

Gray Scale and Rotation Invariant Texture Classification with Local Binary Patterns

Timo Ojala, Matti Pietikäinen and Topi Mäenpää

Machine Vision and Media Processing Unit

Infotech Oulu, University of Oulu

P.O.Box 4500, FIN - 90014 University of Oulu, Finland

{skidi, mkp, topiollli}@ee.oulu.fi

<http://www.ee.oulu.fi/research/imag/texture>

Abstract. This paper presents a theoretically very simple yet efficient approach for gray scale and rotation invariant texture classification based on local binary patterns and nonparametric discrimination of sample and prototype distributions. The proposed approach is very robust in terms of gray scale variations, since the operators are by definition invariant against any monotonic transformation of the gray scale. Another advantage is computational simplicity, as the operators can be realized with a few operations in a small neighborhood and a lookup table. Excellent experimental results obtained in two true problems of rotation invariance, where the classifier is trained at one particular rotation angle and tested with samples from other rotation angles, demonstrate that good discrimination can be achieved with the statistics of simple rotation invariant local binary patterns. These operators characterize the spatial configuration of local image texture and the performance can be further improved by combining them with rotation invariant variance measures that characterize the contrast of local image texture. The joint distributions of these orthogonal measures are shown to be very powerful tools for rotation invariant texture analysis.

1 Introduction

Real world textures can occur at arbitrary rotations and they may be subjected to varying illumination conditions. This has inspired few studies on gray scale and rotation invariant texture analysis, which presented methods for incorporating both types of invariance [2,14]. A larger number of papers have been published on plain rotation invariant analysis, among others [4,5,6,7,8,12], while [3] proposed an approach to encompass invariance with respect to another important property, spatial scale, in conjunction with rotation invariance.

Both Chen and Kundu [2] and Wu and Wei [14] approached gray scale invariance by assuming that the gray scale transformation is a linear function. This is a somewhat strong simplification, which may limit the usefulness of the proposed methods. Chen and Kundu realized gray scale invariance by global normalization of the input image using histogram equalization. This is not a general solution, however, as global histogram equalization can not correct intrimage (local) gray scale variations. Another problem of many approaches to rotation invariant texture analysis is their computational complexity (e.g. [3]), which may render them impractical.

In this study we propose a theoretically and computationally simple approach which is robust in terms of gray scale variations and which is shown to discriminate

rotated textures efficiently. Extending our earlier work [9,10,11], we present a truly gray scale and rotation invariant texture operator based on local binary patterns. Starting from the joint distribution of gray values of a circularly symmetric neighbor set of eight pixels in a 3x3 neighborhood, we derive an operator that is by definition invariant against any monotonic transformation of the gray scale. Rotation invariance is achieved by recognizing that this gray scale invariant operator incorporates a fixed set of rotation invariant patterns.

The novel contribution of this work is to use only a limited subset of ‘uniform’ patterns instead of all rotation invariant patterns, which improves the rotation invariance considerably. We call this operator LBP_8^{riu2} . The use of only ‘uniform’ patterns is motivated by the reasoning that they tolerate rotation better because they contain fewer spatial transitions exposed to unwanted changes upon rotation. This approximation is also supported by the fact that these ‘uniform’ patterns tend to dominate in deterministic textures, which is demonstrated using a sample image data. Further, we propose operator called LBP_{16}^{riu2} , which enhances the angular resolution of LBP_8^{riu2} by considering a circularly symmetric set of 16 pixels in a 5x5 neighborhood.

These operators are excellent measures of the spatial structure of local image texture, but they by definition discard the other important property of local image texture, contrast, since it depends on the gray scale. We characterize contrast with rotation invariant variance measures named VAR_8 and VAR_{16} , corresponding to the circularly symmetric neighbor set where they are computed. We present the joint distributions of these complementary measures as powerful tools for rotation invariant texture classification. As the classification rule we employ nonparametric discrimination of sample and prototype distributions based on a log-likelihood measure of the (dis)similarity of histograms.

The performance of the proposed approach is demonstrated with two problems used in recent studies on rotation invariant texture classification [4,12]. In addition to the original experimental setups we also consider more challenging cases, where the texture classifier is trained at one particular rotation angle and then tested with samples from other rotation angles. Excellent experimental results demonstrate that the texture representation obtained at a specific rotation angle generalizes to other rotation angles. The proposed operators are also computationally attractive, as they can be realized with a few operations in a small neighborhood and a lookup table.

The paper is organized as follows. The derivation of the operators and the classification principle are described in Section 2. Experimental results are presented in Section 3 and Section 4 concludes the paper.

2 Gray Scale and Rotation Invariant Local Binary Patterns

We start the derivation of our gray scale and rotation invariant texture operator by defining texture T in a local 3x3 neighborhood of a monochrome texture image as the joint distribution of the gray levels of the nine image pixels:

$$T = p(g_0, g_1, g_2, g_3, g_4, g_5, g_6, g_7, g_8) \quad (1)$$

where g_i ($i=0,\dots,8$), correspond to the gray values of the pixels in the 3×3 neighborhood according to the spatial layout illustrated in Fig. 1. The gray values of diagonal pixels ($g_2, g_4, g_6,$ and g_8) are determined by interpolation.




Fig. 1. The circularly symmetric neighbor set of eight pixels in a 3×3 neighborhood.

2.1 Achieving Gray Scale Invariance

As the first step towards gray scale invariance we subtract, without losing information, the gray value of the center pixel (g_0) from the gray values of the eight surrounding pixels of the circularly symmetric neighborhood ($g_i, i=1,\dots,8$) giving:

$$T = p(g_0, g_1 - g_0, g_2 - g_0, g_3 - g_0, g_4 - g_0, g_5 - g_0, g_6 - g_0, g_7 - g_0, g_8 - g_0) \quad (2)$$

Next, we assume that differences $g_i - g_0$ are independent of g_0 , which allows us to factorize Eq.(2):

$$T = p(g_0) p(g_1 - g_0, g_2 - g_0, g_3 - g_0, g_4 - g_0, g_5 - g_0, g_6 - g_0, g_7 - g_0, g_8 - g_0) \quad (3)$$

In practice an exact independence is not warranted, hence the factorized distribution is only an approximation of the joint distribution. However, we are willing to accept the possible small loss in information, as it allows us to achieve invariance with respect to shifts in gray scale. Namely, the distribution $p(g_0)$ in Eq.(3) describes the overall luminance of the image, which is unrelated to local image texture, and consequently does not provide useful information for texture analysis. Hence, much of the information in the original joint gray level distribution (Eq.(1)) about the textural characteristics is conveyed by the joint difference distribution [10]:

$$T = p(g_1 - g_0, g_2 - g_0, g_3 - g_0, g_4 - g_0, g_5 - g_0, g_6 - g_0, g_7 - g_0, g_8 - g_0) \quad (4)$$

Signed differences $g_i - g_0$ are not affected by changes in mean luminance, hence the joint difference distribution is invariant against gray scale shifts. We achieve invariance with respect to the scaling of the gray scale by considering just the signs of the differences instead of their exact values:

$$T = p(s(g_1 - g_0), s(g_2 - g_0), s(g_3 - g_0), s(g_4 - g_0), \dots, s(g_8 - g_0)) \quad (5)$$

where

$$s(x) = \begin{cases} 1, & x \geq 0 \\ 0, & x < 0 \end{cases} \tag{6}$$

If we formulate Eq.(5) slightly differently, we obtain an expression similar to the LBP (Local Binary Pattern) operator we proposed in [9]:

$$LBP_8 = \sum_{i=1}^8 s(g_i - g_0) 2^{i-1} \tag{7}$$

The two differences between LBP_8 and the LBP operator of [9] are: 1) the pixels in the neighbor set are indexed so that they form a circular chain, and 2) the gray values of the diagonal pixels are determined by interpolation. Both modifications are necessary to obtain the circularly symmetric neighbor set, which allows for deriving a rotation invariant version of LBP_8 . For notational reasons we augment LBP with subscript 8 to denote that the LBP_8 operator is determined from the 8 pixels in a 3x3 neighborhood. The name ‘Local Binary Pattern’ reflects the nature of the operator, i.e. a local neighborhood is thresholded at the gray value of the center pixel into a binary pattern. LBP_8 operator is by definition invariant against any monotonic transformation of the gray scale, i.e. as long as the order of the gray values stays the same, the output of the LBP_8 operator remains constant.

2.2 Achieving Rotation Invariance

The LBP_8 operator produces 256 (2^8) different output values, corresponding to the 256 different binary patterns that can be formed by the eight pixels in the neighbor set. When the image is rotated, the gray values g_i will correspondingly move along the perimeter of the circle around g_0 . Since we always assign g_1 to be the gray value of element (0,1), to the right of g_0 , rotating a particular binary pattern naturally results in a different LBP_8 value. This does not apply to patterns 00000000_2 and 11111111_2 which remain constant at all rotation angles. To remove the effect of rotation, i.e. to assign a unique identifier to each rotation invariant local binary pattern we define:

$$LBP_8^{ri36} = \min\{ROR(LBP_8, i) \mid i = 0,1,\dots,7\} \tag{8}$$

where $ROR(x, i)$ performs a circular bit-wise right shift on the 8-bit number x i times. In terms of image pixels Eq.(8) simply corresponds to rotating the neighbor set clockwise so many times that a maximal number of the most significant bits, starting from g_8 , are 0. We observe that LBP_8^{ri36} can have 36 different values, corresponding to the 36 unique rotation invariant local binary patterns illustrated in Fig. 2, hence the superscript $ri36$. LBP_8^{ri36} quantifies the occurrence statistics of these patterns corresponding

to certain microfeatures in the image, hence the patterns can be considered as feature detectors. For example, pattern #0 detects bright spots, #8 dark spots and flat areas, and #4 edges. Hence, we have obtained the gray scale and rotation invariant operator LBP_8^{ri36} that we designated as LBPROT in [11].




Fig. 2. The 36 unique rotation invariant binary patterns that can occur in the eight pixel circularly symmetric neighbor set. Black and white circles correspond to bit values of 0 and 1 in the 8-bit output of the LBP_8 operator. The first row contains the nine ‘uniform’ patterns, and the numbers inside them correspond to their unique LBP_8^{riu2} values.

2.3 Improved Rotation Invariance with ‘Uniform’ Patterns

However, our practical experience has showed that LBP_8^{ri36} as such does not provide a very good discrimination, as we also concluded in [11]. There are two reasons:

1) the performance of the 36 individual patterns in discrimination of rotated textures varies greatly: while some patterns sustain rotation quite well, other patterns do not and only confuse the analysis. Consequently, using all 36 patterns leads to a suboptimal result (addressed in this section).

2) crude quantization of the angular space at 45° intervals (addressed in Section 2.4).

The varying performance of individual patterns attributes to the spatial structure of the patterns. To quantify this we define an uniformity measure $U(\text{‘pattern’})$, which corresponds to the number of spatial transitions (bitwise 0/1 changes) in the ‘pattern’. For example, patterns 00000000_2 and 11111111_2 have U value of 0, while the other seven patterns in the first row of Fig. 2 have U value of 2, as there are exactly two 0/1 transitions in the pattern. Similarly, other 27 patterns have U value of at least 4.

We argue that the larger the uniformity value U of a pattern is, i.e. the larger number of spatial transitions occurs in the pattern, the more likely the pattern is to change to a different pattern upon rotation in digital domain. Based on this argument we designate patterns that have U value of at most 2 as ‘uniform’ and propose the following

operator for gray scale and rotation invariant texture description instead of LBP_8^{ri36} :

$$LBP_8^{riu2} = \begin{matrix} 8 \\ s(g_i - g_0) & \text{if } U(LBP_8) \leq 2 \\ i = 1 & \\ 9 & \text{otherwise} \end{matrix} \quad (9)$$

Eq.(9) corresponds to giving an unique label to the nine ‘uniform’ patterns illustrated in the first row of Fig. 2 (label corresponds to the number of ‘1’ bits in the pattern), the 27 other patterns being grouped under the ‘miscellaneous’ label (9). Superscript $riu2$ corresponds to the use of *rotation invariant ‘uniform’* patterns that have U value of at most 2.

The selection of ‘uniform’ patterns with the simultaneous compression of ‘nonuniform’ patterns is also supported by the fact that the former tend to dominate in deterministic textures. This is studied in more detail in Section 3 using the image data of the experiments. In practice the mapping from LBP_8 to LBP_8^{riu2} , which has 10 distinct output values, is best implemented with a lookup table of 256 elements.

2.4 Improved Angular Resolution with a 16 Pixel Neighborhood

We noted earlier that the rotation invariance of LBP_8^{riu2} is hampered by the crude 45° quantization of the angular space provided by the neighbor set of eight pixels. To address this we present a modification, where the angular space is quantized at a finer resolution of 22.5° intervals. This is accomplished with the circularly symmetric neighbor set of 16 pixels illustrated in Fig. 3. Again, the gray values of neighbors which do not fall exactly in the center of pixels are estimated by interpolation. Note that we increase the size of the local neighborhood to 5x5 pixels, as the eight added neighbors would not provide too much new information if inserted into the 3x3 neighborhood. An additional advantage is the different spatial resolution, if we should want to perform multiresolution analysis.




Fig. 3. The circularly symmetric neighbor set of 16 pixels in a 5x5 neighborhood.

Following the derivation of LBP_8 , we first define the 16-bit version of the rotation variant LBP:

$$LBP_{16} = \sum_{i=1}^{16} s(h_i - h_0) 2^{i-1} \tag{10}$$

The LBP_{16} operator has 65536 (2^{16}) different output values and 243 different rotation invariant patterns can occur in the circularly symmetric set of 16 pixels. Choosing again the ‘uniform’ rotation invariant patterns that have at most two 0/1 transitions, we define LBP_{16}^{riu2} , the 16-bit version of LBP_8^{riu2} :

$$LBP_{16}^{riu2} = \begin{cases} \sum_{i=1}^{16} s(h_i - h_0) & \text{if } U(LBP_{16}) \leq 2 \\ 17 & \text{otherwise} \end{cases} \tag{11}$$

Thus, the LBP_{16}^{riu2} operator has 18 distinct output values, of which values from 0 (pattern 0000000000000000₂) to 16 (pattern 1111111111111111₂) correspond to the number of 1 bits in the 17 unique ‘uniform’ rotation invariant patterns, and value 17 denotes the ‘miscellaneous’ class of all ‘nonuniform’ patterns. In practice the mapping from LBP_{16} to LBP_{16}^{riu2} is implemented with a lookup table of 2^{16} elements.

2.5 Rotation Invariant Variance Measures of the Contrast of Local Image Texture

Generally, image texture is regarded as a two dimensional phenomenon that can be characterized with two orthogonal properties, spatial structure (pattern) and contrast (the ‘amount’ of local image texture). In terms of gray scale and rotation invariant texture description these two are an interesting pair: whereas spatial pattern is affected by rotation, contrast is not, and *vice versa*, whereas contrast is affected by the gray scale, spatial pattern is not. Consequently, as long as we want to restrict ourselves to pure gray scale invariant texture analysis, contrast is of no interest, as it depends on the gray scale.

The LBP_8^{riu2} and LBP_{16}^{riu2} operators are true gray scale invariant measures, i.e. their output is not affected by any monotonic transformation of the gray scale. They are excellent measures of the spatial pattern, but by definition discard contrast. If we under stable lighting conditions wanted to incorporate the contrast of local image texture as well, we can measure it with rotation invariant measures of local variance:

$$VAR_8 = \frac{1}{8} \sum_{i=1}^8 (g_i - \bar{g})^2, \text{ where } \bar{g} = \frac{1}{8} \sum_{i=1}^8 g_i \tag{12}$$

$$VAR_{16} = \frac{1}{16} \sum_{i=1}^{16} (h_i - \bar{h})^2, \text{ where } \bar{h} = \frac{1}{16} \sum_{i=1}^{16} h_i \tag{13}$$

VAR_8 and VAR_{16} are by definition invariant against shifts in gray scale. Since LBP and VAR are complementary, their joint distributions $\text{LBP}_8^{\text{riu}2}/\text{VAR}_8$ and $\text{LBP}_{16}^{\text{riu}2}/\text{VAR}_{16}$ are very powerful rotation invariant measures of local image texture.

2.6 Nonparametric Classification Principle

In the classification phase a test sample S was assigned to the class of the model M that maximized the log-likelihood measure:

$$L(S, M) = \prod_{b=1}^B S_b \log M_b \quad (14)$$

where B is the number of bins, and S_b and M_b correspond to the sample and model probabilities at bin b , respectively. This nonparametric (pseudo-)metric measures likelihoods that samples are from alternative texture classes, based on exact probabilities of feature values of pre-classified texture prototypes. In the case of the joint distributions $\text{LBP}_8^{\text{riu}2}/\text{VAR}_8$ and $\text{LBP}_{16}^{\text{riu}2}/\text{VAR}_{16}$, the log-likelihood measure (Eq.(14)) was extended in a straightforward manner to scan through the two-dimensional histograms.

Sample and model distributions were obtained by scanning the texture samples and prototypes with the chosen operator, and dividing the distributions of operator outputs into histograms having a fixed number of B bins. Since $\text{LBP}_8^{\text{riu}2}$ and $\text{LBP}_{16}^{\text{riu}2}$ have a completely defined set of discrete output values, they do not require any additional binning procedure, but the operator outputs are directly accumulated into a histogram of 10 ($\text{LBP}_8^{\text{riu}2}$) or 18 ($\text{LBP}_{16}^{\text{riu}2}$) bins.

Variance measures VAR_8 and VAR_{16} have a continuous-valued output, hence quantization of their feature space is required. This was done by adding together feature distributions for every single model image in a total distribution, which was divided into B bins having an equal number of entries. Hence, the cut values of the bins of the histograms corresponded to the $(100/B)$ percentile of the combined data. Deriving the cut values from the total distribution and allocating every bin the same amount of the combined data guarantees that the highest resolution of quantization is used where the number of entries is largest and *vice versa*. The number of bins used in the quantization of the feature space is of some importance, as histograms with a too modest number of bins fail to provide enough discriminative information about the distributions. On the other hand, since the distributions have a finite number of entries, a too large number of bins may lead to sparse and unstable histograms. As a rule of thumb, statistics literature often proposes that an average number of 10 entries per bin should be sufficient. In the experiments we set the value of B so that this condition was satisfied.

3 Experiments

We demonstrate the performance of our operators with two different texture image data that have been used in recent studies on rotation invariant texture classification [4,12]. In both cases we first replicate the original experimental setup as carefully as

possible, to get comparable results. Since the training data included samples from several rotation angles, we also present results for a more challenging setup, where the samples of just one particular rotation angle are used for training the texture classifier, which is then tested with the samples of the other rotation angles.

However, we first report classification results for the problem that we used in our recent study on rotation invariant texture analysis [11]. There we achieved an error rate of 39.2% with the LBP_8^{ri36} (LBPROT) operator when using 64x64 samples, while LBP_8^{riu2} and LBP_{16}^{riu2} operators provide error rates of 25.5% and 8.0%, respectively. These improvements underline the benefits of using ‘uniform’ patterns and finer quantization of the angular space.

Before going into the experiments we use the image data to take a quick look at the statistical foundation of LBP_8^{riu2} and LBP_{16}^{riu2} . In the case of LBP_8^{riu2} we choose nine ‘uniform’ patterns out of the 36 possible patterns, merging the remaining 27 under the ‘miscellaneous’ label. Similarly, in the case of LBP_{16}^{riu2} we consider only 7% (17 out of 243) of the possible rotation invariant patterns. Taking into account a minority of the possible patterns, and merging a majority of them, could imply that we are throwing away most of the pattern information. However, this is not the case, as the ‘uniform’ patterns tend to be the dominant structure.

For example, in the case of the image data of Experiment #2, the nine ‘uniform’ patterns of LBP_8^{riu2} contribute from 88% up to 94% of the total pattern data, averaging 90.9%. The most frequent individual pattern is symmetric edge detector 00001111_2 with about 25% share, followed by 00000111_2 and 00011111_2 with about 15% each. As expected, in the case of LBP_{16}^{riu2} the 17 ‘uniform’ patterns contribute a smaller proportion of the image data, from 70% up to 84% of the total pattern data, averaging 76.3%. The most frequent pattern is again symmetric edge detector 0000000011111111_2 with about 9.3% share.

3.1 Experiment #1

In their comprehensive study Porter and Canagarajah [12] presented three feature extraction schemes for rotation invariant texture classification, employing the wavelet transform, a circularly symmetric Gabor filter and a Gaussian Markov Random Field with a circularly symmetric neighbor set. They concluded that the wavelet-based approach was the most accurate and exhibited the best noise performance, having also the lowest computational complexity.

Image Data and Experimental Setup. Image data included 16 texture classes from the Brodatz album [1] shown in Fig. 4. For each texture class there were eight 256x256 images, of which the first was used for training the classifier, while the other seven images were used to test the classifier. Rotated textures were created from these source images using a proprietary interpolation program that produced images of 180x180 pixels in size. If the rotation angle was a multiple of 90 degrees (0° or 90° in the case of present ten rotation angles), a small amount of artificial blur was added to the original images to simulate the effect of blurring on rotation at other angles.

In the original experimental setup the texture classifier was trained with several 16x16 subimages extracted from the training image. This fairly small size of training samples increases the difficulty of the problem nicely. The training set comprised rotation angles 0° , 30° , 45° , and 60° , while the textures for classification were presented at rotation angles 20° , 70° , 90° , 120° , 135° , and 150° . Consequently, the test data included 672 samples, 42 (6 angles x 7 images) for each of the 16 texture classes. Using a Mahalanobis distance classifier Porter and Canagarajah reported 95.8% accuracy for the rotation invariant wavelet-based features as the best result.




Fig. 4. Texture images of Experiment #1 printed at particular orientations. Textures were presented at ten different angles: 0° , 20° , 30° , 45° , 60° , 70° , 90° , 120° , 135° , and 150° . Images are 180x180 pixels in size.

Experimental Results. We started replicating the original experimental setup by dividing the 180x180 images of the four training angles (0° , 30° , 45° , and 60°) into

121 disjoint 16x16 subimages. In other words we had 7744 training samples, 484 (4 angles x 121 samples) in each of the 16 texture classes. We first computed the histogram of the chosen operator for each of the 16x16 samples. We then added the histograms of all samples belonging to a particular class into one big model histogram for this class, since the histograms of single 16x16 samples would be too sparse to be reliable models. Also, using 7744 different models would result in computational overhead, for in the classification phase the sample histograms are compared to every model histogram. Consequently, we obtained 16 reliable model histograms containing 108900 (LBP_8^{riu2} and VAR_8 with a 1 pixel border produce 15^2 entries for a 16x16 sample) or 94864 (LBP_{16}^{riu2} and VAR_{16} have a 2 pixel border) entries.

The performance of the operators was evaluated with the 672 testing images. The sample histogram contained 32041/31684 entries, hence we did not have to worry about their stability. Classification results (the percentage of misclassified samples from all classified samples) for the four individual operators and the two joint distributions are given in Table 1.

Table 1: Error rates (%) for the original experimental setup, where training is done with rotations 0° , 30° , 45° , and 60° .

OPERATOR	BINS	ERROR	OPERATOR	BINS	ERROR
LBP_8^{riu2}	10	11.76	VAR_8	128	4.46
LBP_{16}^{riu2}	18	1.49	VAR_{16}	128	11.61
LBP_8^{riu2}/VAR_8	10/16	1.64	LBP_{16}^{riu2}/VAR_{16}	18/16	0.15

As expected, LBP_{16}^{riu2} clearly outperforms its 8-bit version LBP_8^{riu2} . LBP_8^{riu2} has difficulties in discriminating strongly oriented textures *straw* (66.7% error, 28 samples misclassified as *grass*), *rattan* (64.3%, 27 samples misclassified as *wood*) and *wood* (33.3% error, 14 samples misclassified as *rattan*), which contribute 69 of the 79 misclassified samples. Interestingly, in all 79 cases the model of the true class ranks second right after the nearest model of a false class that leads to misclassification. The distribution of rotation angles among the misclassified samples is surprisingly even, as all six testing angles contribute from 10 to 16 misclassified samples (16, 16, 14, 10, 13, 10). LBP_{16}^{riu2} does much better, classifying all samples correctly except ten *grass* samples that are assigned to *leather*. Again, in all ten cases the model of the true class *grass* ranks second.

We see that combining the LBP operators with the VAR measures, which do not do too badly by themselves, improves the performance considerably. In the case of LBP_8^{riu2}/VAR_8 the 1.64% error is caused by 11 *straw* samples erroneously assigned to class *grass*. In ten of the 11 misclassifications the model of the *straw* class ranks second, once third. LBP_{16}^{riu2}/VAR_{16} falls one sample short of a faultless result, as a *straw* sample at 90° angle is labeled as *grass*.

Note that we voluntarily discarded the knowledge that training samples come from four different rotation angles, merging all sample histograms into a single model for each texture class. Hence the final texture model is an ‘average’ of the models of the four training angles, which actually decreases the performance to a certain extent. If we had used four separate models, one for each training angle, for example LBP_{16}^{riu2}/VAR_{16} would have provided a perfect classification result, and the error rate of LBP_{16}^{riu2} would have decreased by 50% to 0.74%.

Even though a direct comparison to the results of Porter and Canagarajah may not be meaningful due to the different classification principle, the excellent results for LBP_{16}^{riu2} and LBP_{16}^{riu2}/VAR_{16} demonstrate their suitability for rotation invariant texture classification.

Table 2 presents results for a more challenging experimental setup, where the classifier is trained with samples of just one rotation angle and tested with samples of other nine rotation angles. We trained the classifier with the 121 16x16 samples extracted from the designated training image, again merging the histograms of the 16x16 samples of a particular texture class into one model histogram. The classifier was tested with the samples obtained from the other nine rotation angles of the seven source images reserved for testing purposes, totaling 1008 samples, 63 in each of the 16 texture classes. Note that in each texture class the seven testing images are physically different from the one designated training image, hence this setup is a true test for the texture operators’ ability to produce a rotation invariant representation of local image texture that also generalizes to physically different samples.

Table 2: Error rates (%) when training is done at just one rotation angle, and the average error rate over the ten angles.

OPERATOR	BINS	TRAINING ANGLE										AVERAGE
		0°	20°	30°	45°	60°	70°	90°	120°	135°	150°	
LBP_8^{riu2}	10	31.5	13.7	15.3	23.7	15.1	15.6	30.6	15.8	23.7	15.1	20.00
LBP_{16}^{riu2}	18	3.8	1.0	1.4	0.9	1.6	0.9	2.4	1.4	1.2	2.3	1.68
VAR_8	128	7.5	3.4	5.4	6.0	4.4	3.1	6.1	5.8	5.4	4.4	5.12
VAR_{16}	128	10.1	15.5	13.8	9.5	12.7	14.3	9.0	10.4	9.3	11.5	11.62
LBP_8^{riu2}/VAR_8	10/16	0.9	5.8	4.3	2.7	4.8	5.6	0.7	4.0	2.7	4.4	3.56
LBP_{16}^{riu2}/VAR_{16}	18/16	0.0	0.5	0.6	0.6	0.6	0.4	0.0	0.5	0.5	0.3	0.40

Training with just one rotation angle allows a more conclusive analysis of the rotation invariance of our operators. For example, it is hardly surprising that LBP_8^{riu2} provides highest error rates when the training angle is a multiple of 45°. Due to the crude quantization of the angular space the presentations learned at 0°, 45°, 90°, or 135° do not generalize that well to other angles.

LBP_{16}^{riu2} provides a solid performance with an average error rate of 1.68%. If we look at the ranks of the true class in the 169 misclassifications, we see that in every case the model of the true class ranks second. There is a strong suspicion that the subpar results for training angles 0° and 90° are due to the artificial blur added to the original images at angles 0° and 90° . The effect of the blur can also be seen in the results of the joint distributions LBP_8^{riu2}/VAR_8 and LBP_{16}^{riu2}/VAR_{16} , which achieve best performance when the training angle is either 0° or 90° , the 16-bit operator pair providing in fact a perfect classification in these cases. Namely, when training is done with some other rotation angle, test angles 0° and 90° contribute most of the misclassified samples, actually all of them in the case of LBP_{16}^{riu2}/VAR_{16} . Nevertheless, the results for LBP_{16}^{riu2} and LBP_{16}^{riu2}/VAR_{16} are quite excellent.

3.2 Experiment #2

Haley and Manjunath [4] proposed a method based on a complete space-frequency model for rotation-invariant texture classification. They developed a polar analytic form of a two-dimensional Gabor wavelet, and used a multiresolution family of these wavelets to compute texture microfeatures. Rotation invariance was achieved by transforming Gabor features into rotation invariant features using autocorrelation and DFT magnitudes and by utilizing rotation invariant statistics of rotation dependent features. Classification results were presented for two groups of textures, of which we use the set of textures available in the WWW [13].

Image Data and Experimental Setup. The image data comprised of the 13 textures from the Brodatz album shown in Fig. 5. For each texture 512x512 images digitized at six different rotation angles (0° , 30° , 60° , 90° , 120° , and 150°) were included. The images were divided into 16 disjoint 128x128 subimages, totaling 1248 samples, 96 in each of the 13 classes. Half of the subimages, separated in a checkerboard pattern, were used to estimate the model parameters, while the other half was used for testing. Using a multivariate Gaussian discriminant, Haley and Manjunath reported 96.8% classification accuracy.

Experimental Results. We first replicated the original experiment by computing the histograms of the training half of the 128x128 samples, which served as our model histograms. Since a 128x128 sample produces a sufficient number of entries (16129/15876) for its histogram to be stable, we did not combine individual histograms. Consequently, we had 624 model histograms in total, 48 (6 angles x 8 images) models for each of the 13 texture classes.

Since the training data includes all rotation angles, this problem is not particularly interesting in terms of rotation invariant texture classification and we restrict ourselves to merely reporting the error rates in Table 3. Both halves of the mosaic partitioning served as the training data in turn, the other being used as test samples, and as the final result we provide the average of the error rates of these two cases. We see that the results obtained with the joint distributions compare favorably to the 3.2% error rate reported by Haley and Manjunath.




Fig. 5. Texture images of Experiment #2 printed at particular rotation angles. Each texture was digitized at six angles: 0°, 30°, 60°, 90°, 120°, and 150°. Images are 512x512 pixels in size.

Table 3: Error rates (%) for the original experimental setup, where training data includes all rotation angles.

OPERATOR	BINS	ERROR	OPERATOR	BINS	ERROR
LBP_8^{riu2}	10	5.05 (5.61 & 4.49)	VAR_8	128	2.48 (2.56 & 2.40)
LBP_{16}^{riu2}	18	3.12 (3.04 & 3.21)	VAR_{16}	128	3.21 (2.56 & 3.85)
LBP_8^{riu2}/VAR_8	10/16	0.40 (0.16 & 0.64)	LBP_{16}^{riu2}/VAR_{16}	18/16	0.48 (0.16 & 0.64)

Again, we constructed a true test of rotation invariant texture classification, where the classifier is trained with the samples of just one rotation angle and tested with the samples of other five rotation angles. We trained the classifier with the 128x128 samples extracted from the 512x512 images of a particular rotation angle, obtaining 208 models in total, 16 for each of the 13 texture classes. The classifier was then evaluated with the 128x128 samples extracted from the 512x512 images of the other five rotation angles, totaling 1040 test samples.

Table 4: Error rates (%) when training is done at just one rotation angle and the average error rate over the six rotation angles.

OPERATOR	BINS	TRAINING ANGLE						AVERAGE
		0°	30°	60°	90°	120°	150°	
LBP_8^{riu2}	10	20.2	13.7	13.7	17.7	17.0	8.8	15.18
LBP_{16}^{riu2}	18	10.4	8.2	8.5	8.6	8.3	6.9	8.46
VAR_8	128	7.9	6.2	4.2	5.2	3.9	3.5	5.14
VAR_{16}	128	7.6	6.3	4.6	3.8	3.7	4.7	5.11
LBP_8^{riu2}/VAR_8	10/16	2.1	2.5	0.8	0.5	1.2	0.5	1.25
LBP_{16}^{riu2}/VAR_{16}	18/16	1.9	1.0	0.5	0.3	0.2	0.3	0.69

From the error rates in Table 4 we observe that using just one rotation angle for training indeed increases the difficulty of the problem quite nicely. If we take a closer look at the confusion matrices of LBP_{16}^{riu2} (8.4% average error rate), we see that about half (246/528) of the misclassifications are due to the samples of the strongly oriented texture *wood* being erroneously assigned to *straw*. The training angle does not seem to affect the classification accuracy too much, as roughly an equal result is obtained in all six cases.

The complementary nature of LBP and VAR operators shows in the excellent results for their joint distributions. LBP_{16}^{riu2}/VAR_{16} achieves a very low average error rate of 0.69%, which corresponds to just about 7 misclassifications out of 1040 samples. Of the 43 misclassifications in total, false assignments of *wool* samples to *pigskin* contribute 16 and of *grass* samples to *leather* 11. It is worth noting that the performance is not sensitive to the quantization of the VAR feature space, as following average error rates are obtained by LBP_{16}^{riu2}/VAR_{16} with different numbers of bins: 1.31% (18/2), 0.71% (18/4), 0.64% (18/8), 0.69% (18/16), 0.71% (18/32), and 0.74% (18/64).

4 Discussion

We presented a theoretically and computationally simple but efficient approach for gray scale and rotation invariant texture classification based on local binary patterns and nonparametric discrimination of sample and prototype distributions. Excellent

experimental results obtained in two problems of true rotation invariance, where the classifier was trained at one particular rotation angle and tested with samples from other rotation angles, demonstrate that good discrimination can be achieved with the occurrence statistics of simple rotation invariant local binary patterns. The proposed approach is very robust in terms of gray scale variations, since the operators are by definition invariant against any monotonic transformation of the gray scale. This should make our operators very attractive in situations where varying illumination conditions are a concern, e.g. in visual inspection. Computational simplicity is another advantage, as the operators can be realized with a few comparisons in a small neighborhood and a lookup table. This facilitates a very straightforward and efficient implementation, which may be mandatory in time critical applications. If the stability of the gray scale is not something to be worried about, performance can be further improved by combining the LBP_8^{riu2} and LBP_{16}^{riu2} operators with rotation invariant variance measures VAR_8 and VAR_{16} that characterize the contrast of local image texture. As we observed in the experiments, the joint distributions of these orthogonal operators are very powerful tools for rotation invariant texture analysis.

Regarding future work, in this study we reported results for two rotation invariant LBP operators having different spatial configuration of the circularly symmetric neighbor set, which determines the angular resolution. As expected, LBP_{16}^{riu2} with its more precise quantization of the angular space provides clearly better classification accuracy. Nothing prevents us from using even larger circularly symmetric neighbor sets, say 24 or 32 pixels with a suitable spatial predicate, which would offer even better angular resolution. Practical implementation will not be as straightforward, though, at least not for the 32-bit version. Another interesting and related detail is the spatial size of the operators. Some may find our experimental results surprisingly good, considering how small the support of our operators is for example in comparison to much larger Gabor filters that are often used in texture analysis. However, the built-in support of our operators is inherently larger than 3×3 or 5×5 , as only a limited subset of patterns can reside adjacent to a particular pattern. Still, our operators may not be suitable for discriminating textures where the dominant features appear at a very large scale. This can be addressed by increasing the spatial predicate, as the operators can be generalized to any neighborhood size. Further, operators with different spatial resolutions can be combined for multiscale analysis, and ultimately, we would want to incorporate scale invariance, in addition to gray scale and rotation invariance. Another thing deserving a closer look is the use of a problem or application specific subset of rotation invariant patterns, which may in some cases provide better performance than ‘uniform’ patterns. Patterns or pattern combinations are evaluated with some criterion, e.g. classification accuracy on a training data, and the combination providing the best accuracy is chosen. Since combinatorial explosion may prevent from an exhaustive search through all possible subsets, suboptimal solutions such as stepwise or beam search should be considered. We also reported that when there are classification errors, the model of the true class very often ranks second. This suggests that classification could be carried out in stages, by selecting features which best discriminate among remaining alternatives.

Acknowledgments

The authors wish to thank Dr. Nishan Canagarajah and Mr. Paul Hill from the University of Bristol for providing the texture images of Experiment #1. The financial support provided by the Academy of Finland is gratefully acknowledged.

Note

Texture images used in this study, together with other imagery used in our published work, can be downloaded from <http://www.ee.oulu.fi/research/imag/texture>.

References

1. Brodatz, P.: Textures: a Photographic Album for Artists and Designers. Dover, New York (1966)
2. Chen, J.-L., Kundu, A.: Rotation and Gray Scale Transform Invariant Texture Identification Using Wavelet Decomposition and Hidden Markov Model. *IEEE Trans. Pattern Analysis and Machine Intelligence* **16** (1994) 208-214
3. Cohen, F.S., Fan, Z., Patel, M.A.: Classification of Rotated and Scaled Texture Images Using Gaussian Markov Random Field Models. *IEEE Trans. Pattern Analysis and Machine Intelligence* **13** (1991) 192-202
4. Haley, G.M., Manjunath, B.S.: Rotation-Invariant Texture Classification Using a Complete Space-Frequency Model. *IEEE Trans. Image Processing* **8** (1999) 255-269
5. Kashyap, R.L., Khotanzad, A.: A Model-Based Method for Rotation Invariant Texture Classification. *IEEE Trans. Pattern Analysis and Machine Intelligence* **8** (1986) 472-481
6. Lam, W.-K., Li, C.-K.: Rotated Texture Classification by Improved Iterative Morphological Decomposition. *IEE Proc. - Vision, Image and Signal Processing* **144** (1997) 171-179
7. Madiraju, S.V.R., Liu, C.C.: Rotation Invariant Texture Classification Using Covariance. *Proc. Int. Conf. Image Processing* **1** (1995) 262-265, Washington, D.C.
8. Mao, J., Jain, A.K.: Texture Classification and Segmentation Using Multiresolution Simultaneous Autoregressive Models. *Pattern Recognition* **25** (1992) 173-188
9. Ojala, T., Pietikäinen, M., Harwood, D.: A Comparative Study of Texture Measures with Classification Based on Feature Distributions. *Pattern Recognition* **29** (1996) 51-59
10. Ojala, T., Valkealahti, K., Oja, E., Pietikäinen, M.: Texture Discrimination with Multidimensional Distributions of Signed Gray Level Differences. *Pattern Recognition* (2000), in press
11. Pietikäinen, M., Ojala, T., Xu, Z.: Rotation-Invariant Texture Classification Using Feature Distributions. *Pattern Recognition* **33** (2000) 43-52
12. Porter, R., Canagarajah, N.: Robust Rotation-Invariant Texture Classification: Wavelet, Gabor Filter and GMRF Based Schemes. *IEE Proc. - Vision, Image and Signal Processing* **144** (1997) 180-188
13. The USC-SIPI Image Database, Volume 1: Rotated Textures. Signal and Image Processing Institute of the University of Southern California, Los Angeles. <http://sipi.usc.edu/services/database/database.cgi?volume=rotate>
14. Wu, W.-R., Wei, S.-C.: Rotation and Gray-Scale Transform-Invariant Texture Classification Using Spiral Resampling, Subband Decomposition and Hidden Markov Model. *IEEE Trans. Image Processing* **5** (1996) 1423-1434