

Rotation-Invariant and Scale-Invariant Steerable Pyramid Decomposition for Texture Image Retrieval

Javier A. Montoya-Zegarra^{1,2}
Computer Engineering Department¹
Faculty of Engineering
San Pablo Catholic University
Av. Salaverry 301, Vallecito
Arequipa, Peru
jmontoyaz@gmail.com

Neucimar J. Leite², Ricardo da S. Torres²
Institute of Computing²
State University of Campinas
Av. Albert Einstein, 1216
Campinas, Sao Paulo, Brazil
{neucimar,rtorres}@ic.unicamp.br

Abstract

This paper proposes a new rotation-invariant and scale-invariant representation for texture image retrieval based on Steerable Pyramid Decomposition.

*By calculating the mean and standard deviation of decomposed image subbands, the texture feature vectors are extracted. To obtain rotation or scale invariance, the feature elements are aligned by considering either the **dominant orientation** or **dominant scale** of the input textures.*

Experiments were conducted on the Brodatz database aiming to compare our approach to the conventional Steerable Pyramid Decomposition, and a recent proposal for texture characterization based on Gabor Wavelets with regard to their retrieval effectiveness. Results demonstrate the superiority of the proposed method in rotated and scaled image datasets.

1 Introduction

Texture-based image retrieval has been an active research topic over the last years. Although several methods have achieved high retrieval rates [3, 7, 8], some of them were evaluated under controlled scenarios, where there is a lack of image distortions, such as rotations and scales. Furthermore, methods that achieved rotation-invariant texture characterizations, where mostly proposed for classification applications. In general, these methods take advantage of a *priori* knowledge about the texture patterns, so that their classification rates can be improved [9, 11, 14].

In this context, the next challenge in texture image retrieval applications consists therefore in achieving rotation-, and scale-invariant feature representations for *non-controlled* environments.

This paper addresses these problems by proposing a new texture descriptor based on Steerable Pyramids. Roughly speaking, a Steerable Pyramid is a method in which images are decomposed into a set of multi-scale, and multi-orientation image subbands, where the basis functions are directional derivative operators [4]. Our motivation in using Steerable Pyramids relies not only on the fact that they have demonstrated discriminability properties for texture characterization [5], but also that unlike other image decomposition methods, the feature coefficients are less modified under the presence of image rotations, or even scales. The proposed descriptor computes the mean and standard deviation of the decomposed image subbands. To obtain rotation or scale invariance, extracted feature vectors are aligned by considering either the **dominant orientation** or **dominant scale** of the input textures.

The main contributions of this paper are: (1) the proposal of a new texture feature representation based on Steerable Pyramid Decomposition to facilitate rotation-invariant, and scale-invariant applications, and (2) a comprehensive evaluation that demonstrates the discriminating power of our approach in characterizing different classes of textures under the presence of image distortions (scales and orientations).

The outline of this paper is as follows. In the next section, we briefly review the fundamentals of the Steerable Pyramid Decomposition. Section 3 describes how texture images are characterized to obtain rotation-invariant, and scale-invariant representations. The experimental setup in our study is presented in section 4. In section 5, experimental results on several datasets are given, and are used to demonstrate the retrieval effectiveness improvement of our approach. Finally, some conclusions are drawn in section 6.

2 Steerable Pyramid Decomposition

The Steerable Pyramid Decomposition is a linear multi-resolution image decomposition method, by which an image is subdivided into a collection of subbands localized at different scales and orientations [4]. Using a high-, and low-pass filter (H_0, L_0) the input image is initially decomposed into two subbands: a high-, and a low-pass subband, respectively. Further, the low-pass subband is decomposed into K -oriented band-pass portions B_0, \dots, B_{K-1} , and into a lowpass subband L_1 . The decomposition is done recursively by subsampling the lower low-pass subband (L_S) by a factor of 2 along the rows and columns. Each recursive step captures different directional information at a given scale. Considering the polar-separability of the filters in the Fourier domain, the first low-, and high-pass filters, are defined as [13]:

$$\begin{aligned} L_0(r, \theta) &= L\left(\frac{r}{2}, \theta\right) / 2 \\ H_0(r, \theta) &= H\left(\frac{r}{2}, \theta\right) \end{aligned} \quad (1)$$

where r, θ are the polar frequency coordinates. L, H are raised cosine low-, and high-pass transfer functions:

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

$$B_k(r, \theta) = H(r)G_k(\theta), \quad k \in [0, K-1] \quad (3)$$

$B_k(r, \theta)$ represents the K directional bandpass filters used in the iterative stages, with radial and angular parts, defined as:

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

$$G_k(\theta) = \begin{cases} \alpha_K \left(\cos\left(\theta - \frac{\pi k}{K}\right)\right)^{K-1} & \left|\theta - \frac{\pi k}{K}\right| < \frac{\pi}{2} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where $\alpha_k = 2^{(k-1)} \frac{(K-1)!}{\sqrt{K[2^{(K-1)}]!}}$.

A first level image decomposition using Steerable Pyramid is shown in Figure 1.

3 Texture feature representation

This section describes the proposed modification of Steerable Pyramid Decomposition to obtain the rotation-invariant, and scale-invariant representations, which are used to characterize the texture images.

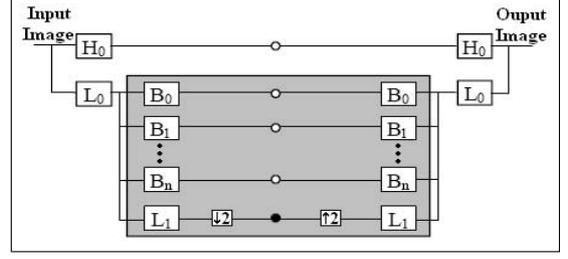


Figure 1. First level image decomposition based on Steerable Pyramid Decomposition.

3.1 Texture representation

Roughly speaking, texture images can be seen as a set of basic repetitive primitives characterized by their spatial homogeneity [2]. By applying statistical measures, this information is extracted, and used to capture the relevant image content into feature vectors. More precisely, we use the mean (μ_{mn}) and standard deviation (σ_{mn}) of the energy distribution of the filtered images (S_{mn}), by considering the presence of homogeneous regions in texture images. Given an image $I(x, y)$, its Steerable Pyramid Decomposition is defined as:

$$S_{mn}(x, y) = \sum_{x_1} \sum_{y_1} I(x_1, y_1) B_{mn}(x - x_1, y - y_1) \quad (6)$$

where B_{mn} denotes the directional bandpass filters at stage $m = 0, 1, \dots, S-1$, and orientation $n = 0, 1, \dots, K-1$. The energy distribution ($E(m, n)$) of the filtered images at scale m , and at orientation n is defined as:

$$E(m, n) = \sum_x \sum_y |S_{mn}(x, y)| \quad (7)$$

Additionally, the mean (μ_{mn}) and standard deviation (σ_{mn}) of the energy distributions are found as follows:

$$\mu_{mn} = \frac{1}{MN} E_{mn}(x, y) \quad (8)$$

$$\sigma_{mn} = \sqrt{\frac{1}{MN} \sum_x \sum_y (|S_{mn}(x, y)| - \mu_{mn})^2} \quad (9)$$

The corresponding feature vector (\vec{f}) is defined by using the mean and standard deviation as feature elements. It is denoted as:

$$\vec{f} = [\mu_{00}, \sigma_{00}, \mu_{01}, \sigma_{01}, \dots, \mu_{S-1K-1}, \sigma_{S-1K-1}] \quad (10)$$

3.2 Rotation-invariant representation

Rotation-invariant representation is achieved by computing the dominant orientation of the texture images followed by feature alignment. The **dominant orientation** (DO) is defined as the orientation with the highest total energy across the different scales considered during image decomposition [1]. It is computed by finding the highest accumulated energy for the K different orientations considered during image decomposition:

$$DO_i = \max \left\{ E_0^{(R)}, E_1^{(R)}, \dots, E_{K-1}^{(R)} \right\} \quad (11)$$

where i is the index where the dominant orientation appeared, and:

$$E_n^{(R)} = \sum_{m=0}^{S-1} E(m, n), \quad n = 0, 1, \dots, K-1. \quad (12)$$

Note that each $E_n^{(R)}$ covers a set of filtered images at different scales but at same orientation.

Finally, rotation-invariance is obtained by shifting circularly feature elements within the same scales, so that first elements at each scale correspond to dominant orientations. Let \vec{f} be a feature vector obtained by using a Pyramid Decomposition with $S = 2$ scales, and $K = 3$ orientations:

$$\vec{f} = \begin{bmatrix} \mu_{00}, \sigma_{00}, \mu_{01}, \sigma_{01}, \mu_{02}, \sigma_{02}; \\ \mu_{10}, \sigma_{10}, \mu_{11}, \sigma_{11}, \mu_{12}, \sigma_{12} \end{bmatrix} \quad (13)$$

Now suppose that the dominant orientation appears at index $i = 1$ ($DO_{i=1}$), thus the rotation-invariant feature vector, after feature alignment, is represented as follows:

$$\vec{f}^R = \begin{bmatrix} \mu_{01}, \sigma_{01}, \mu_{02}, \sigma_{02}, \mu_{00}, \sigma_{00}; \\ \mu_{11}, \sigma_{11}, \mu_{12}, \sigma_{12}, \mu_{10}, \sigma_{10} \end{bmatrix} \quad (14)$$

3.3 Scale-invariant representation

Similarly, scale-invariant representation is achieved by finding the scale with the highest total energy across the different orientations (**dominant scale**). For this purpose, the dominant scale (DS) at index i is computed as follows:

$$DS_i = \max \left\{ E_0^{(S)}, E_1^{(S)}, \dots, E_{S-1}^{(S)} \right\} \quad (15)$$

where $E_m^{(S)}$ denotes the accumulated energies across the S different scales:

$$E_m^{(S)} = \sum_{n=0}^{K-1} E(m, n), \quad m = 0, 1, \dots, S-1. \quad (16)$$

Note that each $E_m^{(S)}$ covers a set of filtered images at different orientations for each scale. Let \vec{f} be, again, the feature vector obtained by using a Pyramid Decomposition with $S = 2$ scales, and $K = 3$ orientations:

$$\vec{f} = \begin{bmatrix} \mu_{00}, \sigma_{00}, \mu_{01}, \sigma_{01}, \mu_{02}, \sigma_{02}; \\ \mu_{10}, \sigma_{10}, \mu_{11}, \sigma_{11}, \mu_{12}, \sigma_{12} \end{bmatrix} \quad (17)$$

By supposing that the dominant scale was found at index $i = 2$ (second scale in the image decomposition), its scale-invariant version, after feature alignment, is defined as:

$$\vec{f}^S = \begin{bmatrix} \mu_{10}, \sigma_{10}, \mu_{11}, \sigma_{11}, \mu_{12}, \sigma_{12}; \\ \mu_{00}, \sigma_{00}, \mu_{01}, \sigma_{01}, \mu_{02}, \sigma_{02} \end{bmatrix} \quad (18)$$

For both rotation-invariant, and scale-invariant representations, the feature alignment is based on the assumption that to compare similarity between texture images, they should be aligned according their dominant orientations/scales.

3.4 Similarity Measure

Similarity between 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 [10]:

$$d(i, j) = \sum_m \sum_n d_{mn}(i, j) \quad (19)$$

where:

$$d_{mn}(i, j) = \left| \frac{\mu_{mn}^i - \mu_{mn}^j}{\alpha(\mu_{mn})} \right| + \left| \frac{\sigma_{mn}^i - \sigma_{mn}^j}{\alpha(\sigma_{mn})} \right|, \quad (20)$$

$\alpha(\mu_{mn})$ and $\alpha(\sigma_{mn})$ denote the standard deviations of the respective features over the entire dataset. They are used for feature normalization purposes.

4 Experimental setup

4.1 Datasets

To evaluate the effectiveness of our approach, we selected thirteen texture images obtained from the standard Brodatz dataset. Before being digitized, each of the 512×512 texture images was rotated at different degrees [12]. Figure 2 displays the non-rotated version of each of the texture images.

To test the rotation-invariance, and scale-invariance of the method, three different image datasets were generated: non-distorted, rotated, and scaled. The *non-distorted* image dataset was constructed just from input textures with no

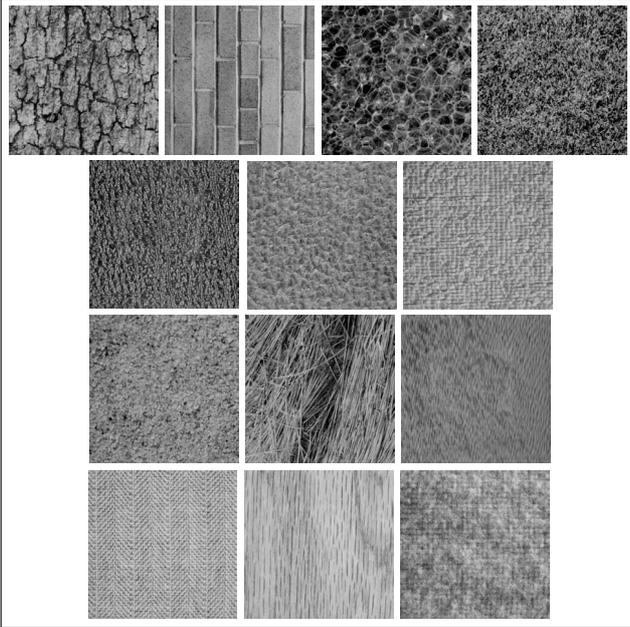


Figure 2. Texture images from the Brodatz dataset used in our experiments. From left to right, and from top to bottom, they include: Bark, Brick, Bubbles, Grass, Leather, Pigskin, Raffia, Sand, Straw, Water, Weave, Wood, and Wool.

rotation and scale changes. Each texture image was partitioned into sixteen 128×128 non-overlapping subimages. Thus, this dataset comprises 208 (13×16) different images. The second image dataset is referred to as *rotated* image dataset, and was generated by selecting the four 128×128 innermost subimages from texture images at 0, 30, 60, and 120 degrees. A total number of 208 images were generated ($13 \times 4 \times 4$). Finally, in the *scaled* image dataset, the 512×512 non-rotated textures were first partitioned into four 256×256 non-overlapping subimages. Each partitioned subimage was further scaled by using four different factors, ranging from 0.6 to 0.9 with 0.1 interval. This led to 208 ($13 \times 4 \times 4$) scaled images.

4.2 Retrieval Effectiveness Evaluation

In our experiments, a simulated query is represented by any of the 208 images in a dataset. The relevant images for each query are defined as the 15 remaining subimages from the same input texture. In this context, a total number of 43056 (207×208) queries were performed in each dataset.

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.

5 Experimental Results

Three series of experiments were conducted to evaluate the retrieval effectiveness our method. In the first ones (Section 5.1), we evaluate the discriminating power of the conventional Steerable Pyramid Decomposition [15] in characterizing texture images, and how its retrieval effectiveness is affected by the presence of scaled and rotated versions of texture patterns. The second and third series of experiments are used to evaluate the rotation-, and scale-invariant properties of our approach (Sections 5.2, 5.3, respectively). Comparisons with the conventional Pyramid Decomposition [15], and with a recent proposal for rotation, and scale-invariance texture retrieval based on Gabor Wavelets [6] are further discussed.

We used a Steerable Pyramid having four orientations ($K = 4$), and two levels of decomposition ($S = 2$), leading thus to a feature vector of 16 elements ($4 \times 2 \times 2$). Our experiments agree with [3] in that the most relevant textural information is contained in those two levels, since little retrieval improvement is achieved by varying the number of scales from two to three levels during image decomposition. Furthermore, our motivations in using small size feature vectors are: (1) to show that the retrieval effectiveness of our approach is not compromised, and (2) to facilitate image retrieval applications where data storage capacity is a limitation.

5.1 Effectiveness of conventional Steerable Pyramid Feature Representation

The retrieval effectiveness of the conventional Steerable Pyramid is shown in Figure 3. It compares, for each class of texture, the average rate of retrieving the relevant images for the *non-distorted* image dataset, the *rotated* image dataset *without* rotation-invariant representation, and the *scaled* image dataset *without* scale-invariant representation.

In the case of the *non-distorted* dataset, it can be noticed that the Steerable Pyramid presents good retrieval accuracy for almost all classes. However, the effectiveness is low for textures that have, either no strong direction (Bark, Bubbles), or for textures whose subimages present visual dissimilarities among each other (Brick). Additionally, we can see that for the *rotated* image dataset, the retrieval effectiveness decays rapidly for images with a defined direction (Brick, Leather, Pigskin, Raffia, Straw, Water, Wood, Wool). This happens because feature coefficients of self-similar subimages are rotated (shifted). Moreover, note that, in the *scaled* image dataset, the retrieval effectiveness is less affected than in the case of the *rotated* image dataset. However, images whose fine patterns become more visible,

when increasing the scale factor, are most affected (Leather, Raffia, Sand, and Weave). In this case, the “zoomed” micro-patterns produce enough discriminatory information, so that images belonging to the same texture class cannot be judged as being similar.

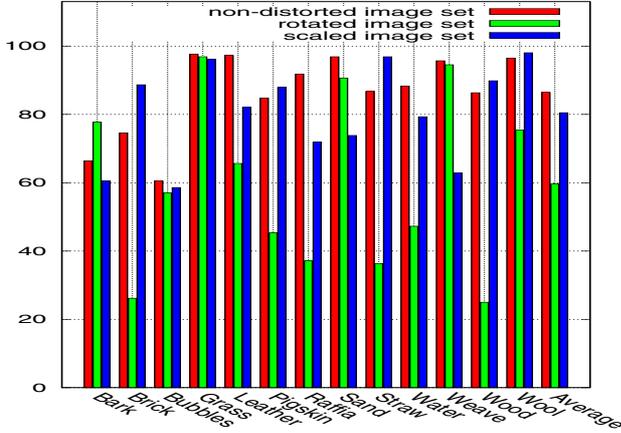


Figure 3. Average retrieval rates for *non-distorted* image dataset, *rotated* image dataset without rotation-invariant representation, and the *scaled* image dataset without scale-invariant representation using the conventional Steerable Pyramid Decomposition having $S = 2$ scales and $K = 4$ orientations.

5.2 Effectiveness of Rotation Invariance Representation

To show both the improvement over the conventional Steerable Pyramid, and the capability of our method for characterizing rotated texture images, we used the *rotated* image dataset. From Figure 4, we can see that for almost all classes the retrieval accuracy was dramatically increased. This observation is more notorious in images having strong direction (Brick, Leather, Pigskin, Raffia, Straw, Water, Wood), since our feature alignment considers the dominant orientation in the images. Further, Figure 5 shows that our approach outperforms the recent proposal for rotation-invariant characterization using based on Gabor Wavelets [6]. However, the latter method presents better retrieval accuracy for images that have a uniform scattered region, i.e, some unique primitives appear at any location in the textures, including, for example, image borders (Water, Wood). In this case, Gabor Wavelets perform better, since at least one filter in the filterbank covers that subregion.

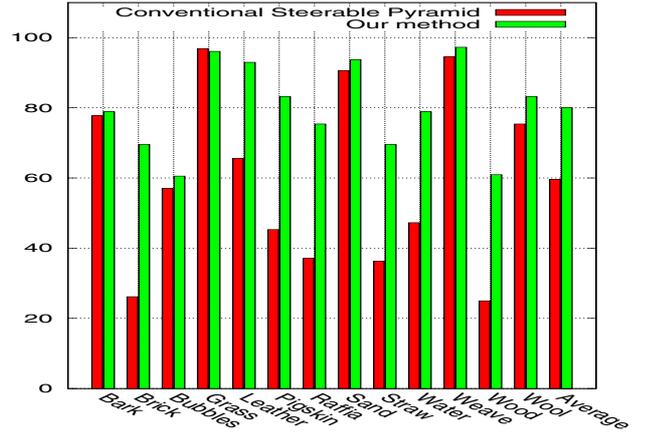


Figure 4. Average retrieval rates per texture class for *rotated* image dataset without and with rotation-invariant representation, respectively.

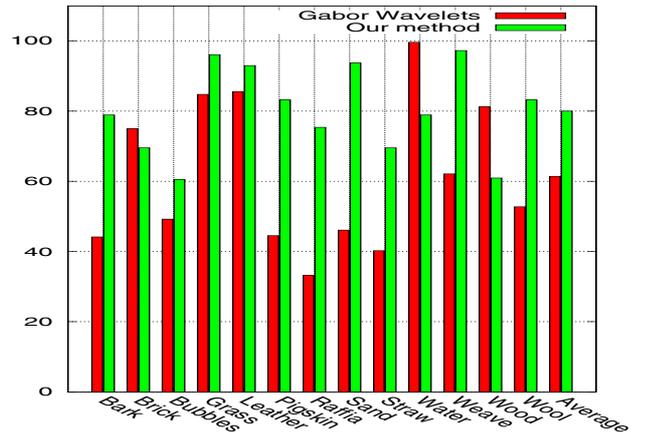


Figure 5. Comparison of rotation-invariant retrieval performances for the *rotated* image dataset using the recent Gabor-based approach [6], and our approach.

5.3 Effectiveness of Scale Invariance Representation

In this subsection, our approach is compared with the conventional Steerable Pyramid Decomposition, and Gabor Wavelets for the *scaled* image dataset (Figures 6, 7). From Figure 6, we can notice that the highest retrieval accuracy improvement is achieved in textures characterized by the presence of micro-patterns (Pigskin, Sand, Weave). By us-

ing different levels of scale decompositions, these micro-patterns are highlighted, and therefore are more distinctive even for textures belonging to the same texture class. Furthermore, feature vectors are aligned according to the dominant scale in texture images, improving the effectiveness of the retrieval process.

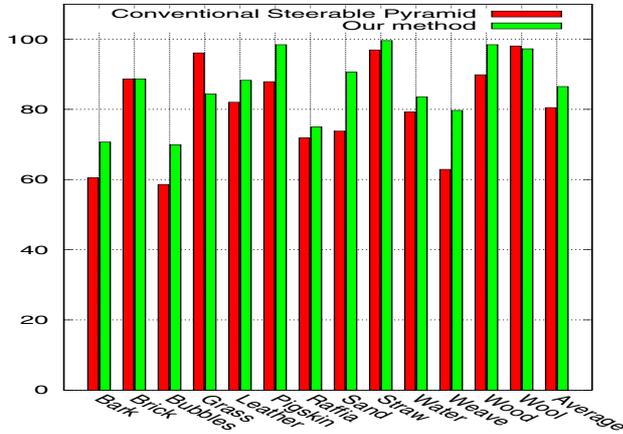


Figure 6. Average retrieval rates per texture class for *scaled* image dataset without and with rotation-invariant representation, respectively.

In Figure 7, we compare the retrieval effectiveness of our approach with the one using the modified Gabor Wavelets. Our method yields better retrieval effectiveness for all classes except for the Leather texture pattern.

Note that more discriminatory information is obtained by downsampling rows and columns (Steerable Pyramid approach) than by using filterbanks at few scales (Gabor Wavelets approach).

5.4 Results Summarization

A summary of our experimental results is provided in Table 1. It compares the average rates of retrieving the 15 most similar images for each of the input textures in the *rotated*, and *scaled* image datasets. We can see that our method (M-SPyd) improves retrieval rates over the conventional Steerable Pyramid (C-SPyd) and Gabor Wavelets (GWs) on both image datasets. Note that the retrieval effectiveness was improved in almost 20% for the rotated image dataset, whereas for the scaled image dataset it was improved from 80.5% to 86.51% over the conventional Steerable Pyramid.

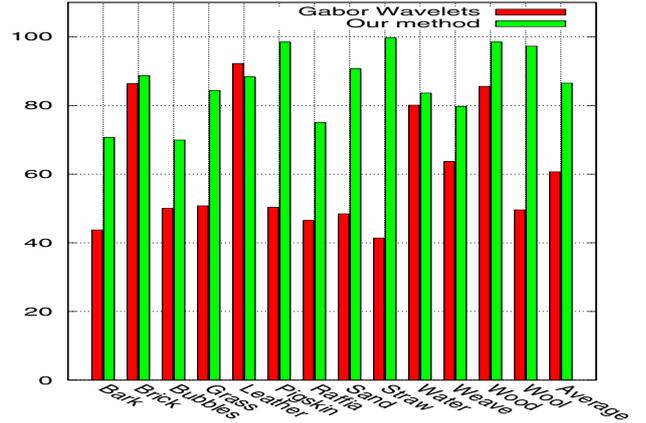


Figure 7. Comparison of scale-invariant retrieval performances for the *scaled* image dataset using the recent Gabor-based approach [6], and our approach.

	M-SPyd	C-SPyd	GWs
Rotated image dataset	80.02%	59.62%	61.42%
Scaled image dataset	86.51%	80.5%	60.7%
Feature dimensionality	16	16	16

Table 1. Average retrieval rates and feature dimensionalities for the methods used in our experiments.

5.5 Image Retrieval Examples

Figure 8 displays four retrieval examples of the three evaluated methods. The two first queries were selected from the rotated-image datasets. They correspond to the classes for which our method and the Gabor Wavelets achieved the highest retrieval accuracy (Raffia and Water, respectively). Similarly, the remaining two image queries were selected to illustrate the scale-invariance capabilities. They belong to the Sand and Leather classes.

The input queries are shown on the first columns, and their relevant retrieved images are displayed in increasing order on the following columns. Furthermore, for each query image, the retrieved relevant images shown on the first, second, and third rows, are obtained by considering respectively the conventional Steerable Pyramid, our approach, and Gabor Wavelets.

From the retrieved images, we can see that our method outperforms the conventional Steerable Pyramid for all query images. Additionally, for the cases where Gabor Wavelets achieved highest retrieval accuracy (second, and fourth queries), our method retrieved almost all relevant images.

6 Conclusions

In this paper, a new method for facilitating rotation-invariant, and scale-invariant image retrieval applications was introduced. Our approach exploits the discriminability properties of the Steerable Pyramid Decomposition for texture characterization, and by taking into account either the **dominant orientation** or **dominant scale** of the input textures, a new feature descriptor is proposed.

We performed a total number of 43056 different queries in three image datasets. Experimental results obtained in those datasets are very encouraging, since the new method demonstrated retrieval rates improvements for both rotated, and scaled image datasets. More specifically, the rotation-invariant and scale-invariant average retrieval rates of our method are increased respectively from 59.62% to 80.02% and from 80.5% to 86.51% over the conventional Steerable Pyramid. It is worth mentioning, that the average retrieval rates of Gabor Wavelets presented less retrieval accuracy being 61.42% and 60.7% for the rotated and scaled image datasets, respectively.

Future work on this matter will include extending this method for both rotated and scaled image characterization. Furthermore, we plan to use this approach for problems concerning texture segmentation and recognition.

7 Acknowledgments

Javier Montoya would like to thank professor Serge Belongie for his valuable discussions and suggestions, and is also grateful to CNPq for the financial support (134990/2005-6). This work is partially supported by Microsoft Escience and Tablet PC Technology and Higher Education projects, CNPq (Proc. 477039/2006-5 and 311309/2006-2), FAPESP, CAPES, and FAEPEX.

References

- [1] S. Arivazhagan, L. Ganesan, and S. P. Priyal. Texture classification using gabor wavelets based rotation invariant features. *Pattern Recognition Letters*, 27:1976–1982, 2006.
- [2] A. del Bimbo. *Visual Information Retrieval*. Morgan Kaufmann, 1st. edition, 1999.
- [3] M. N. Do and M. Vetterli. Wavelet-based texture retrieval using generalized gaussian density and kullback-leibler distance. *IEEE Transactions on Image Processing*, 11(2):146–158, 2002.
- [4] W. T. Freeman and E. H. Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(9):891–906, 1991.
- [5] H. Greenspan, S. Belongie, R. Goodman, and P. Perona. Rotation invariant texture recognition using a steerable pyramid. *Proceedings of the 12th IAPR. International Conference on Pattern Recognition*, 2:162–167, 1994.
- [6] J. Han and K. K. Ma. Rotation-invariant and scale-invariant gabor features for texture image retrieval. *Journal of Vision and Image Computing*, 2006. Article in Press (Available online 28 December 2006).
- [7] P. W. Huang, S. K. Dai, and P. L. Lin. Texture image retrieval and image segmentation using composite sub-band gradient vectors. *J. Visual Communication and Image Representation*, 17(5):947–957, 2006.
- [8] M. Kokare, B. N. Chatterji, and P. K. Biswas. Cosine-modulated wavelet based texture features for content-based image retrieval. *Pattern Recognition Letters*, 25(4):391–398, 2004.
- [9] A. Laine and J. Fan. Texture classification by wavelet packet signatures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1186–1191, 1993.
- [10] B. S. Manjunath and W. Y. Ma. Texture features for browsing and retrieval of image data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(8):837–842, 1996.
- [11] R. Manthalkar, P. K. Biswas, and B. N. Chatterji. Rotation invariant texture classification using even symmetric gabor filters. *Pattern Recognition Letters*, 24(12):2061–2068, 2003.
- [12] U. of Southern California Signal and I. P. Institute. Rotated textures. <http://sipi.usc.edu/services/database/Database.html> (Accessed on 1 March 2007).
- [13] 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.
- [14] T. Randen and J. H. Husoy. Filtering for texture classification: A comparative study. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(4):291–310, 1999.
- [15] E. P. Simoncelli and W. T. Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. *Proceedings of IEEE International Conference on Image Processing*, 13(9):891–906, 1995.

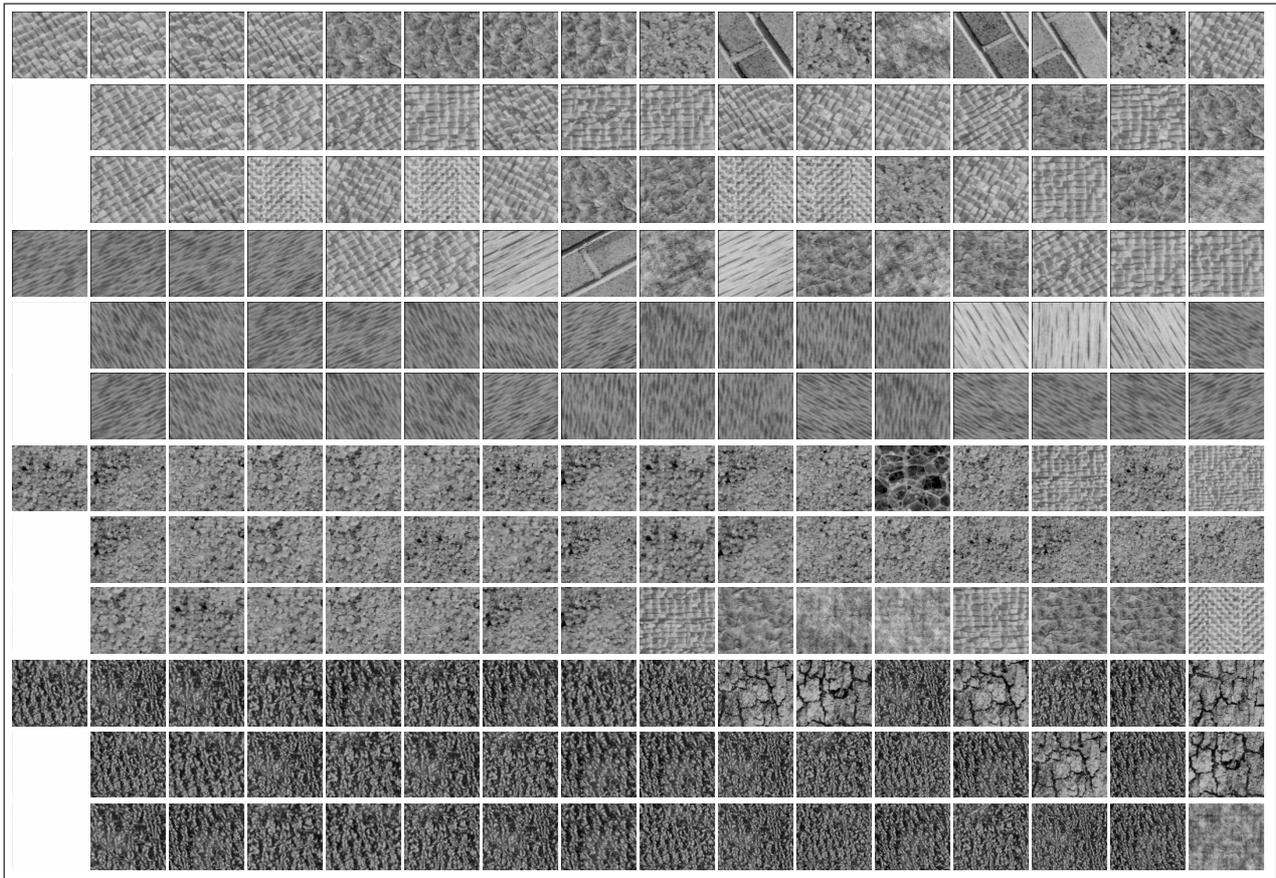


Figure 8. Image retrieval examples for each of the three methods: conventional Steerable Pyramid, our approach, and Gabor Wavelets (first, second, and third rows, respectively). Each query image is shown on the first column, and its 15 top matches are displayed in increasing order.