance of the two retinex algorithms, with a mean improvement of 59.8% for retinex 1 and of 113.6% for retinex 2 (see Table 1). This is due to the different characteristics of these algorithms, as explained in Sec. 3: the multilevel algorithm (retinex 2) computes the ratio-threshold-reset-product at various subsampling levels, while the simplified LUT version (retinex 1) computes the same operations at only one given subsampling resolution, the 1/8 here, losing significant spatial color correlation as a consequence. For example, we can see in Fig. 4(a) that the SFU database is undersampled to 1/8, the broad black background that results produces very dark pictures, further decreasing the color dynamics. Retinex 2, operating at higher resolution levels, maintains the color dynamics of the whole images.

Given the improvement in results after retinex prefiltering, we must ask ourselves why the searching the UEA database is less successful than searching the SFU. Two possible reasons are that the UEA images, designed to test different color constancy methods and not database retrieval algorithms, have less color dynamics than the SFU images, and that many of the pictures have similar color palettes.

If we compare the variance of target search results on the databases of ceramics and paintings before and after applying the retinex algorithms, we see that it is greatly reduced by the application of retinex filters, and that both algorithms produce similar results. On the contrary, on the UEA and SFU databases, although the overall results improve with the application of retinex, the variance of the score increases when retinex 1 is applied to prefilter the SFU database.

Since the color features used for retrieval totally, or at least partially, disregard spatial information and the database images have a limited palette and a similar color content, some nonrelevant filtered images may be ranked in top positions. This can be seen, for example, in Fig. 8. Figure 8(a) shows the top five images retrieved taking the first image to query the UEA database and using the SCH (Sec. 4.5) retrieval method. Both the query and the database images were preprocessed with the retinex 2 algorithm. The rank and distance from the query are reported under each image. Among the relevant images we find a nonrelevant one, the visual appearance and color content of which are quite close to those of the relevant images. While results are not perfect, however, retinex prefiltering has still greatly improved them, as we can see in Fig. 8(b), which shows the ranks and distances for the same images, retrieved without retinex prefiltering. Retinex effectively discounts the differences in illuminant in the images of the same object, rendering them much more similar to each other greatly reducing the distances from the query and consequently improving the ranking.

6 Conclusions

We examined the results of retinex-based preprocessing in color-based image retrieval. Two different implementations of the algorithm were used in searching four databases with four commonly implemented color-based retrieval strategies.

The results demonstrate that retinex preprocessing clearly improves the effectiveness of the retrieval strategies applied.

The two retinex algorithms used here produced very similar results on the painting and ceramics databases (on which we simulated a change in imaging conditions), while on the SFU and the UEA databases, acquired under different lighting conditions (and therefore the product of more complex and unpredictable color transformations), the latest version of McCann's algorithm outperforms the Brownian LUT retinex designed by some of the authors of this paper. The comparison may be biased by the tuning of some of the retinex parameters, and perhaps the relationship between the search strategy and retinex parameters could be better, and, possibly, automatically adjusted. A more thorough investigation is planned.

We must not forget that color alone cannot suffice to index large image databases, not even those acquired under standard lighting conditions: in real image searches, the combination of color with other visual features is a must.

In some cases of retrieval, where it might be necessary to preserve the images as they are since their color cast is part of their semantic content, (the picture of a sunset, for example), the application of retinex, or any other color normalization method, would not be useful. Integrating retinex preprocessing in a general purpose search engine would make it possible to exploit image classification in order to identify those special cases. Or the user could be given the choice of whether or not to apply retinex preprocessing, just as the retrieval strategy can be selected.[33,41,42]

We plan to continue our testing on a large image database of public domain, and have taken the first steps in that direction.

References

1. Y. Rui and T. S. Huang, ''Image retrieval: current technologies, promising directions, and open issues,'' *J. Visual Commun. Image Represent* **10**, 39–62 (1999).
2. A. Del Bimbo, *Visual Information Retrieval*, Morgan Kaufmann Publishers, San Francisco (1999).
3. D. Androutsos, ''A novel vector-based approach to color image retrieval using vector angular-based distance measure,'' *Comput. Vis. Image Underst.* **1–2**, 46–58 (1999).
4. B. V. Funt and G. D. Finlayson, ''Color constant color indexing,'' *IEEE Trans. Pattern Anal. Mach. Intell.* **17**, 522–529 (1995).
5. T. Gevers and A. W. M. Smeulders, ''A comparative study of several color models for color image invariant retrieval,'' in *Proc. 1st Int. Workshop on Image Database & Multimedia Search*, pp. 17–26, Amsterdam, Netherlands (1996).
6. T. Gevers and A. W. M. Smeulders, ''Color based object recognition,'' *Pattern Recogn.* **32**, 453–464 (1999).
7. G. D. Finlayson, S. S. Chatterjee, and B. V. Funt, ''Color angular indexing and image retrieval,'' in *Image Description and Retrieval*, Vicario, Ed., Plenum, New York (1998).
8. M. K. Mandal, T. Aboulnasr, and S. Panchanathan, ''Illuminant invariant image indexing using moments and wavelets,'' *J. Electron. Imaging* **7**(2), 282–293 (1998).
9. B. A. Wandell, ''Standard surface-reflectance model and illuminant estimation,'' *J. Opt. Soc. Am. A* **6**, 576–584 (1989).
10. P. M. Hubel and S. Hordley, ''Color by correlation,'' in *Proc. 5th Color Imaging Conf.*, pp. 6–10, Scottsdale, AZ (1997).
11. B. V. Funt and V. Cardei, ''Committee-based color constancy,'' *J. Opt. Soc. Am. A* **11**(11), 3011–3020 (1994).
12. E. Land and J. McCann, ''Lightness and retinex theory,'' *J. Opt. Soc. Am.* **61**(1), 1–11 (1971).
13. G. D. Finlayson, B. Schiele, and J. L. Crowley, ''Comprehensive colour image normalization,'' in *Proc. 5th Eur. Conf. on Computer Vision* pp. 475–490 (1998).
14. B. V. Funt, K. Barnard, and L. Martin, ''Is machine colour constancy good enough?'' in *Proc. 5th Eur. Conf. on Computer Vision* pp. 445–459 (1998).
15. G. Finlayson and G. Schaefer, ''Color indexing across devices and viewing conditions,'' in *Proc. IEEE 2nd Int. Workshop on Content-Based Multimedia Indexing*, pp. 215–221, Brescia, Italy (2001).
16. B. V. Funt, F. Ciurea, and J. McCann, ''Retinex in Matlab,'' in *Proc. 8th Color Imaging Conf.*, pp. 112–121, Scottsdale, AZ (2000).
17. D. Marini and A. Rizzi, ''A computational approach to color adaptation effects,'' *Image Vis. Comput.* **18**(13), 1005–1014 (2000).
18. J. J. McCann, ''Lesson learned from mondrians applied to real images and color gamuts,'' *IS&T Rep.* **14**, 1–13 (1999).
19. D. J. Jobson, Z. Rahman, and G. A. Woodel, ''Properties and performance of a center/surround retinex,'' *IEEE Trans. Image Process.* **6**(3), 451–462 (1997).
20. B. A. Wandell, ''Analysis of the retinex theory of color vision,'' *J. Opt. Soc. Am. A* **3**, 1651–1661 (October 1986).
21. A. Moore, J. Allman, and R. Goodman, ''A real-time neural system for color constancy,'' *IEEE Trans. Neural Netw.* **2**, 237–246 (1991).
22. D. Marini, A. Rizzi, and L. De Carli, ''Multiresolution retinex: comparison of algorithms,'' in *Proc. 1st Int. Conf. on Color Graphics Image Processing*, Cépaduès-éditions, pp. 106–110, St. Etienne, France (2000).
23. D. Marini, A. Rizzi, and M. Rossi, ''Color constancy measurement for synthetic image generation,'' *J. Electron. Imaging* **8**(4), 394–403 (October 1999).
24. D. Marini and A. Rizzi, ''A color appearance approach to image database visual retrieval,'' EI2000, IS&T/SPIE's Electronic Imaging 2000, San Jose, California, USA, January 24–28 (2000). *Proc. SPIE* **3964**, 186–195 (2000).
25. D. Saupe and H. O. Peitsen, ''Algorithms for random fractals,'' in *The Science of Fractal Images*, D. Saupe and H. O. Peitsen, Eds., Springer, New York (1988).
26. R. Schettini, G. Ciocca, and S. Zuffi, ''A survey on methods for colour image indexing and retrieval in image databases,'' in *Color Imaging Science: Exploiting Digital Media*, R. Luo and L. MacDonald, Eds., J. Wiley, New York (2001).
27. W. Y. Ma and H. J. Zhang, ''Content-based image indexing and retrieval,'' Chap. 13 in *The Handbook of Multimedia Computing*, B. Furht, Ed., CRC Press (1998).
28. M. Stricker and A. Dimai, ''Color indexing with weak spatial constraints,'' in *Proc. SPIE* **2670**, 29–40 (1996).
29. M. Swain and D. Ballard, ''Color indexing,'' *Int. J. Comput. Vis.* **7**(1), 11–32 (1991).
30. J. R. Smith and S. F. Chang, ''Tools and techniques for color image retrieval,'' *Proc. SPIE* **2670**, 426–437 (1996).
31. G. Pass, R. Zabih, and J. Miller, ''Comparing images using color coherence vectors,'' in *Proc. 4th ACM Multimedia Conf.* pp. 65–73 (1996).
32. L. Cinque, S. Levialdi, and A. Pellicano, ''Color-based image retrieval using spatial-chromatic histograms,'' *IEEE Multimedia Syst.* **99**(2), 969–973 (1999).
33. G. Ciocca, I. Gagliardi, and R. Schettini, ''Quicklook2: an integrated multimedia system,'' *J. Vis. Lang. Comput.* **12**(1), 81–103 (2001).
34. The SFU database can be found at http://www.cs.sfu.ca/~colour/data/objects_under_different_lights/index.html.
35. The UEA database can be found at http://www.sys.uea.ac.uk/~gs/research/CATSI/database.html.
36. The Cora Donation, Medieval and Renaissance Ceramics of the International Museum of Ceramics of Faenza, http://jargo.itim.mi.cnr.it/Ceramica/.
37. M. Nielsen, M. Stokes, ''The creation of the sRGB ICC profile,'' in *Proc. 6th Color Imaging Conf.*, pp. 253–257, Scottsdale, AZ (1998).
38. M. D. Fairchild, *Color Appearance Models*, Addison-Wesley, Reading, MA (1998).
39. R. W. G. Hunt, *Measuring Colour*, 3rd ed., Fountain Press, Kingston-upon-Thames, England (1998).
40. G. D. Finlayson, G. Schaefer, and G. Y. Tain, ''The UEA uncalibrated colour image database,'' Technical Report SYS-C00/, School of Information System, University of East Anglia, Norwich, UK (2000).
41. R. Schettini, G. Ciocca, A. Valsasna, C. Brambilla, and M. De Ponti, ''A hierarchical classification strategy for digital documents,'' *Pattern Recogn.* **35**, 1759–1769 (2002).
42. R. Schettini, A. Valsasna, C. Brambilla, and M. De Ponti, ''An indoor/outdoor/close-up photo classifier,'' in *Proc. IX Color Imaging Conf.*, pp. 35–40, Scottsdale, AZ (2001).

Gianluigi Ciocca received his degree (laurea) in computer science from the University of Milan in 1998 and has since been a fellow with the Institute of Multimedia Information Technologies, Italian National Research Council, where his research has focused on the development of systems for the management of image and video databases and the development of new methodologies and algorithms for automatic indexing.

Daniele Marini is an associate professor and graduated in physics in 1972. Since 1978 his research with the Department of Information Sciences, Università di Milano, has encompassed several areas of graphics and image processing, with specific reference to visual simulation, realistic visualization, classification, image recognition, and compression. In Italy he pioneered image synthesis. He contributed to the foundation of the journal *PIXEL* and he was one of the founders of the Aicographics Association. In 1982 he created Eidos, the first Italian company specialized in advanced image processing untill 1988. Since 1997 he has been a member of the National University Council. In 1998 he was appointed supervisor and coordinator of the initiatives on multimedia at Triennale di Milano. He has published more than 140 papers as well as three books. He is currently teaching computer graphics and image processing for the graduation programs on informatics at the Università degli Studi di Milano.

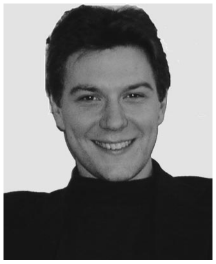

Alessandro Rizzi received his degree in computer science from University of Milano and his PhD degree in information engineering from the University of Brescia. He has been professor of information systems and computer graphics. He is currently an assistant professor with the University of Milano teaching human computer interaction. His main research topic is the use of color information in computer vision with particular attention to color adaptation mechanisms.

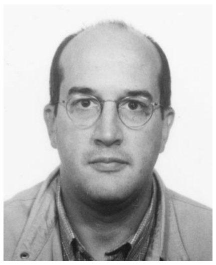

Raimondo Schettini is an associate professor at DISCO, University of Milano Bicocca. He has been associated with Italian National Research Counci (CNR) since 1987. In 1994 he moved to the Institute of Multimedia Information Technologies, where he is currently in charge of the Imaging and Vision Lab. He has been team leader in several research projects and published more than 130 refereed papers on image processing, analysis and reproduction, and on image content-based indexing and retrieval. He is member of the CIE TC 8/3. He has been General Co-Chairman of the 1st Workshop on Image and Video Content-based Retrieval (1998), and general co-chair of the EI Internet Imaging Conferences (2000-2002). He was general-co-chair of the First European Conference on Color in Graphics, Imaging and Vision (CGIV'2002).

Silvia Zuffi received her degree in electronic engineering in 1995 from the University of Bologna, Italy. She was a software engineer until 1997, when she joined the Movement Analysis Laboratory at the Istituti Ortopedici Rizzoli, Bologna, Italy. Since 1999 she has been with the Italian National Research Council (CNR) in the Institute for Multimedia Information Technologies. Her research interests include color reproduction and device characterization.