Combining local binary patterns and local color contrast for texture classification under varying illumination

Claudio Cusano,[1,†] Paolo Napoletano,[2,*,†] and Raimondo Schettini[2,†]

[1]*Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Via Ferrata 1, 27100 Pavia, Italy*
[2]*Department of Informatics, Systems and Communication, University of Milan-Bicocca, Viale Sarca 336, 20126 Milano, Italy*
**Corresponding author: napoletano@disco.unimib.it*

Received November 14, 2013; revised March 20, 2014; accepted April 21, 2014; posted May 6, 2014 (Doc. ID 201282); published June 12, 2014

This paper presents a texture descriptor for color texture classification specially designed to be robust against changes in the illumination conditions. The descriptor combines a histogram of local binary patterns (LBPs) with a novel feature measuring the distribution of local color contrast. The proposed descriptor is invariant with respect to rotations and translations of the image plane and with respect to several transformations in the color space. We evaluated the proposed descriptor on the Outex test suite, by measuring the classification accuracy in the case in which training and test images have been acquired under different illuminants. The results obtained show that our descriptor outperforms the original LBP approach and its color variants, even when these are computed after color normalization. Moreover, it also outperforms several other color texture descriptors in the state of the art.

OCIS codes: (100.2960) Image analysis; (100.5010) Pattern recognition; (100.3008) Image recognition, algorithms and filters; (330.1690) Color.
http://dx.doi.org/10.1364/JOSAA.31.001453

1. INTRODUCTION

The combination of color and texture features is very common in texture classification [1]. These two visual cues are often treated as orthogonal. In fact, studies about human perception suggest that texture and color information are processed separately [2].

The importance of color in texture classification has been thoroughly explored by Bianconi *et al.* [3]. They proposed a taxonomy for the classification of the color/texture descriptors in the state of the art. They also evaluated these descriptors on a standard data set. The results show that color is indeed a very important feature for texture classification. In fact, they obtained performance very close to the maximum just by using marginal histograms in the sRGB color space. The evaluation of Bianconi *et al.*, however, does not take into account changes in illumination. When lighting conditions are allowed to change, the increased intra-class variability may easily cause a drop in classification accuracy. The reader can refer to the work of Drbohlav and Leonardis [4] for an in-depth comparative analysis of texture methods under varying viewpoints and illumination, as well as to the work of Kandaswamy *et al.* [5] for a comparison among texture analysis schemes under nonideal conditions in general.

In practice, variable lighting conditions can be dealt with by simply ignoring the contribution of color and by focusing exclusively on the texture. This simple approach is supported by the study of Mäenpää and Pietikäinen [1], in which a variety of strategies for the combination of color and texture features have been experimented. They clearly confirmed that color features can be extremely effective for the classification of textures when the lighting conditions are fixed. However, under varying illumination conditions they found that the best performance is achieved by using histograms of local binary patterns (LBPs) computed on the gray-scale image.

Besides the trivial solution of completely disregarding color, there are two main approaches that address the issue of changes in lighting conditions: the first consists in performing the feature extraction step after a chromatic normalization procedure such as gray world [6], Retinex [7], gray edge [8,9], or comprehensive color image normalization [10]. The second consists in extracting invariant features from the images.

An example of the first approach has been proposed by Khan *et al.* [11]. They considered a diagonal/offset model for illumination variations, on the basis of which they deduced an image normalization transformation. Then, Gabor filters are extracted from the normalized images. The features they obtained provided very good results in the context of texture classification across illumination changes.

The second approach has been followed by Finlayson *et al.*, who proposed rank-based features obtained from invariant color representations [12]. Seifi *et al.*, instead, proposed to characterize color textures by analyzing the rank correlation between pixels located in the neighborhood from each other. They obtained a correlation measure that is related to the colors of the pixels, and is not sensitive to illumination changes [13].

In this paper we present a texture descriptor that properly combines a histogram of LBPs with a feature encoding the distribution of local color contrast (LCC). LBPs [14] are among the most effective texture features in the state of the art. They are computed on the gray-scale image, and have been designed to be robust against changes in the lighting conditions, such as those causing variations in the local contrast.

The second feature, the distribution of LCC, is designed to preserve a useful part of color information and, at the same

time, to discard the part that is most affected by changes in illumination. Our previous investigation on this topic showed us not only the effectiveness of measuring the local variability of the color of the pixels, but also how well this kind of information complements the LBP approach [15,16]. This fact distinguishes our approach from previous attempts to incorporate color information in the LBP framework [17,18], which resulted in an excess of sensitivity with respect to changes in lighting conditions.

To evaluate our descriptor we used a data set that has been previously used in several other works on this topic [1,11,19,20]. This is the Outex-14 data set, a collection of texture images that have been acquired under three different illuminants [21].

To prove the effectiveness of the proposed approach, we did the following:

- we evaluated the influence of the parameters on the performance of the descriptor;
- we compared the performance of the proposed descriptor with that obtained with several LBP-based features, with and without the use of color normalization methods;
- we compared the performance of our descriptor with that of several other methods from the state of the art.

2. PROPOSED DESCRIPTOR

In this work we present a texture descriptor that is designed to be robust with respect to changes in illumination. Differently from most works in the state of the art, where this kind of robustness is achieved by completely disregarding color, we propose a measure of LCC that preserves a useful part of color information and, at the same time, discards the part that is often affected by changes in illumination.

The descriptor is formed by concatenating two feature vectors: the first consists in a histogram of LBPs, while the second represents the distribution of LCC.

A. Local Binary Patterns

Despite their remarkable discriminative power, the definition of LBPs is quite simple [22]. Briefly, LBPs are computed by thresholding the value of the neighbors of each pixel with the value of the pixel itself. More in detail, given the number n of neighbors and the neighborhood radius r, for each pixel the $\mathrm{LBP}_{n,r}$ code is computed as follows:

$$\mathrm{LBP}_{n,r}(\hat{g}) = \sum_{p=0}^{n-1} s(g_p - \hat{g}) 2^p, \tag{1}$$

where $\hat{g}$ is the gray level of the current pixel, $g_0, \ldots, g_{n-1}$ are the gray levels of its neighbors, and s is the step function defined as $s(x) = 1$ if $x \geq 0$; $s(x) = 0$ otherwise. The n neighbors considered lie on a circle of radius r centered on the current pixel: the gray value g_p is obtained by interpolating the image at a displacement $(r\cos(2\pi p/n), r\sin(2\pi p/n))$.

To form a feature vector, the occurrences of the local patterns are counted into a histogram. More precisely, the feature extraction procedure is the following: (i) input images are converted to gray-scale, (ii) Eq. (1) is applied for each pixel to compute the local patterns, (iii) the occurrences of the patterns are counted in a histogram, and (iv) the histogram is normalized to sum one. Conversion to gray-scale can be performed in several ways [23,24].

In practice, not all the 2^n patterns are equally significant. Only patterns describing a somewhat regular neighborhood are considered. These patterns are called "uniform" and are defined as those for which there are at most two 0/1 transitions between adjacent bits in the code. All nonuniform patterns are accounted for in a single bin of the histogram.

In this work we set $n = 16$ and $r = 2$. Therefore, the LBP histogram has 243 bins (242 for the uniform patterns and one for the nonuniform patterns; see Pietikäinen *et al.* [22]).

B. Local Color Contrast

Color information is usually represented as color vectors $\mathbf{c} = (c_1, c_2, c_3)$ consisting of triplets of values corresponding to the components along the axis of a given color space. For the sake of simplicity we considered only the RGB color space, but the method described here can be easily computed in other color spaces. The LCC is measured by comparing the color at a given location with the average color in a surrounding neighborhood. To make it robust against changes in the color of the illuminant, the LCC θ is computed in terms of the angular difference between the two color vectors:

$$\theta = \arccos\left(\frac{\langle \hat{\mathbf{c}}, \bar{\mathbf{c}} \rangle}{\|\hat{\mathbf{c}}\| \cdot \|\bar{\mathbf{c}}\|}\right), \tag{2}$$

where $\hat{\mathbf{c}}$ is the color of the current pixel, $\bar{\mathbf{c}}$ is average color of the neighborhood, and $\langle \cdot, \cdot \rangle$ and $\| \cdot \|$ denote the inner product and the Euclidean norm. When the denominator is null, we simply set $\theta = 0$.

Inspired by the LBP approach, the average color is computed on a circular neighborhood of radius r and cardinality n:

$$\bar{\mathbf{c}} = \frac{1}{n} \sum_{p=0}^{n-1} \mathbf{c}_p, \tag{3}$$

where the vector $\mathbf{c}_p$ is obtained by a bilinear resampling of the color image at the location of $\hat{\mathbf{c}}$ displaced by $(r\cos(2\pi p/n), r\sin(2\pi p/n))$. Figure 1 shows a diagram of the computation of the LCC θ. A LCC map is formed by computing the LCC at all the locations in the image. For instance, Fig. 2 shows an example of the LCC map obtained with

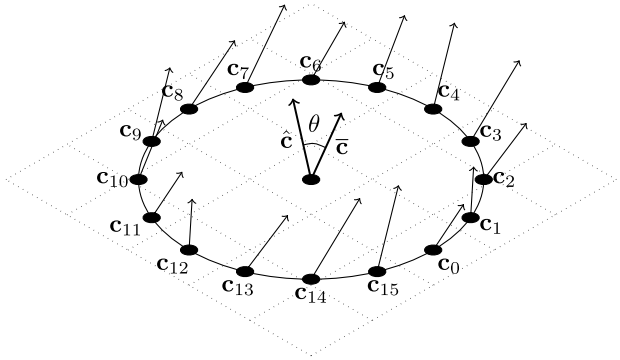

Fig. 1. Diagram representing the computation of the LCC corresponding to the neighborhood defined by the parameters $n = 16$ and $r = 2$. The color vector $\bar{\mathbf{c}}$ is computed by averaging the 16 neighbors. Then, the local contrast is measured by the angle θ between the average $\bar{\mathbf{c}}$ and the central color vector $\hat{\mathbf{c}}$.

different neighborhoods. It can be noticed that the values are always positive: this is due to the choice of the RGB color space, where the components always assume positive values. Therefore the angles between two RGB vectors necessarily fall in the range $[0, \pi/2]$. We used histograms to represent the distribution of the LCC. The angular difference is uniformly quantized in Q bins. Actually, the cases in which θ exceeds $\pi/4$ are exceptional. Therefore, we decided to quantize only the range $[0, \pi/4]$. The very few cases in which the upper bound is exceeded are accounted for in the last bin of the histogram. Figure 2 also shows histograms obtained with $Q = 256$ and three different neighborhoods. The figure shows that the larger the neighborhood, the higher the color contrast values. This is expected, since for large neighborhoods the average color tends to diverge more from the color of the central pixel.

The feature proposed here has been designed by taking into account our previous investigation where we introduced and evaluated the intensity color contrast descriptor (ICCD) [15]. This was based on the measure of color contrast defined as the angular difference between the color vector of the pixel and the average of the normalized color vectors in a square neighborhood. The average color is defined as

$$\bar{\mathbf{c}} = \frac{1}{(2W+1)^2} \sum_{i=-W}^{W} \sum_{j=-W}^{W} \frac{\mathbf{c}_{ij}}{\|\mathbf{c}_{ij}\|}, \tag{4}$$

where W determines the size of the neighborhood, and $\mathbf{c}_{ij}$ denotes the color of the pixel at a displacement (i, j) with respect to the central pixel. Null vectors are ignored in the summation. In that work we observed that (i) the unit normalization of the color vectors makes the descriptor sensitive to image noise and (ii) better performance is obtained for small neighborhoods. In the new LCC descriptor we removed the normalization and we reduced to the minimum the size of the neighborhood. The older ICCD assigned a uniform weight to all the neighbors, while LCC weights more the brighter pixels, for which color information is more reliable. We also realized that a circular shape makes the LCC invariant with respect to rotations on the image plane. Moreover, since LBPs and LCC share the same neighborhood, it would be easier to design multiscale classification strategies (this topic, however, has not been addressed in this work). As a consequence, with respect to ICCD, the LCC has one parameter less (the size of the neighborhood), rotation invariance, and a slightly better performance.

In summary, the distribution of LCC enjoys the following properties:

- it is invariant with respect to rotations and translations of the image plane, due to the circular shape of the neighborhood and to the use of a histogram representation;
- it is invariant to uniform scalings in the color space, due to the normalization of the vectors in Eq. (2);
- it is invariant to rotations in the color space, because of the use of an angular distance.

The last two properties make the distribution of LCC robust with respect to changes in the illumination conditions. Moreover, the final descriptor can be computed very quickly; in fact feature extraction requires a number of operations that is proportional to the total number of pixels M, and to the size n of the neighborhood ($O(Mn)$), which corresponds to the same computational complexity of LBP-like features.

Note that the LCC histogram cannot be considered purely as a color feature. In fact it completely disregards the information about the "dominant color" (i.e., uniform patches have zero color contrast, independently of their color).

C. Combining of LBPs and LCC

The final descriptor is obtained by combining the LBP and LCC features. A comparison of the combination strategies for color and texture features has been discussed by Khan *et al.* [25]. They distinguished among two main approaches:

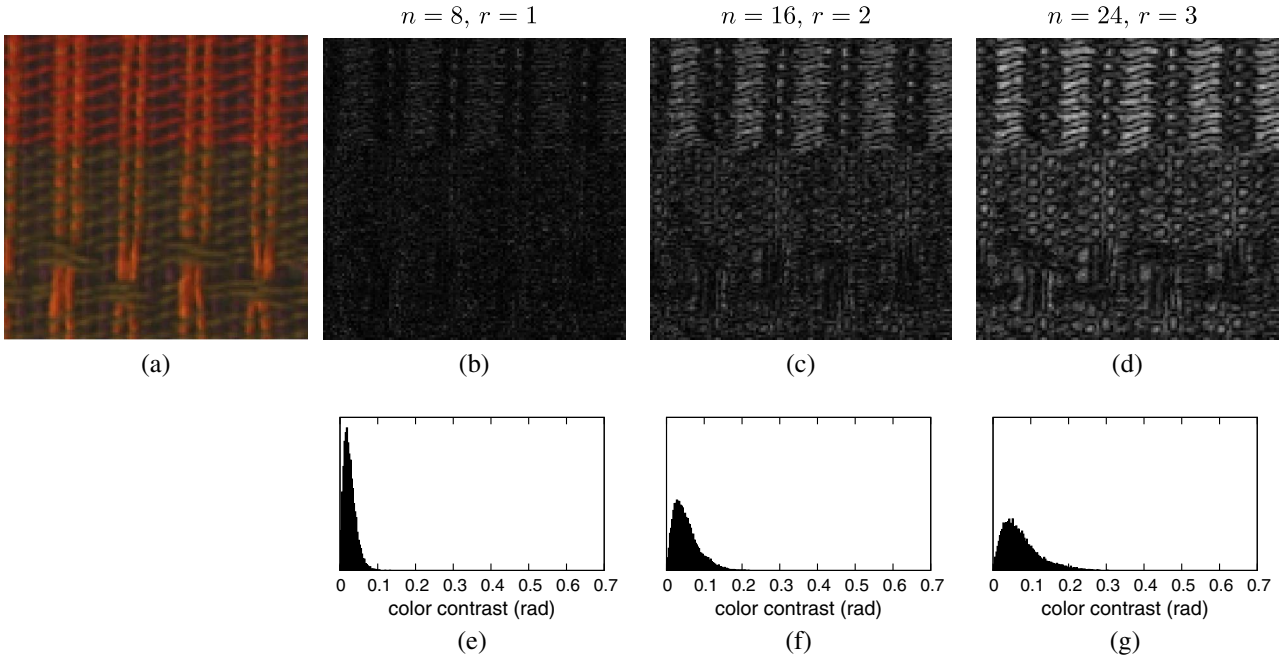

Fig. 2. (b)–(d) LCC maps and (e)–(f) LCC distributions obtained from a sample image (a). The neighborhoods considered are, from left to right, $n = 8$ and $r = 1$, $n = 16$ and $r = 2$, and $n = 24$ and $r = 3$. In the maps, brighter pixels stand for higher values of LCC.

early and late fusion (they also proposed a more sophisticated *Portmanteau* fusion, not considered here). Early fusion involves binding the two visual cues at the pixel level. Late fusion consists in processing the visual cues independently and in combining them at the image level. They concluded that late fusion provides superior performance in those cases in which one of the visual cues changes significantly. Therefore, since we are clearly addressing one of these cases (changes in the illuminant color, for instance, may leave unchanged the achromatic texture) we decided to adopt the late fusion strategy.

The combined descriptor D is formed by concatenating the histogram of LBPs (H_{LBP}) and the histogram of the color contrast values (H_{LCC}). The relative importance of the two feature vectors may depend on the specific task at hand. Therefore, we decided to introduce a final parameter w that determines a linear scaling to be applied to the second feature vector before the final combination:

$$D = \begin{pmatrix} H_{\mathrm{LBP}} \\ w \times H_{\mathrm{LCC}} \end{pmatrix}. \tag{5}$$

When w is zero, only the LBP histogram is considered. As w increases, the contribution of LCC grows, and that of LBPs becomes progressively more marginal. The optimal value of w clearly depends on the amount of variation of illumination conditions.

3. EXPERIMENTS

In the experiments, we first tuned the descriptor and evaluated its sensitivity with respect to the parameters. Then, we compared its performance with that obtained by replacing the LBP feature described in Section 2.A, with other LBP-based features. A further comparison has been made with the results that can be achieved by computing the texture descriptors after the application of color normalization methods. Finally we compared the performance of our descriptor against that of several other methods from the state of the art.

For additional results on the same data set, the reader can refer to the works of Bianconi *et al.* [3] and to that of Mäenpää and Pietikäinen [1].

A. Data Set and Classification Problems

To study the robustness of the proposed features to lighting variations we considered the Outex-14 test suite, which is part of the Outex collection [21] (the data set is available at the address http://www.outex.oulu.fi). The suite contains images that depict textures of 68 different classes acquired under three different light sources, each positioned differently: the 2856 K incandescent CIE A (denoted as "inca"), the 2300 K horizon sunlight (denoted as "horizon"), and the 4000 K fluorescent TL84 (denoted as "TL84"). Texture images have been obtained by subdividing texture photographs in 20 subimages of size 128×128 pixels. The photographs have been acquired with a Sony DXC-755P camera calibrated using the "inca" illuminant. Images are encoded in linear camera RGB space [1]. The suite contains, for each light source, 1360 images: 680 for training and 680 for test. Figure 3 shows the intra-class variability for some of the 68 classes across the three illuminants.

We experimented with different subsets of texture images from the Outex-14 data set obtained by considering all the possible combinations of illuminants for the training and the test sets. For instance, considering the training and test set from "inca" we obtained the subset "inca/inca" (i/i), while considering the training set from "inca" and test set from "horizon" or from "TL84" we obtained the subsets "inca/horizon" (i/h) or "inca/TL84" (i/t). In this way, we obtained nine different classification problems.

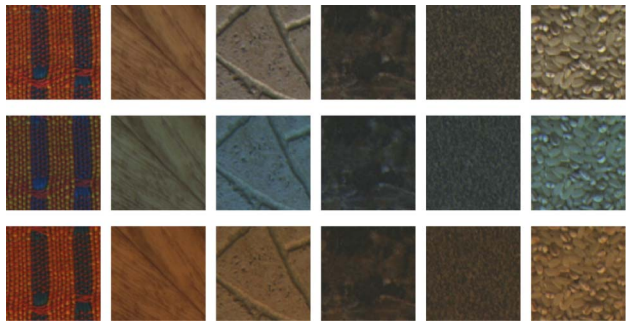

Fig. 3. Examples of images in six of the 68 classes. Each column contains images from a different class. The first row contains images under "inca", the second row contains images under "TL84", and the last row contains images under "horizon" (best viewed in color).

Note that the authors of the data set indicate the "inca" illuminant for the training images, and the union of the images taken under the other two illuminants for the test set. Other works about texture analysis under variable illumination detailed the results by separating the performance with the "horizon" and the "TL84" test sets (our i/h and i/t subsets). It is also important to point out that the Outex collection contains another suite, the Outex-13, that is composed of the subset of images, contained in the Outex-14, acquired under "inca" and thus corresponding to the subset i/i.

In all the experiments we used the nearest neighbor classification strategy: given an image in the test set, its distance with respect to all the training images is computed. The prediction of the classifier is the class of the closest training set image. For this purpose, after some preliminary tests in which we evaluated the most common distance measures, we decided to use the χ^2 distance

$$\chi^2(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \sum_{i=1}^{N} \frac{(x_i - y_i)^2}{x_i + y_i}, \tag{6}$$

where $\mathbf{x}$ and $\mathbf{y}$ are two feature vectors of nonnegative values. The χ^2 distance has been demonstrated to provide good performance when applied to histograms [22]. Performance is reported as classification rates (i.e., the ratio between the number of correctly classified images and the number of test images).

B. Tuning of the Parameters

To evaluate the goodness of the proposed descriptor we first studied how much the classification performance is affected by its parameters. In particular, we computed the classification rates obtained by varying the weight w that modulates the contribution of LBP and LCC features, and the number of bins Q that forms the LCC histogram.

Figure 4(a) shows the results obtained by combining LCC with LBPs computed on the luminance image (obtained as described in the OutTex Web site at the address

http://www.outex.oulu.fi/index.php?page=image_acquisition) with $Q = 256$ and varying the combination weight w. For $w = 0$ the plot reports the results obtained only by using LBPs without any contribution from LCC. Two classification rates are considered: one is the average of the three fixed-illuminant conditions (i/i, h/h and t/t), the other is the average of the six conditions with variable illuminant (i/h, i/t, h/i, h/t, t/i, t/h). In general, the addition of the color contrast feature to the LBP descriptor helps to better discriminate across texture classes. However, when the LCC contribution is too much the performances tend to decrease. It can also be noted that this behavior is more evident when the illuminant changes. Since we are particularly interested in the case of variable illumination, in the following we will consider the value $w = 0.15$ that corresponds to the best performance (77.1%) obtained in the case of variable illumination.

Figure 4(b) shows a similar analysis for the number of bins of the LCC histogram. The best average performance is achieved at $Q = 512$ with an accuracy of 85.6%, and at $Q = 128$ with an accuracy of 77.1% for the subsets without and with illuminant variation, respectively. Note that the performance obtained for values of Q ranging from 32 to 1024 is almost uniform. We also observed a similar behavior for other values of the combination factor w (not reported here). In the following we will consider $Q = 256$. This value is compatible with the number of bins of the most frequently used LBP configuration in these experiments ($LBP^{u}_{16,2}$), which is 243.

C. Color Variants of LBP Features

LBP features can be extended in several ways to take into account color information. The easiest one is to compute them on each color channel independently, and to concatenate the obtained histograms into a single feature vector [1]. This approach can be directly applied to the red, green, and blue channels of the input RGB image, or to the components of any other color space (we considered the RGB, HSV, CIE-Lab, and Ohta's $I_1I_2I_3$ color spaces, where the conversion to CIE-Lab as been obtained as described in [1]). Within this approach, it is quite possible that potentially useful information is lost, such as the information about co-occurrences of local patterns across different color channels.

Opponent color LBP (OCLBP) [17] tries to incorporate cross-channel information in the LBP framework: LBPs are extracted from each color channel independently, and then for pairs of color channels so that the center pixel is taken from one channel and the neighboring pixels from the other. In total, six histograms are computed: three for the R, G, and B channels, and three for the combinations R/G, R/B, and G/B. The histograms are concatenated to form the feature vector.

Connah and Finlayson [18] used 3D joint histograms of LBP values computed from the three R,G,B channels. In order to reduce the number of combinations of different patterns on the three channels, the patterns considered are restricted to the $LBP_{8,1}$ operator with rotation invariance and uniform patterns. Therefore, their final histogram has $10 \times 10 \times 10 = 1000$ components. We will refer to this variant as "Connah and Finlayson" (CF).

Table 1 reports the performance obtained by these variants of LBP features. It is clear that color information is very useful when training and test images are taken under the same illuminant. In this case, the OCLBP descriptor allows us to achieve a classification rate of about 93.3%, which is by far the best performance. However, changes in illumination make cross-channel information completely unreliable, to the point that OCLBP obtains only 10.2% of classification accuracy. For variable lighting conditions the best strategy seems to be to just ignore color information, and to compute the LBP histogram on the luminance image (71.9% of classification accuracy).

Table 2 shows the results obtained by concatenating the color variants of LBPs to the LCC histogram. When training and test illuminants are the same, the introduction of LCC increases the average performance of all the descriptors except OCLBP. The best performance with illuminant variation is obtained by the combination of LCC with LBPs computed on the luminance image with a classification accuracy of 77.1%, which is more than 5% better than the second best combination (71.9% for LBP-L, in Table 1).

Figure 5 summarizes the results obtained with illuminant variation by reporting for each method the minimum, average, and maximum classification rates obtained with the six combinations of training and test illuminants. The figure clearly shows how, with the exception of the CF method, the inclusion of LCC consistently improves the classification rate. The introduction of LCC also makes the results more robust in terms of worst-case performance. For instance, for LBP-L

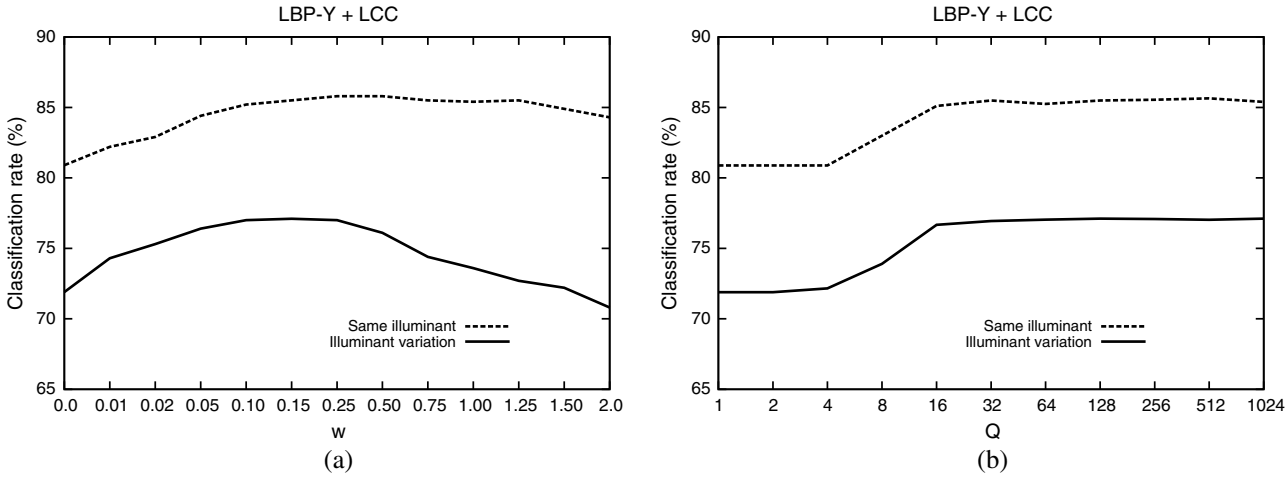

Fig. 4. Performance of the combination of $LBP^{u}_{16,2}$ and LCC obtained by varying (a) combination weight w and (b) number of bins Q of the LCC histogram. In the first plot Q was fixed to 256, and in the second w was fixed to 0.15.

Table 1. Evaluation of Several Variants of LBPs[a]

	Same Illuminant (%)				Illuminant Variation (%)						
Features	i/i	t/t	h/h	avg	i/t	i/h	t/i	t/h	h/i	h/t	avg
LBP-L	81.0	79.0	82.7	80.9	68.5	68.2	71.6	**77.9**	68.8	**76.2**	**71.9**
LBP-RGB	86.3	86.0	88.2	86.9	**71.3**	48.8	**72.4**	46.9	58.2	57.4	59.2
LBP-HSV	88.1	88.5	88.2	88.3	43.5	58.5	45.4	44.9	67.1	47.5	51.2
LBP-Lab	88.5	84.9	89.3	87.6	42.4	**77.9**	47.7	40.7	**77.5**	43.2	54.9
LBP-$I_1I_2I_3$	86.3	83.4	89.4	86.4	53.5	58.0	54.7	54.6	66.8	54.6	57.0
OCLBP	**93.1**	**96.2**	**90.7**	**93.3**	8.5	15.3	8.2	2.7	23.8	2.8	10.2
CF	73.8	75.7	73.1	74.2	42.8	25.7	42.1	19.7	32.5	28.4	31.9

[a]For each column, the best result is reported in bold.

Table 2. Evaluation of Several Variants of LBPs Combined with the LCC Descriptor[a]

	Same Illuminant (%)				Illuminant Variation (%)						
Features	i/i	t/t	h/h	avg	i/t	i/h	t/i	t/h	h/i	h/t	avg
LCC + LBP – L	85.3	85.2	86.2	85.5	76.8	75.4	**80.9**	**79.0**	74.7	**75.7**	**77.1**
LCC + LBP – RGB	88.8	87.8	90.3	89.0	**78.1**	59.9	79.6	58.4	64.6	61.9	67.1
LCC + LBP – HSV	89.3	**89.4**	**91.9**	**90.2**	52.7	63.5	57.1	54.0	73.8	56.5	59.6
LCC + LBP – Lab	**90.2**	87.4	91.5	89.7	51.0	**81.8**	56.2	51.3	**81.3**	50.0	62.0
LCC + LBP – $I_1I_2I_3$	87.1	85.2	91.6	87.9	60.4	66.8	65.3	57.5	75.4	61.6	64.5
LCC + OCLBP	35.6	39.4	36.0	37.0	18.4	18.2	18.7	11.5	19.7	9.6	16.0
LCC + CF	28.2	30.9	32.8	30.6	16.8	16.3	16.6	10.0	17.2	8.4	14.2

[a]For each column, the best result is reported in bold.

the worst combination is i/h, corresponding to a classification rate of 68.2%. The addition of LCC raised the worst-case classification rate up by 6.5% (74.7% for h/i). In most cases, the worst-case performance obtained with LCC is better than the average classification rate obtained without it.

D. Color Normalization Methods

An alternative to the use of illumination invariant features is to apply a color normalization method before feature extraction. In order to include this strategy in the comparison, we considered two variants of the Retinex method [26,27], the gray world [6], two variants of edge-based algorithm, the gray edge [8] and the weighted gray-edge method [9], and the comprehensive normalization [10].

The first Retinex method, denoted here as *McCann99*, is described in [26], and the second, denoted as *Frankle–McCann*, is described in [28]. Specifically, we used the implementations of both methods described in [27] that improve the computational efficiency while preserving the underlying principles.

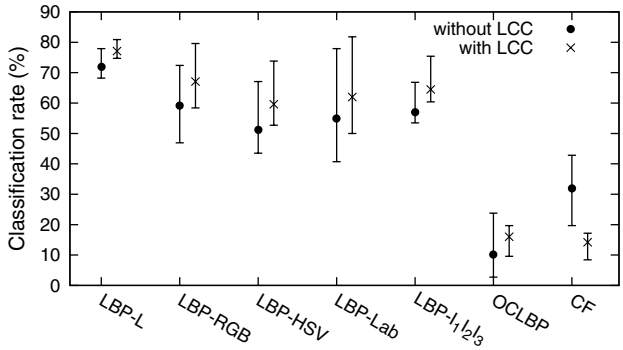

Fig. 5. Average classification rates obtained in the case of variable illumination. The bars indicate the minimum and maximum performance obtained on the six combinations of training and test illuminants.

The Retinex theory is based on the assumption that the perception of the color of an object is not determined by the spectral composition of the light stimulus but is determined by the result of the retina and visual cortex processing [29]. The algorithms considered compute the lightness at a given image pixel by comparing the pixel's value to that of other pixels. The strategy in which the pixel comparison is performed makes the main difference between these algorithms. The *McCann99* algorithm considers a multiresolution pyramid of the input image. Comparisons are initially calculated at a given scale, for instance, the top level of the pyramid, and then propagated down the other levels. The *Frankle–McCann* algorithm does not use the multiresolution approach and makes single pixel comparisons with variable separations. A single pixel eventually averages different products from all other pixels.

The gray-world algorithm is based on the assumption that the average value of the R, G, and B components of a given image should converge to a gray color given a sufficient amount of color variations [6].

Edge-based color constancy methods are based on the assumption that the average edge difference in a scene is achromatic [8]. Different edge types exist in real-world images, such as material, shadow, and highlight edges. The basic version of the gray-edge algorithm estimates the illuminant by using all types of image derivatives [8]. Due to the fact that different edge types may have a distinctive influence on the illuminant estimation, the weighted gray edge gives more importance to some types of edges, such as those corresponding to shadows and specularity [9].

The comprehensive normalization is an iterative procedure that removes image dependency on lighting geometry and illumination color. It is a very simple method that iterates two normalization steps until a stable state is reached. The RGB

Table 3. Average Accuracy of Several Variants of LBPs Combined with Different Preprocessing Methods Experimented on Subsets without Illuminant Variations[a]

	Same Illuminant (%)						
Preprocessing	LBP-L	LBP-RGB	LBP-HSV	LBP-Lab	LBP – $I_1I_2I_3$	OCLBP	CF
—	80.9	86.9	**88.3**	87.6	**86.4**	**93.3**	74.2
Retinex McCann	80.3	86.4	83.0	87.3	84.7	81.0	73.7
Retinex Frankle	79.2	86.2	82.0	87.7	85.1	80.0	72.7
Gray world	80.0	86.7	84.3	87.9	84.1	88.3	69.9
Gray edge	81.5	86.9	85.8	88.0	82.4	85.4	73.9
Weighted gray edge	**82.2**	86.9	86.2	**88.1**	82.3	83.9	74.4
Comprehensive norm.	42.7	73.1	70.0	76.0	67.2	78.7	48.7
CAT02	81.5	**88.5**	88.2	83.3	85.5	93.0	**79.0**

[a]For each column, the best result is reported in bold.

Table 4. Average Accuracy of Several Variants of LBPs Combined with Different Preprocessing Methods Experimented on Subsets with Illuminant Variations[a]

	Illuminant Variation (%)						
Preprocessing	LBP-L	LBP-RGB	LBP-HSV	LBP-Lab	LBP-$I_1I_2I_3$	OCLBP	CF
—	71.9	59.2	51.2	54.9	57.0	10.2	31.7
Retinex McCann	69.6	57.2	46.9	54.4	54.9	8.8	29.4
Retinex Frankle	70.4	57.2	47.5	54.5	54.8	9.0	29.9
Gray world	69.5	63.3	**55.8**	62.0	58.3	**51.2**	**35.5**
Gray edge	69.8	59.2	39.7	64.7	57.1	23.5	31.9
Weighted gray edge	69.0	59.2	38.5	**64.9**	56.9	22.5	31.9
Comprehensive norm.	22.8	35.1	32.0	37.1	29.7	33.6	18.3
CAT02	**72.6**	**63.6**	43.3	64.6	**61.9**	28.3	34.0

[a]For each column, the best result is reported in bold.

pixel normalization removes dependence on lighting geometry, while the R, G, and B channel normalization removes dependence on the illuminant color.

Finally, we also included in the comparison an ideal correction of the illuminant obtained by exploiting the knowledge of the acquisition device and of the illuminants under which the pictures have been taken. To do so, we used a chromatic adaptation transform based on the Von Kries model [30]. The use of a CAT is possible since the authors of the data set made available the color conversion matrix from RGB camera values to CIE XYZ tristimulus values, together with the spectral power distributions of the three illuminants. The CAT chosen is the CAT02 [31], but different choices could be made [32,33]. The equal energy CIE E standard illuminant has been chosen as the destination inside the CAT.

Tables 3 and 4 report the performance obtained by these color normalization algorithms combined with several variants of LBP features and evaluated on the subsets without and with illuminant variations. Although the use of color normalization gives, in most of the cases, a slight improvement of performance, for none of the LBP variants is such an improvement larger than that obtained with the combination with the proposed LCC histogram.

E. Comparison with Other Methods Robust to Illuminant Conditions

Table 5 reports a comparison between the best performance obtained in the previous experiments and some methods in the state of the art. Precisely, we considered the OCLBP that achieved the maximum performance on the i/i subset (namely the Outex-13), and the proposed descriptor (LBP-L combined with the LCC) that achieved the maximum average on the Outex-14.

Furthermore, we considered three other methods in the state of the art. The first is a combination of three variants of LBPs, that is, $LBP^{u}_{8,1}$, $LBP^{u}_{16,3}$, and $LBP^{u}_{24,5}$, applied to the Lab color space [1]. The last two methods use the $Gabor_{4,6}$ features after an image color normalization. More precisely, the method described in [20] was applied to both Outex 13 and 14, and features were calculated on the luminance and RGB color space. The method described in [19] has been applied only to Outex-14.

On Outex-14, our descriptor outperformed all the other alternatives considered. On Outex-13 (i.e., without illuminant variations) OCLBP obtained the best performance.

Kandaswamy *et al.* also reported the results obtained by applying their method to an *ad hoc* color space obtained by the combination of the luminance, the red and green channels

Table 5. Comparison of Best LBPs with Best LBPs + LCC and Other Methods in the State of the Art[a]

	Outex-13		Outex-14	
Method	i/i	i/t	i/h	avg
OCLBP [17]	**93.1**	8.5	15.3	12.0
$LBP^{u}_{16,2}$ – L + ICCD [15]	89.0	—	—	75.6
$LBP^{u}_{8,1} + LBP^{u}_{16,3} + LBP^{u}_{24,5}$ on Lab [1]	87.8	—	—	67.8
Khan *et al.* on L [20]	—	69.2	73.3	71.3
Khan *et al.* on Colors [20]	84.8	74.4	53.6	64.0
Kandaswamy *et al.* on L [19]	—	73.5	**75.9**	74.7
$LBP^{u}_{16,2}$ – L + LCC (proposed)	85.3	**76.8**	75.4	**76.1**

[a]For each column, the best result is reported in bold.

(denoted as RGL). This choice is not general, and depends on the particular characteristics of the Outex-14 data set. By exploiting the RGL color space, they achieved a significantly higher classification accuracy, 76.3% on i/t and 77.8% on i/h, obtaining an average accuracy of 77.1%. We decided to consider the RGL color space as well, for the computation of the LCC feature. The results we obtained confirm that RGL is indeed a good color space for the classification of Outex images. More in detail, we obtained a classification accuracy of 76.9% on i/t and 79.1% on i/h, for an average of 78.0%.

4. DISCUSSION AND CONCLUSIONS

On the basis of their results on the Outex-14 data set, Mäenpää and Pietikäinen concluded their seminal work with the observation that, under varying illumination conditions, gray-scale texture features work clearly better than color-based features [1]. In this paper we demonstrated that color information can be processed in such a way that the resulting descriptor clearly outperforms intensity-based approaches on the same Outex-14 test suite.

More in detail, we proposed a descriptor that is formed by combining a histogram of LBPs with a feature encoding the distribution of a measure of the LCC. An extensive evaluation on the Outex data set showed that its performance is superior to that of the other gray-scale and color descriptors in the state of the art, especially in the case of variable illumination conditions. We also considered the usage of variants of LBP features after the application of color normalization methods. None of the combinations of features and normalization methods achieved good results in the case of variable illumination, even in the case in which *a priori* information about the illuminant is exploited.

More investigation should be devoted to the issue of the proper combination of the two parts of the proposed descriptor. In fact, the optimal value of the parameter w depends on the variability of the illumination condition. We plan to study how to measure such a variability and how to adaptively determine the best value for w. To do so a new data set is required, since none of the existing ones provide images taken under a large number of controlled lighting conditions [34].

The descriptor proposed here has been mainly designed to be robust with respect to variations of the color of the illuminant. However, we speculate that even bigger challenges may be posed by variations of the illuminant position and by the possible presence of multiple illuminants. In fact, in a preliminary work we verified that when only the illuminant's color is allowed to change there are simple descriptors that are pretty much invariant [35]. In the future, we plan to further investigate all the sources of variation in the appearance of color textures.

†The authors contributed equally to this work.

REFERENCES

1. T. Mäenpää and M. Pietikäinen, "Classification with color and texture: jointly or separately?" Pattern Recogn. **37**, 1629–1640 (2004).
2. A. Poirson and B. Wandell, "Pattern-color separable pathways predict sensitivity to simple colored patterns," Vis. Res. **36**, 515–526 (1996).
3. F. Bianconi, R. Harvey, P. Southam, and A. Fernández, "Theoretical and experimental comparison of different approaches for color texture classification," J. Electron. Imaging **20**, 043006 (2011).
4. O. Drbohlav and A. Leonardis, "Towards correct and informative evaluation methodology for texture classification under varying viewpoint and illumination," Comput. Vis. Image Underst. **114**, 439–449 (2010).
5. U. Kandaswamy, S. A. Schuckers, and D. Adjeroh, "Comparison of texture analysis schemes under nonideal conditions," IEEE Trans. Image Process. **20**, 2260–2275 (2011).
6. G. Finlayson and E. Trezzi, "Shades of gray and colour constancy," in *Color and Imaging Conference* (Society for Imaging Science and Technology, 2004), pp. 37–41.
7. E. Land and J. McCann, "Lightness and retinex theory," J. Opt. Soc. Am. **61**, 1–11 (1971).
8. J. van de Weijer, T. Gevers, and A. Gijsenij, "Edge-based color constancy," IEEE Trans. Image Process. **16**, 2207–2214 (2007).
9. A. Gijsenij, T. Gevers, and J. van de Weijer, "Improving color constancy by photometric edge weighting," IEEE Trans. Pattern Anal. Mach. Intell. **34**, 918–929 (2012).
10. G. Finlayson, B. Schiele, and J. Crowley, "Comprehensive colour image normalization," in *European Conference on Computer Vision* (Springer, 1998), pp. 475–490.
11. R. Khan, D. Muselet, and A. Trémeau, "Classical texture features and illumination color variations," in *Proceedings of Third International Conference on Machine Vision* (IEEE, 2010), pp. 280–285.
12. G. Finlayson, S. Hordley, G. Schaefer, and G. Yun Tian, "Illuminant and device invariant colour using histogram equalisation," Pattern Recogn. **38**, 179–190 (2005).
13. M. Seifi, X. Song, D. Muselet, and A. Tremeau, "Color texture classification across illumination changes," in *Conference on Colour in Graphics, Imaging, and Vision* (Society for Imaging Science and Technology, 2010), pp. 332–337.
14. T. Ojala, M. Pietikäinen, and T. Mänepää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," IEEE Trans. Pattern Anal. Mach. Intell. **24**, 971–987 (2002).
15. C. Cusano, P. Napoletano, and R. Schettini, "Illuminant invariant descriptors for color texture classification," in *Computational Color Imaging*, Vol. **7786** of Lecture Notes in Computer Science (Springer, 2013), pp. 239–249.
16. C. Cusano, P. Napoletano, and R. Schettini, "Intensity and color descriptors for texture classification," Proc. SPIE **8661**, 866113 (2013).
17. C.-H. Chan, J. Kittler, and K. Messer, "Multispectral local binary pattern histogram for component-based color face verification," in *First IEEE International Conference on Biometrics: Theory, Applications, and Systems* (IEEE, 2007), pp. 1–7.
18. D. Connah and G. Finlayson, "Using local binary pattern operators for colour constant image indexing," in *Proceedings of European Conference on Color in Graphics, Imaging, and Vision* (Society for Imaging Science and Technology, 2006), pp. 60–64.
19. U. Kandaswamy, D. Adjeroh, S. Schuckers, and A. Hanbury, "Robust color texture features under varying illumination conditions," IEEE Trans. Syst. Man Cyber. Part B **42**, 58–68 (2012).
20. R. Khan, D. Muselet, and A. Trémeau, "Texture classification across illumination color variations," Int. J. Comput. Theory Eng. **5**, 65–70 (2013).
21. T. Ojala, T. Mäenpää, M. Pietikäinen, J. Viertola, J. Kyllönen, and S. Huovinen, "Outex-new framework for empirical evaluation of texture analysis algorithms," in *16th International Conference on Pattern Recognition*, Vol. **1** (IEEE, 2002), pp. 701–706.
22. M. Pietikäinen, A. Hadid, G. Zhao, and T. Ahonen, "Local binary patterns for still images," in *Computer Vision Using Local Binary Patterns*, Vol. **40** of Computational Imaging and Vision (Springer, 2011), pp. 13–47.
23. M. Drew, D. Connah, G. Finlayson, and M. Bloj, "Improved colour to greyscale via integrability correction," in *IS&T/SPIE Electronic Imaging* (International Society for Optics and Photonics, 2009), p. 72401B.
24. A. Alsam and M. Drew, "Fast multispectral2gray," J. Imaging Sci. Technol. **53**, 60401 (2009).
25. F. Khan, J. van de Weijer, S. Ali, and M. Felsberg, "Evaluating the impact of color on texture recognition," in *Proceedings of International Conference on Computer Analysis of Images and Patterns* (Springer, 2013), pp. 154–162.

26. J. McCann, "Lessons learned from mondrians applied to real images and color gamuts," in *Color and Imaging Conference* (Society for Imaging Science and Technology, 1999), pp. 1–8.
27. B. Funt, F. Ciurea, and J. McCann, "Retinex in MATLAB," J. Electron. Imaging **13**, 48–57 (2004).
28. J. Frankle and J. McCann, "Method and apparatus for lightness imaging," U.S. patent 4,384,336 (May 17, 1983).
29. E. Land, "Recent advances in retinex theory," Vis. Res. **26**, 7–21 (1986).
30. J. von Kries, "Chromatic adaptation," [originally published in Festschrift der Albrecht-Ludwigs-Universität (Fribourg, Germany, 1902), pp. 145–148], in *Sources of Color Vision*, L. D. MacAdam, ed. (MIT, 1970), pp. 109–126.
31. N. Moroney, M. Fairchild, R. Hunt, C. Li, M. Luo, and T. Newman, "The CIECAM02 color appearance model," in *Color and Imaging Conference* (Society for Imaging Science and Technology, 2002), pp. 23–27.
32. M. Luo, "A review of chromatic adaptation transforms," Rev. Progr. Coloration Rel. Top. **30**, 77–92 (2000).
33. S. Bianco and R. Schettini, "Two new Von Kries based chromatic adaptation transforms found by numerical optimization," Color Res. Appl. **35**, 184–192 (2010).
34. S. Hossain and S. Serikawa, "Texture databases—a comprehensive survey," Pattern Recogn. Lett. **34**, 2007–2022 (2013).
35. S. Bianco, C. Cusano, P. Napoletano, and R. Schettini, "On the robustness of color texture descriptors across illuminants," in *17th International Conference on Image Analysis and Applications (ICIAP)*, Vol. **8157** of Lecture Notes in Computer Science (Springer, 2013), pp. 652–662.