Combining global with local texture information for image retrieval applications

Javier A. Montoya-Zegarra[1,2], Jan Beeck[1]
Computer Engineering Department[1]
San Pablo Catholic University
Av. Salaverry 301, Vallecito, Arequipa, Peru
{jmontoyaz,jcbeeck}@gmail.com

Neucimar Leite,Ricardo Torres,Alexandre Falcão[2]
Institute of Computing[2]
University of Campinas
Caixa Postal 6176, Campinas, SP, Brazil
{neucimar,afalcao,rtorres}@ic.unicamp.br

Abstract

This paper proposes a new texture descriptor to guide the search and retrieval in image databases. It extracts rich information from global and local primitives of textured images. At a higher level, the global macro-features in textured images are characterized by exploiting the multi-resolution properties of the Steerable Pyramid Decomposition. By doing this, the global texture configurations are highlighted. At a finer level, the local arrangements of texture micro-patterns are encoded by the Local Binary Pattern operator.

Experiments were carried out on the standard Vistex dataset aiming to compare our descriptors against popular texture extraction methods with regard to their retrieval accuracies. The comparative evaluations allowed us to show the superior descriptive properties of our feature representation methods.

1 Introduction

Texture is one of the most important low-level image features used for both human perception and recognition [12]. Despite of the advances in texture analysis, an effective texture representation is still an open task due to the small inter-class variations among texture images as well as due to the presence of image distortions such as affine transformations or changes in contrast/illumination.

Some of these challenges are faced in this work. More specifically, we are interested in providing effective mechanisms for capturing relevant texture information. Pioneer works concluded that relevant texture primitives are localized at different resolutions [1, 4]. At coarse resolutions, the particular macro-regularities in textures appear, whereas at finer levels, the local arrangements of micro-patterns become more notorious. These observations help us to consider that, intuitively, a robust texture representation mechanism should be able to analyze texture images both globally and locally.

Most of the previous works in texture image retrieval consider only global image patterns (texture macro-regularities) [2, 7, 10]. Thus, their image representations

lack of the local information available in texture micro-regularities. More recent approaches that use global and local texture information are the methods of Jain et al. [5] and Zhang et al. [13]. However, since they consider biometric images (fingerprints and faces, respectively), different needs appear.

To capture texture macro-patterns, our method is benefited from the multi-resolution properties of the Steerable Pyramid [3, 10]. Previous studies have demonstrated that this image representation presents good discriminative properties for texture characterization [8]. The texture micro-patterns are then extracted from the steerable multi-resolutions subbands by applying the Local Binary Pattern (LBP) operator [9]. Some important key characteristics of this operator include its computational efficiency as well as its high discriminative properties at local regions.

The remainder of this paper is structured as follows. Next Section introduces the proposed texture representation method, whilst Section 3 presents the experimental setup conducted in our study. In Section 4, the experimental results on the MIT Media Laboratory VisTex [6] dataset are given and are used to demonstrate the retrieval effectiveness improvement of our approach. Comparisons with other approaches are also discussed. Finally, conclusions and future research directions are drawn in Section 5.

2 Texture Feature Representation

The general overview of the proposed method is shown in Figure 1. Its main components can be summarized as follows:

1. Find the Global Texture Representation of the images by extracting the magnitudes of each of the Steerable Pyramid subbands (Subsection 2.1).

2. Compute the Local Texture Representation of the magnitude subbands by considering either the Non-uniform or Uniform Local Binary Pattern operator (Subsection 2.2).
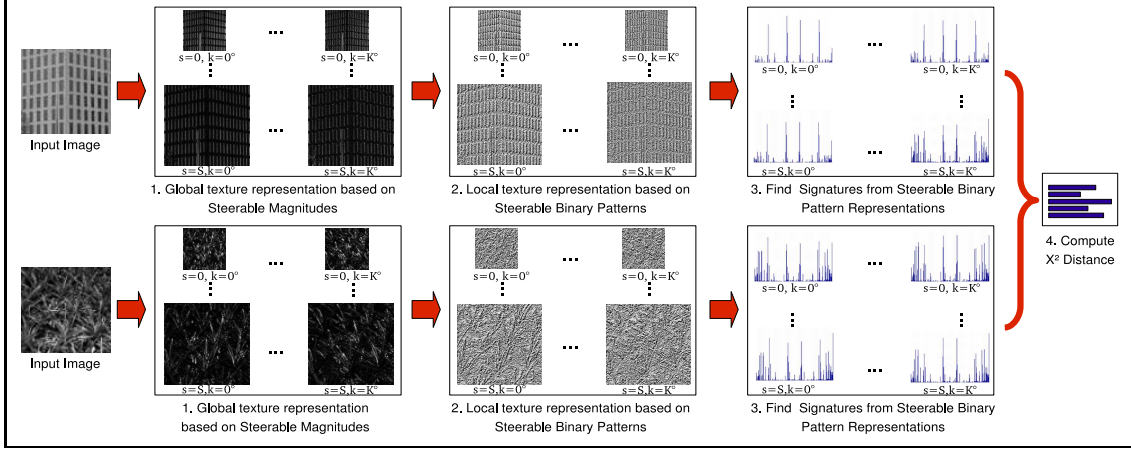
Figure 1. Overview of the proposed texture feature representation.

3. Extract the global and local signatures by first computing and then by concatenating the histograms of the previous encoded subbands (Subsection 2.3).

4. Compute the similarity between the texture images by comparing their corresponding signatures using the Chi-square distance (Subsection 2.4).

2.1 Global Texture Representation based on Steerable-Pyramid Magnitudes

The basis of our descriptor relies on the Steerable Pyramid Decomposition. It consists in a linear multi-resolution image decomposition method by which an image is subdivided into a collection of sub-bands localized at different scales and orientations [3]. A Steerable Pyramid is implemented by recursively splitting a given image into a set of oriented subbands and a lowpass residual bank [10]. The response on the lowpass subband is used to iterate the decomposition. By using a set of basis filters, which are translations and rotations of a single function, the oriented filters are computed by linear combinations of those basis filters. The spectral representation of the bandpass filter at scale $s = \{1, \ldots, S\}$ and orientation $k = \{1, \ldots, K\}$ is computed as:

$$B_{sk}(u,v) = \begin{cases} B_s(u,v), & \theta_{k-1} \leq \phi_k(u,v) \leq \theta_k \\ 0, & otherwise \end{cases} \quad (1)$$

where $\theta_k = \frac{k*PI}{K}$ and

$$\phi_k(u,v) = \begin{cases} arctan(\frac{v}{u}), & 0 < \theta_k \leq \frac{\pi}{2} \\ \pi + arctan(\frac{v}{u}), & \frac{\pi}{2} < \theta_k \leq \pi \end{cases} \quad (2)$$

The bandpass filter at scale s is computed as a combination of a cosine low-, and high-pass transfer functions:

$$B_s(u,v) = L_s(u,v)H_s(u,v) \quad (3)$$

having $r = \sqrt{u^2 + v^2}$, the low-pass filter is given by

$$L(r) = \begin{cases} 2 & r \leq \frac{\pi}{4} \\ 2\cos\left(\frac{\pi}{2}\log_2\left(\frac{4r}{\pi}\right)\right) & \frac{\pi}{4} < r < \frac{\pi}{2} \\ 0 & r \geq \frac{\pi}{2} \end{cases} \quad (4)$$

whilst the high-pass filter is defined as

$$H(r) = \begin{cases} 1 & r \geq \frac{\pi}{2} \\ \cos\left(\frac{\pi}{2}\log_2\left(\frac{2r}{\pi}\right)\right) & \frac{\pi}{4} < r < \frac{\pi}{2} \\ 0 & r \leq \frac{\pi}{4} \end{cases} \quad (5)$$

Given an image $I(x,y)$, its Steerable-Pyramid Decomposition is defined as:

$$S_{sk}(x,y) = \sum_{p=-\infty}^{p=\infty} \sum_{q=-\infty}^{q=\infty} I(p,q) * B_{sk}(x-p, y-q) \quad (6)$$

where $*$ denotes the convolution operator. In our method, we considered four scales ($S = 4$) and six orientations ($K = 6$). With these configurations, a total number of 24 Steerable Filters are generated. After computing the magnitude of each pixel in the filter outputs, 24 Steerable-Pyramid Magnitude Representations are generated.

2.2 Local Texture Representation based on Steerable Binary Patterns

The Local Binary Pattern (LBP) operator is an important method for describing texture images. It characterizes different types of micro-patterns such as: constant areas, edges, points, etc. [9]. One of its important key properties include its monotonic gray-level invariance and its computational efficiency. In our approach, we exploit these properties in order to capture relevant local information from the global features encountered in the Steerable-Pyramid Magnitude Representations. In traditional approaches, in which

just the global information is considered, local image areas of interest may be missed. Next subsections formalize the Steerable Binary Pattern representation by introducing the original LBP operator (Non-uniform LBP) and its extension known as Uniform LBP [9].

2.2.1 Non-uniform Steerable Binary Pattern:

The original version of the LBP operator, also known as Non-uniform LBP, is applied on a 3×3 image neighborhood represented by g_p $(p = 0, 1, 2, \ldots, 7)$ [9]. The pixel values in the selected region are initially thresholded by the value of the center pixel g_c as follows:

$$S(g_p - g_c) = \begin{cases} 1, & g_p \geq g_c \\ 0, & g_p < g_c \end{cases} \tag{7}$$

Then, the LBP pattern of the image neighborhood is obtained by summing the corresponding thresholded values $S(g_p - g_c)$ weighted by a binomial factor of 2^p:

$$LBP = \sum_{p=0}^{7} S(g_p - g_c) \, 2^p \tag{8}$$

This whole process is repeated for the remaining 3×3 overlapping neighborhoods of the input image. Note that, in the proposed method, the LBP operator is applied to each of the Steerable-Pyramid Magnitudes (SPMs). In this sense, the operator $S_{lbp}(x, y, s, k)$ corresponds to the (x, y) pixel location of the LBP, applied to the Steerable-Pyramid Magnitude at scale s and orientation k. We denote this texture representation as the Non-uniform Steerable Binary Pattern (NU-SBP).

2.2.2 Uniform Steerable Binary Pattern:

The original LBP operator was extended to incorporate a fixed set of *uniform* rotation-invariant patterns, and hence its name: Uniform LBP operator [9]. In the original LBP approach, if a given image was rotated, so were its corresponding LBP features.
An uniform pattern consists basically in a circular structure that contains binary spatial transitions. The spatial transitions of the uniform patterns characterize a certain number of different texture primitives. Examples of such primitives include bright spots (circular structure of 0's), dark spots (circular structure of 1's), etc.
To quantify these patterns, an uniformity measure U was introduced and corresponds to the number of bitwise spatial transitions (0/1 changes) in the current pattern. For example, if we consider the bright spot pattern as being the 00000000 bit sequence in a 3×3 neighborhood, its corresponding uniformity measure U will be 0, since no 0/1 transitions appeared. Ojala et al. [9] noticed that, the higher the uniformity measure is the more sensitive the features are

upon rotations. Furthermore, they concluded that the most relevant texture information is associated with patterns having at most a uniformity measure U equal to 2. The Uniform LBP operator is expressed as follows:

$$ULBP = \begin{cases} \sum_{p=0}^{7} S(g_p - g_c), & if \; U(LBP) \leq 2 \\ 9, & otherwise \end{cases} \tag{9}$$

where the uniformity measure U is defined as:

$$\begin{aligned} U(LBP) \quad &= |S(g_{p-1} - g_c) - S(g_0 - g_c)| \\ &+ \sum_{p=0}^{7} |S(g_p - g_c) - S(g_{p-1} - g_c)| \end{aligned} \tag{10}$$

In equation 9, an unique label is assigned to the patterns that have an uniformity value of at most *2*, whilst the non-uniform patterns are assigned under the miscellaneous label *9*. The Uniform LBP operator $(S_{ulbp}(x, y, s, k))$ is applied in each of the Steerable-Pyramid Magnitudes (SPMs). As in the case of the Non-uniform Steerable Binary Pattern, this operator represents the (x, y) pixel location of the ULBP applied in the Steerable-Pyramid Magnitude at scale s and orientation k. This texture representation is denoted as the Uniform Steerable Binary Pattern (U-SBP).

2.3 Steerable Binary Pattern Histogram

In our method, the texture images are modeled by histograms that characterize the occurrence statistics of both global and local texture patterns. Thus, we denote this representation as Steerable Binary Pattern Histogram (SBPH). To compute the SBPH, first the histogram H of the gray level image $I(x, y)$ with $0, 1, \ldots, L$ different gray values is computed as follows:

$$H(l) = \sum_{x=0}^{M} \sum_{y=0}^{N} q_l(I(x, y)), \;\; l = 0, 1, 2, \ldots, L - 1 \tag{11}$$

where l represents a specific gray value, $H(l)$ denotes the number of pixels at gray level l in the input image, and q_l is a mapping function defined as:

$$q_l(f) \begin{cases} 1, & if \; l \in F_l \\ 0, & otherwise \end{cases} \tag{12}$$

the variable F_l is used to represent the values that fall into the l region. The Steerable Binary Pattern Histogram is then found by first computing the histogram of each image representation in the Steerable Binary Pattern:

$$H_{s,k} = \sum_{x=0}^{M} \sum_{y=0}^{N} q_l(S_{lbp}(x, y, s, k)) \tag{13}$$

Then, to build the single global histogram used to represent the texture image, the different histograms are concatenated as a unique sequence:

$$SBPH = (H_{0,0}, \ldots, H_{0,K}; H_{1,1}, \ldots, H_{1,K-1}$$
$$H_{S-2,0}, \ldots, H_{s-2,K-1}; H_{S-1,1}, \ldots, H_{S-1,K-1}) \quad (14)$$

Note that the size of the histogram may vary depending whether it is computed from the Non-uniform Steerable Binary Pattern representations or from the ones belonging to the Uniform Steerable Binary Pattern. In the former, the number of bins in the histogram is equal to 256, whereas in the latter, it is reduced to 59 bins. The smaller size of the uniform histogram is because by considering a uniformity measure (U) of at most 2, only a total number of nine uniform patterns will be generated, which will lead in turn to 58 circularly rotated versions used for invariant analysis and one additional bin to characterize the miscellaneous patterns [9].

2.4 Signature Similarity

Similarity between texture images is obtained by computing the distance of their corresponding feature vectors. The smaller the distance, the more similar the images. Given the query image i, and the target image j in the dataset, the distance between the two patterns is defined as:

$$X^2(H^i, H^j) = \sum_{l=0}^{L} \frac{(H_l^i - H_l^j)^2}{(H_l^i + H_l^j)} \quad (15)$$

where L corresponds to the length of the Steerable Binary Pattern Histogram. If the SBPH was obtained by using the Non-uniform LBP operator, then the length of the feature vector comprises $6144 = 256 \times 24$ elements (Recall that *four* scales and *six* orientations were used for generating the Steerable-Pyramid Magnitude Representations. On the other hand, if the Steerable Binary Pattern Histogram was obtained by means of the Uniform LBP operator, then the length of the feature vector is reduced to just $1416 = 59 \times 24$ elements.

3 Experimental setup

3.1 Dataset

To evaluate the retrieval effectiveness of our approach, 40 texture images were selected from the VisTex standard dataset [6]. The selected images follow the work of Do et al. [2], since this method will be used as a reference in our experiments. The 512×512 images correspond to different natural scenes and only their luminance components are considered. To generate the image dataset, each original texture image was partitioned into sixteen 128×128 non-overlapping subimages. Thus, this dataset comprises 640 (40×16) different images.

In addition, the main peculiarities of the images in this dataset include: (i) small inter-class variations, therefore, although some images may appear very similar to each other, they may belong to different classes, (ii) different types of visual patterns, consequently, the images vary in uniformity, smoothness, directionality, etc., and (iii) homogeneity, since in CBIR applications the relevant images are traditionally represented by subimages of the same source, then the visual properties of the textures should not change much over the entire image.

3.2 Retrieval Effectiveness Evaluation

In our experiments, each image in the dataset was used to simulate a query. The relevant images for each query are defined as the 15 remaining subimages from the same input texture. In this sense, a total number of $408\,960$ (639×640) queries were performed. Furthermore, the retrieval effectiveness was measured in terms of relevant retrieval average rate, i.e., the percentage of relevant images among the top N retrieved images.

4 Experimental Results

To demonstrate the discriminating properties of the proposed method, two series of experiments were conducted. In both cases, our method was compared against two other approaches: the conventional Steerable Pyramid Decomposition [11] and a proposal that models the distribution of Wavelet coefficients using generalized Gaussian density with the Kullback-Leibner measure as similarity measure (GGD & KLD) [2].

In the case of the conventional Steerable Pyramid Decomposition, we used four scales and six orientations to represent the texture images. To build the feature vectors, we computed the mean and standard deviation of the magnitude of each filter response. For the sake of similarity computation, the Weighted-Mean-Variance distance was used [7]. In the case of Wavelet-based generalized Gaussian density model, we used three levels of decomposition of the Daubechies 4-tap filters (D_4). This configuration follows the work of Do et al. [2].

In the first series of experiments, we wanted to know, how well the methods in study performed in the texture dataset. To answer this question, we compared the relevant retrieval average of each technique by varying the number of the top retrieved images from $N = 16$ to $N = 64$, this means that 1% to 4% of the images were used during the retrieval phase. In the ideal case, a perfect retrieval accuracy would be achieved, if the average retrieval rate would be equal to 100% for all image classes after $n-1$ retrievals, where n denotes the number of subimages of the same class (in our case $n = 16$).

The graph in Figure 2 serves us to illustrate these experiments. As can be seen, both Non-uniform and Uniform SBPH perform better than the other two reference ap-
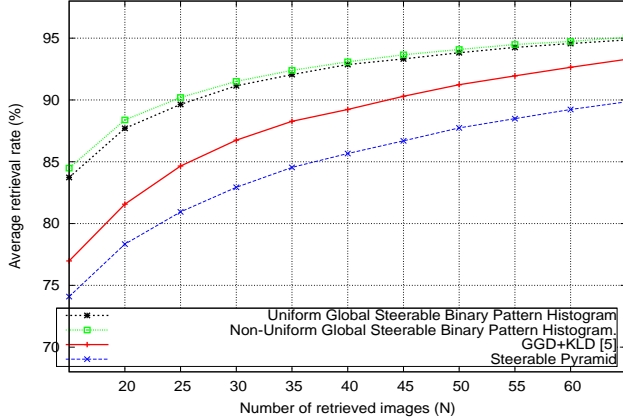
Figure 2. Average retrieval rates per texture class for *rotated* image dataset *without* and *with* rotation-invariant representation, respectively.

proaches. The accuracy of our both descriptors was followed by the method of Do et al. [2] and by the conventional Steerable Pyramid Decomposition [11], respectively. Furthermore, the referred Figure reveals some relevant information:

1. First, by comparing the curves of the Non-uniform and the Uniform SBPH, one can observe that both methods achieve very similar results. The higher retrieval accuracy of the Non-uniform SBPH relies on the fact that by definition, Non-uniform LBPs are capable of capturing more texture information than Uniform Patterns, since more texture primitives are modeled.

   In addition, our experiments agree with [9] in that, although the number of texture primitives is reduced in the Uniform LBPs, the selected patterns provide enough discriminative information for characterizing texture images. This justifies the very slight difference of retrieval accuracy between the two curves.

2. Second, although the Uniform SBPH achieved inferior retrieval accuracy than its Non-uniform counterpart, it is benefited from the reduced length of its feature representations, since they are $23.05\%(1416/6144)$ smaller. Therefore, if data storage capacity is limited, then the Uniform SBPH represents a good choice for texture retrieval applications.

3. Finally, in the case of the conventional Steerable Pyramid Decomposition, almost the double number of retrieved images are required to achieve the same retrieval accuracy of our descriptors, whereas in the case of the Wavelet-based generalized Gaussian density method, this number is slightly smaller.

To this end, we have analyzed how well the descriptors perform in average on the whole dataset. Therefore, in the second series of experiments, we have summarized the retrieval accuracy achieved by the descriptors in each of the 40 texture classes considering the $N = 16$ top retrieved images (See Table 1). From the results, one can observe that the highest retrieval rates increased by both Non-uniform and Uniform SBPH correspond to homogeneous classes, such as Bark $(0, 8, 9)$ and Brick $(1, 4)$. Furthermore, texture classes having fine micro-patterns were also benefited. Examples of such classes include Leaves $(8, 11)$ and Flowers5. These results validate our intuition that for characterizing texture images it is not just enough to capture their global visual primitives (e.g. directionality, granularity, etc.), but also to combine these informations with their fine local patterns such as edges, spots, etc. On the other hand, examples of texture classes, in which our image descriptors were not benefited from the characterization of local patterns, include Food8 and Grass1. By considering that both texture images present non-homogeneous local regions, local gray-level distortions compromise the discriminative capabilities of both Non-uniform and Uniform SBPH.

Finally, the lowest retrieval rates achieved by the four descriptors correspond to the same texture class: Wood1. This observation relies basically on the fact that this class presents different lighting conditions across its surface.

5 Conclusions

In this paper, we have introduced a new texture representation method that encodes the global and local primitives of texture images. At a higher level, the global patterns in the textures are characterized by exploiting the multi-resolution properties of the Steerable Pyramid Decomposition. By doing this, the global texture configurations, such as directionality and coarseness, are highlighted. At a finer level, the local micro-primitives of texture images are captured by applying the discriminative Non-uniform and Uniform LBP operators on the multi-resolution representations. These operators represent the basis for capturing the micro-textures, such as edges, flat areas, etc. By considering these remarks, our descriptors are capable of encoding rich global and local texture information.

Experiments carried out on the Vistex dataset, compared the retrieval accuracy of our descriptors against the ones obtained by the other two techniques, namely the conventional Steerable Pyramid Decomposition [11] and Wavelet-based generalized Gaussian density model [2]. The comparative evaluation with both techniques allowed us, on the one side, to show the descriptive properties of our feature representation methods. On the other, it allowed us to confirm the necessity in capturing the global and local information of texture images to achieve good retrieval performance. This is shown in Figure 2, where the retrieval accuracy of the conventional Steerable Pyramid Decomposi-

	Average Retrieval Rate %					Average Retrieval Rate %			
	M1	M2	M3	M4		M1	M2	M3	M4
Bark0	97.26	97.65	60.15	54.29	Grass1	40.23	42.96	84.37	68.75
Bark6	67.96	67.96	76.17	48.82	Leaves8	96.48	93.35	75.00	66.01
Bark8	75.00	73.43	53.90	78.51	Leaves10	48.04	47.65	53.12	36.32
Bark9	79.68	80.46	35.15	55.85	Leaves11	87.89	85.54	54.68	69.53
Brick1	99.21	99.21	60.93	76.56	Leaves12	91.01	93.35	70.70	81.64
Brick4	77.73	75.78	46.09	69.53	Leaves16	58.20	53.12	42.96	80.85
Brick5	87.50	88.67	65.62	83.59	Metal0	73.04	67.96	82.81	73.43
Buildings9	99.60	99.60	87.89	92.96	Metal2	99.21	96.87	99.60	100.00
Fabric0	84.37	84.37	95.70	87.50	Misc2	84.76	84.37	76.56	79.68
Fabric4	69.14	67.96	68.35	64.45	Sand0	100.00	100.00	82.81	81.64
Fabric7	100.00	100.00	85.93	99.60	Stone1	76.56	75.39	72.26	54.29
Fabric9	95.31	96.09	100.00	84.37	Stone4	94.53	93.75	92.96	77.73
Fabric11	73.04	72.26	89.84	73.43	Terrain10	65.62	61.32	48.04	51.95
Fabric14	100.00	100.00	100.00	100.00	Tile1	66.01	64.84	51.17	51.95
Fabric15	97.26	97.65	93.75	98.43	Tile4	98.82	99.21	67.96	98.04
Fabric17	100.00	100.00	100.00	92.18	Tile7	100.00	100.00	100.00	100.00
Fabric18	97.26	96.09	100.00	98.04	Water5	100.00	100.00	57.03	95.31
Flowers5	91.01	89.06	65.62	57.42	Wood1	33.98	33.59	27.34	37.10
Food0	93.75	92.18	89.06	85.54	Wood2	97.65	98.43	76.56	83.98
Food5	95.31	95.31	81.64	90.23					
Food8	86.71	83.59	91.79	99.21	**Avg.**	**84.48**	**83.73**	**74.09**	**76.97**

Table 1. Average Retrieval Rate for the 40 VisTex Images.

tion was increased from 74.09% to 84.48% when using the Non-uniform SBPH and to 83.73% when using the Uniform SBPH. In the case of the Wavelet-based Gaussian density model, the retrieval accuracy was increased by almost 7%.

An important future research direction is to study how to increase the retrieval performance in non-homogeneous texture images. For this kind of images, the whole descriptors studied seemed to have difficulties in achieving discriminative characterizations. Intuitively, by augmenting the size of the local regions under characterization, the texture primitives may become more stable. Therefore, we further plan to study the configurations of the LBP operators that suit best these requeriments.

References

[1] A. C. Bovik, M. Clark, and W. S. Geisler. Multichannel texture analysis using localized spatial filters. *IEEE PAMI*, 12(1):55–73, 1990.

[2] M. N. Do and M. Vetterli. Wavelet-based texture retrieval using generalized gaussian density and kullback-leibler distance. *IEEE TIP*, 11(2):146–158, 2002.

[3] W. T. Freeman and E. H. Adelson. The design and use of steerable filters. *IEEE PAMI*, 13(9):891–906, 1991.

[4] R. M. Haralick. Statistical and structural approaches to texture. *Proceedings of IEEE*, 67(5):786–804, May 1979.

[5] A. K. Jain, S. Prabhakar, L. Hong, and S. Pankanti. Filterbank-based fingerprint matching. *IEEE TIP*, 9(5):846–859, 2000.

[6] M. M. Laboratory. Vistex: Texture image database (accessed on 1 march 2008). http://vismod.media.mit.edu/vismod/imagery/VisionTexture/.

[7] B. S. Manjunath and W.-Y. Ma. Texture features for browsing and retrieval of image data. *IEEE PAMI*, 18(8):837–842, 1996.

[8] J. A. Montoya-Zegarra, J. P. Papa, N. J. Leite, R. da S. Torres, and A. X. Falcão. Learning how to extract rotation-invariant and scale-invariant features from texture images. *EURASIP JASP*, 2008, 2008.

[9] T. Ojala, M. Pietikäinen, and T. Mäenpää. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE PAMI*, 24(7):971–987, 2002.

[10] J. Portilla and E. P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *Int. J. Comput. Vision*, 40(1):49–70, 2000.

[11] E. P. Simoncelli and W. T. Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. *ICIP*, 13(9):891–906, 1995.

[12] M. Tuceryan and A. K. Jain. Texture analysis. In P. S. P. W. C. H. Chen, L. F. Pau, editor, *The Handbook of Pattern Recognition and Computer Vision*, pages 207–248. World Scientific Publishing Co., 2 edition, 1998.

[13] W. Zhang, S. Shan, W. Gao, X. Chen, and H. Zhang. Local gabor binary pattern histogram sequence (lgbphs): A novel non-statistical model for face representation and recognition. In *ICCV '05, Volume 1*, pages 786–791, 2005.