

A Study of Zernike Invariants for Content-Based Image Retrieval

Pablo Toharia¹, Oscar D. Robles¹, Ángel Rodríguez², and Luis Pastor¹

¹ Dpto. de Arquitectura y Tecnología de Computadores, Ciencias de la Computación e Inteligencia Artificial,

U. Rey Juan Carlos, C/ Tulipán, s/n. 28933 Móstoles. Madrid. Spain
{pablo.toharia,oscardavid.robles,luis.pastor}@urjc.es

² Dpto. de Tecnología Fotónica, U. Politécnica de Madrid,
Campus de Montegancedo s/n, 28660 Boadilla del Monte, Madrid, Spain
arodri@dtf.fi.upm.es

Abstract. This paper presents a study about the application of Zernike invariants to content-based Image Retrieval for 2D color images. Zernike invariants have been chosen because of their good performance for object recognition. Taking into account the good results achieved in previous CBIR experiments with color based primitives using a multiresolution representation of the visual contents, this paper presents the application of a wavelet transform to the images in order to obtain a multiresolution representation of the shape based features studied. Experiments have been performed using two databases: the first one is a small self-made 2D color database formed by 298 RGB images and a test set with 1655 query images that has been used for preliminary tests; the second one is Also experiments using the Amsterdam Library of Object Images (ALOI), a free access database. Experimental results show the feasibility of this new approach.

Keywords: CBIR primitives, Zernike invariants.

1 Introduction

Content-Based Image Retrieval (CBIR) systems are becoming a very useful tool when dealing with a large volume of visual information, due to the maturity level of the techniques involved and posed nowadays and in the past by the research community [1,2,3,4]. Most of these techniques are inherited from computer vision and from database systems to represent and manage the available data. CBIR systems can be classified following different criteria, such as the nature of primitives used for characterizing the image's contents (color, texture, shape, scheme or attributed graphs, etc.), the abstraction levels covered by these primitives (low, medium high), the automation level achieved in the primitive extraction process (automatic, semi-automatic, manual), the classifier used for the retrieval stage (standard metrics, neural networks, SVM, etc.) or the way data are stored, processed and managed in the whole CBIR system (centralized, distributed). All

these topics are different research areas in which continuous advances are made in order to improve the performance of CBIR systems.

In this paper, several new shape primitives to be used in a CBIR system are implemented. The computation of Zernike invariants has been chosen because of their good performance achieved in object recognition applications. It has been a traditional path to follow to adapt primitives used in recognition systems to CBIR ones, since some of the techniques used in the latter have been inherited from the former. However, it must be noticed that the working domain of recognition systems, i. e. the set of objects that the system is able to recognize, is usually very restricted. To our knowledge, the only works in the CBIR field that have used the Zernike Invariants have been Novotni and Klein, applying 3D Zernike Descriptors for 3D object retrieval [5,6], Kim *et al.*, who work with a restricted subject database of binarized images [7], and Lin and Chou [8] who work with a reduced database of color images but only report results about computational cost, not including any measure of recall or precision. On the other hand, the main drawbacks of Zernikes moments are both the computational cost and the approximation errors. Lately, there have been a few works trying to solve or minimize this disadvantages [9,10,11,12].

Based on the good results achieved in previous CBIR experiments with color based primitives using a multiresolution representation of the visual contents [13], in this paper is also studied the application of a wavelet transform to obtain a multiresolution representation of the shape based features searched. The development of wavelet transform theory has spurred new interest in multiresolution methods and has provided a more rigorous mathematical framework. Wavelets give the possibility of computing compact representations of functions or data. Additionally, they allow variable degrees of detail or resolution to be achieved, as well as are attractive from the computational point of view [14,15]. Analysis and diagonal detail coefficients of the image's wavelet transform have been used. The resulting primitive is a feature vector composed of the values computed at each resolution level of the transformed image, obtaining a more robust representation of the image contents.

Section 2 describes some background about the Haar transform, the wavelet transform used in the work herein presented, and the different feature extraction techniques implemented. Section 3 analyzes some implementation details of the implemented primitives. Section 4 shows the experiments performed, presenting the success rate achieved for each primitive as well as discusses these results. Finally, conclusions are presented in Section 5.

2 Description of the Implemented Primitives

2.1 The Haar Transform

Wavelet transforms can be seen as a reformalization of the multiresolution methods of the 80's [16,17]. The information they provide is very similar to that obtained from Fourier-based techniques, with the advantage of working with local

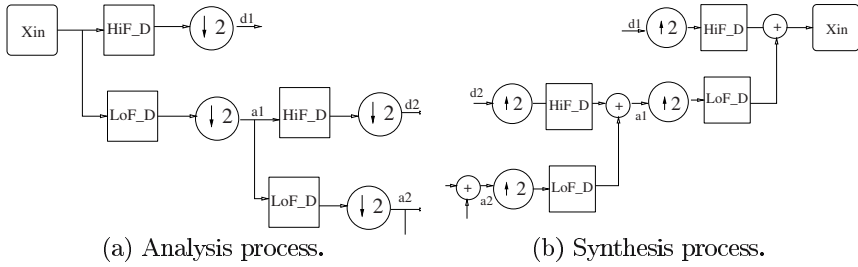


Fig. 1. Scheme of a wavelet transform

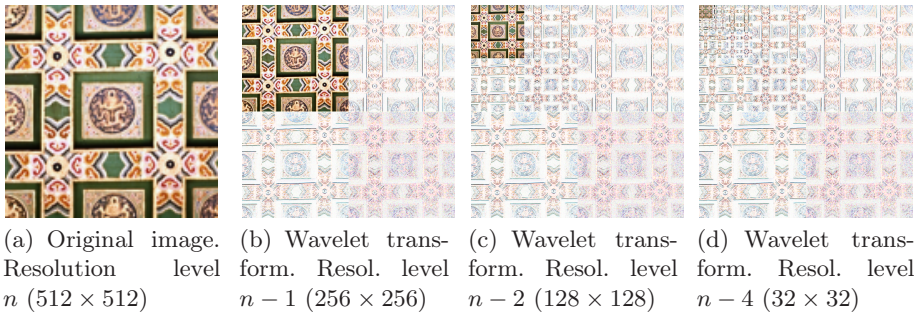


Fig. 2. Non-standard Haar transform of a 2D color image

information using base functions with compact support. Wavelet transform coefficients show variations in object features at different resolution or scale levels [14,18,19,20]. Roughly speaking, the detail coefficients of the wavelet transform can be considered as a high-frequency extraction process of the objects appearing on the images, while the analysis coefficients behave complementary: the lower the resolution level, the more homogeneous the regions they produce. This could be equivalent to a successive application of low-pass filters with a signal subsampling operation (see Fig. 1(a)). The inverse of this process allows the reconstruction of the original signal by the so-called synthesis process (Fig. 1(b)). The result of this transform is a series of images where four regions of coefficients are obtained at each step: analysis coefficients (top-left, Fig. 2(b)), vertical detail coefficients (top-right, Fig. 2(b)), horizontal detail coefficients (bottom-left, Fig. 2(b)) and diagonal detail coefficients (bottom-right, Fig. 2(b)).

The Haar transform will be used as a tool to extract some features of the transformed image that will allow a discrimination process to be performed between the queries and the images stored in the information system. The low order complexity of this transform, $O(n)$, allows an efficient implementation of the whole process.

2.2 Zernike Invariants

Zernike invariants have been selected because of its demonstrated good performance in object recognition problems [21,22].

In 1934, Zernike [23] presented a set of complex polynomials $V_{nm}(x, y)$ that were defined inside a unity radius circle ($x^2 + y^2 \leq 1$) in the following way:

$$V_{nm}(x, y) = V_{nm}(\rho, \theta) = R_{nm}(\rho) e^{jm\theta} \tag{1}$$

where V_{nm} is a complete set of complex polynomials, n is a positive integer value $n \geq 0$ that represents the polynomial degree and m is the angular dependency, ρ and θ are the polar coordinates of the Cartesian coordinates (x, y) and R_{nm} is a set of radial polynomials that have the property of being orthogonal inside the unity circumference. The values of n and m have the following relation:

$$(n - |m|) \bmod 2 = 0 \quad \text{and} \quad |m| \leq n \tag{2}$$

and the radial polynomials have the following expression:

$$R_{nm}(\rho) = \sum_{s=0}^{\frac{m-|n|}{2}} (-1)^s \frac{(m-s)!}{s! (\frac{m+|n|}{2} - s)! (\frac{m-|n|}{2} - s)!} \rho^{m-2s} \tag{3}$$

Starting from Zernike polynomials and projecting the function over the orthogonal basis composed by the polynomials, the moments can be generated in the following way:

$$A_{mn} = \frac{m+1}{\pi} \iint_{x^2+y^2 \leq 1} f(x, y) V_{nm}^*(x, y) dx dy \quad \text{with} \quad x^2 + y^2 \leq 1 \tag{4}$$

The discretization needed to work with digital data can be done straightforwardly:

$$A_{mn} = \frac{m+1}{\pi} \sum_x \sum_y f(x, y) V_{nm}^*(x, y) dx dy \quad \text{with} \quad x^2 + y^2 \leq 1 \tag{5}$$

Figure 3 shows the reconstruction of two example images using Zernike moments of several orders.

From these functions, we compute the modulus to obtain the p different invariant values for each considered case. The invariant values are used to create a vector of p elements ZI_i that collect the shape information of an image i . For example, in the case of polynomials of 10^{th} degree, p would be 36.

2.3 Signature Based on Zernike Invariants

The visual contents of the images are transformed into a vector of some features, named signature, that aims to collect some discriminant information of the original data.

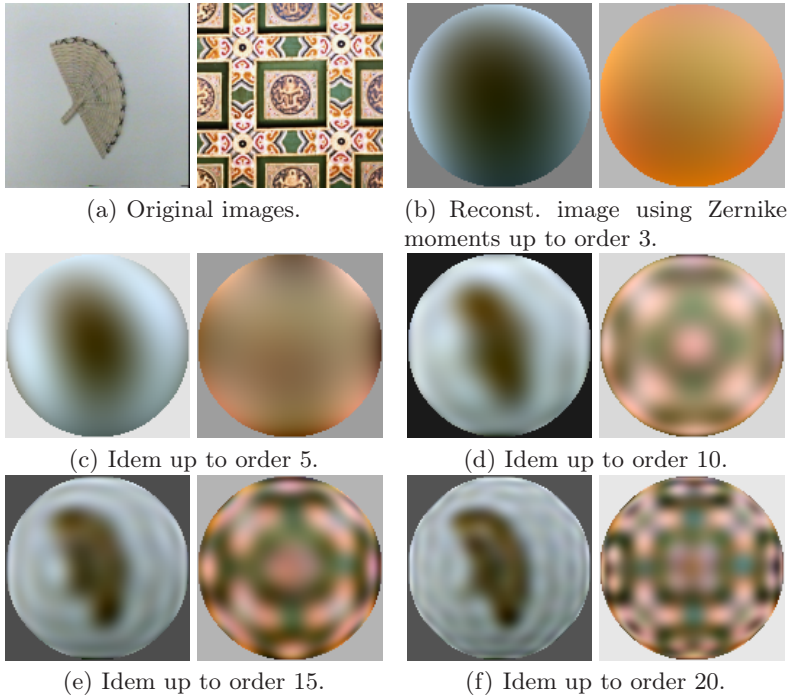


Fig. 3. Original images and reconstructions obtained working with Zernike moments of orders 3, 5, 10, 15 and 20

Firstly, a primitive based on the Zernike Invariants extracted from the original image has been implemented, in order to compare the multiresolution primitives studied in this paper with the results achieved without applying the Haar transform, as it has been traditionally done in object recognition environments.

Also, the invariant values computed from polynomials of several orders have been tested. The signature is generated concatenating the invariants extracted until the maximum polynomial degree considered. A simple vector of scalar values is obtained.

2.4 Signature Based on Zernike Invariants over Analysis Coefficients

The analysis coefficients of the Haar transform have been selected to compute the Zernike invariants over different scales of the original image. Analysis coefficients retain a coarse description of the objects' shape as can be observed in Fig. 2. Once the sequence of lower resolution images is obtained, we compute the Zernike Invariants over the analysis coefficient region at each resolution level. The final step is to compose a vector as explained in the previous section, so as to collect the original image's multiresolution shape information.




Fig. 4. Inscribing Zernike's polynomial definition domain into rectangular images

2.5 Signature Based on Zernike Invariants over Analysis and Diagonal Detail Coefficients

The choice of the diagonal coefficients is due to the fact that the diagonal region, right bottom squares in Figure 2 at each resolution level, gathers high-pass features on both vertical and horizontal directions. This feature should be an additional discriminant element when dealing with images of similar appearance.

3 Implementation Analysis

As it can be deduced from Eqs. 3 and 5, the computation of invariants is a very high demanding task from a computational point of view, so different order polynomials have been tested in order to verify if their responses are significantly different. The only Zernike invariants computed online are those of the query image, so this approach is completely feasible in a real-world CBIR system.

Another issue faced in the implementation is mapping the images' rectangular domain to the circular space where radial polynomials are defined (Eq. 1). The unity radius circle has been inscribed into the image, so its corners have been discarded under the assumption that they do not usually contain relevant information about the scene (Fig. 4).

When working with gray-level images, the original image's wavelet transform gives all the information needed to compose the signature representing the image's contents. Joining together Zernike invariant values for all the considered regions and all the resolutions levels, a vector of features collecting the intensity of the input data is obtained. When dealing with color images, as it is the case in the RGB color space, the process described for the monochrome images must be applied over each one of the color channels.

4 Experimental Results

The main objectives of the tests are to measure and analyze the implemented features recall and precision values achieved by the implemented features. The

classical definitions of recall and precision have been used:

$$\text{Recall} = \frac{\text{True positives}}{\text{True positives} + \text{False negatives}} \quad (6)$$

$$\text{Precision} = \frac{\text{True positives}}{\text{True positives} + \text{False positives}} \quad (7)$$

For the experiments presented in this paper two different databases have been considered: our self-made color database and a public one.

4.1 Self-made Color Database

Experiments setup. For computational purposes, instead of using a larger database as ALOI is, a small one previously available has been used. This database has been used for preliminary tests in order to extract some conclusions about the influence of the moments' order considered in the Zernike invariants extracted for obtaining the signature. The same color database described in [24] has been used. It is formed by 298 RGB two-dimension 128x128 images collected from different sources, like [25] or Internet. The test set was generated introducing images which share the same concept but look quite different from those of the considered database and making affine transformations or selecting regions of interest from the original database images. Following these guidelines, a test set formed by 1655 two-dimension color images was obtained.

The experiments have consisted in querying all the images from the test set in order to retrieve the associated image stored in the corresponding database, which was selected as the representative image of its class in the query. For each input image, the result from the search is a list sorted according to the level of similarity of its signature with the signatures stored in the database. The minimum distance classifier used is based on the Euclidean distance.

The aforementioned list contains the best n matches of the query. Figure 5 presents an example of the system response on a query to the TRECVID database [26]. The most similar object appears on the upper-left corner of the mosaic. In the example presented, the topic searched was "shots of one or more soccer goal posts".

Several tests have been performed considering all the parameters involved in the Zernike invariants combined with the multiresolution representation described above. The influence of Zernike invariants order on the retrieval results has been studied due to the great computational load demand that involves their computation. Orders 3, 5, 10, 15 and 20 have been considered. It must be noticed that they respectively imply 6, 12, 36, 72 and 121 first invariants. The behaviour of the wavelet coefficients over the proposed shape primitive has been considered, taking into account the following configurations: analysis coefficients, analysis plus diagonal detail coefficients.

Results analysis. Table 1 shows recall values computed for all the primitives herein described using the Euclidean distance as minimum distance classifier.

The suffix notation used in the method column in Table 1 is the following:

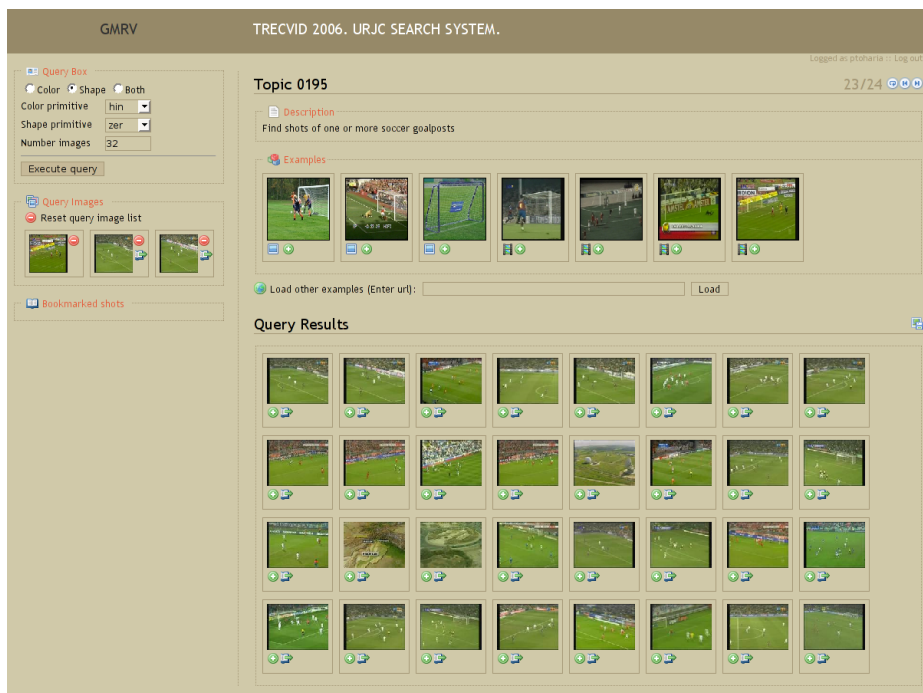


Fig. 5. Example of the output of the CBIR implemented system

Table 1. Recall values achieved over the whole set of Zernike based primitives

Primitive	Recall	Primitive	Recall	Primitive	Recall
ZER_3	0.480	ZMR_CA_3	0.517	ZMR_CA_DD_3	0.517
ZER_5	0.471	ZMR_CA_5	0.514	ZMR_CA_DD_5	0.514
ZER_10	0.474	ZMR_CA_10	0.510	ZMR_CA_DD_10	0.511
ZER_15	0.460	ZMR_CA_15	0.515	ZMR_CA_DD_15	0.515
ZER_20	0.460	ZMR_CA_20	0.490	ZMR_CA_DD_20	0.503

- ZER: Zernike invariants.
- ZMR: Zernike invariants with multiresolution information.
- CA: Analysis coefficients of the Haar transform.
- CA_DD: Analysis and diagonal detail coefficients of the Haar transform.
- number: order of the Zernike moments used.

As it can be seen, several order values for the Zernike Invariants computation have been tested.

An analysis of Table 1 shows how the obtained results are not influenced by the order of Zernike moments since the recall value is quite similar in all cases. It can be seen that the best results are achieved using invariants computed with

Table 2. Basic characteristics of the ALOI's subsets ILL, COL and VIEW

Subset	Variations	Images
ILL	24 illumination variations	24000
COL	12 illumination variations	12000
VIEW	72 illumination variations	72000

low order moments. The reason for this might be the more significant information carried by low order moments (up to order 15) during the reconstruction process, as it can be observed in Figure 3. In that case, those low order moments would be basically determining the main difference between two images. On the other hand, the use of multiresolution approach improves the information provided by the Zernike invariants computed over the original image, although this improvement is not as remarkable as the one obtained when the multiresolution approach has been applied to color primitives [13]. The explanation for this behaviour comes from the fact that low resolution levels are not sharp enough to significantly contribute to the discriminant power of the primitive.

4.2 ALOI Database

Experiments setup. Taking into account the results obtained using our self-made database a new experiments setup with the Amsterdam Library of Object Images (ALOI) [27] database was established. It will allow to do a deeper study about the behaviour of Zernike-based primitives as well as to obtain more exhaustive results.

The ALOI database is a color image collection of 1000 small objects. This database is divided in different subsets depending on the parameter varied at capture time (color, illumination, view-angle, etc.). From the available subsets, the main three of them have been chosen for our purposes: illumination direction (ILL), illumination color (COL), object viewpoint (VIEW). Table 2 shows some basic information about these subsets. More detailed information and examples of the captures can be found on [27].

The ALOI database does not provide a separated query image set. For this reason, the experiments have consisted in using each image to query a target database formed by the rest of the images. In order to evaluate recall and precision in this scenario the number of retrieved images has to be fixed. Different values for this number have been evaluated in order to obtain the behaviour tendency of each primitive.

Another aspect to be taken into account is that the three subsets (COL, ILL and VIEW) have different number of relevant images, so results from different subsets can not be directly compared. For example, retrieving 100 images from VIEW could have 72 relevant images, while retrieving the same number of images from COL could at most have 12 relevant ones. Under this scenario, recall and precision values would be biased. Since the number of relevant images per object in each subset is well known, a percentage of this number can be used instead of the number of retrieved images so as to be able to do that comparison.

Considering the previous experiments over the self-made database, we have selected order 15 Zernike moments to build the Zernike invariants primitive run in the experiments over the ALOI database. From our point of view, they offer a promising trade-off between discriminant power and computational cost that will allow to extract a better performance to the multiresolution approach.

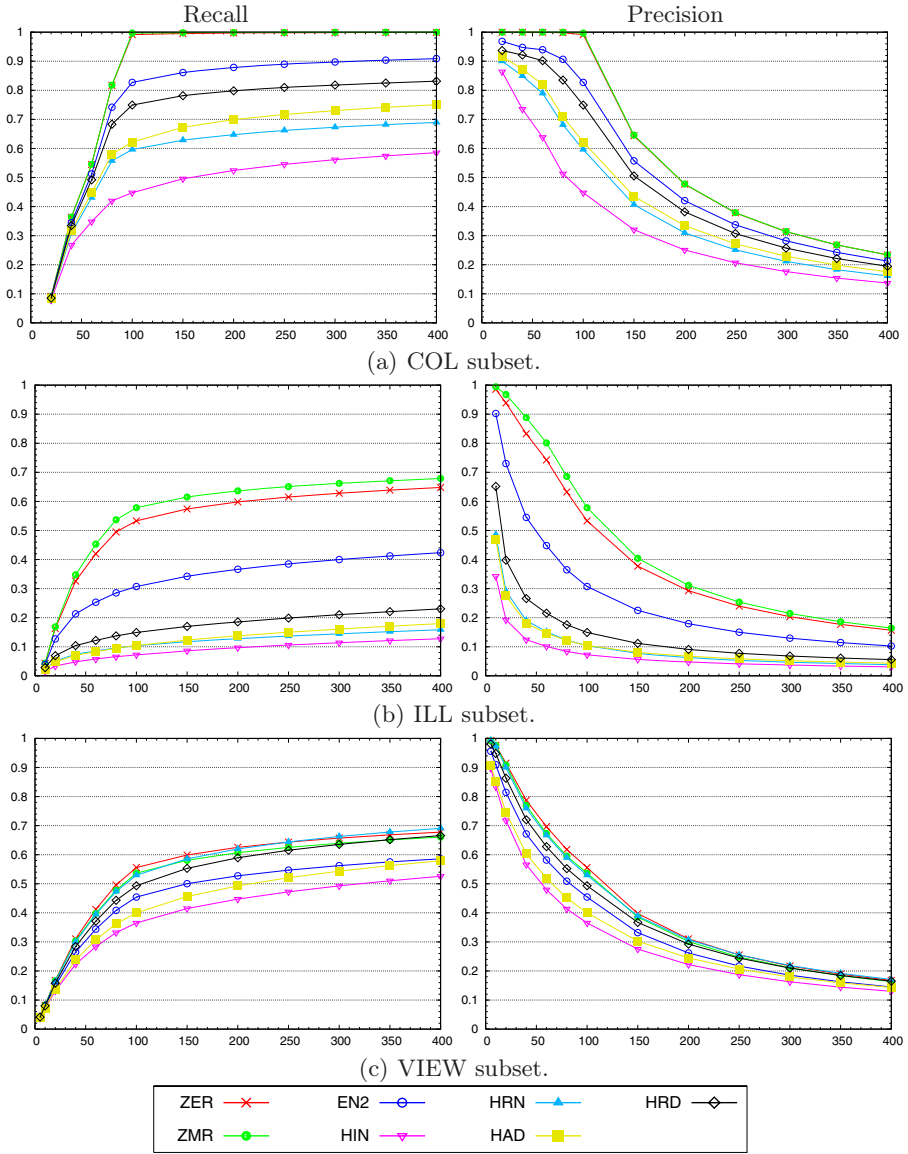


Fig. 6. Precision (right column) and recall (left column) against number of retrieved images (abscissa axis) over the ALOI database

Zernike-based primitives have been compared to some other multiresolution primitives developed and tested in previous works:

- Multiresolution energy-based primitive (EN2) [24].
- Multiresolution histogram primitives [13]:
 - Global multiresolution histogram over analysis coefficients (HIN).
 - Local multiresolution histogram over analysis coefficients (HRN).
 - Global multiresolution histogram over analysis and detail coefficients (HAD).
 - Local multiresolution histogram over analysis and detail coefficients (HRD).

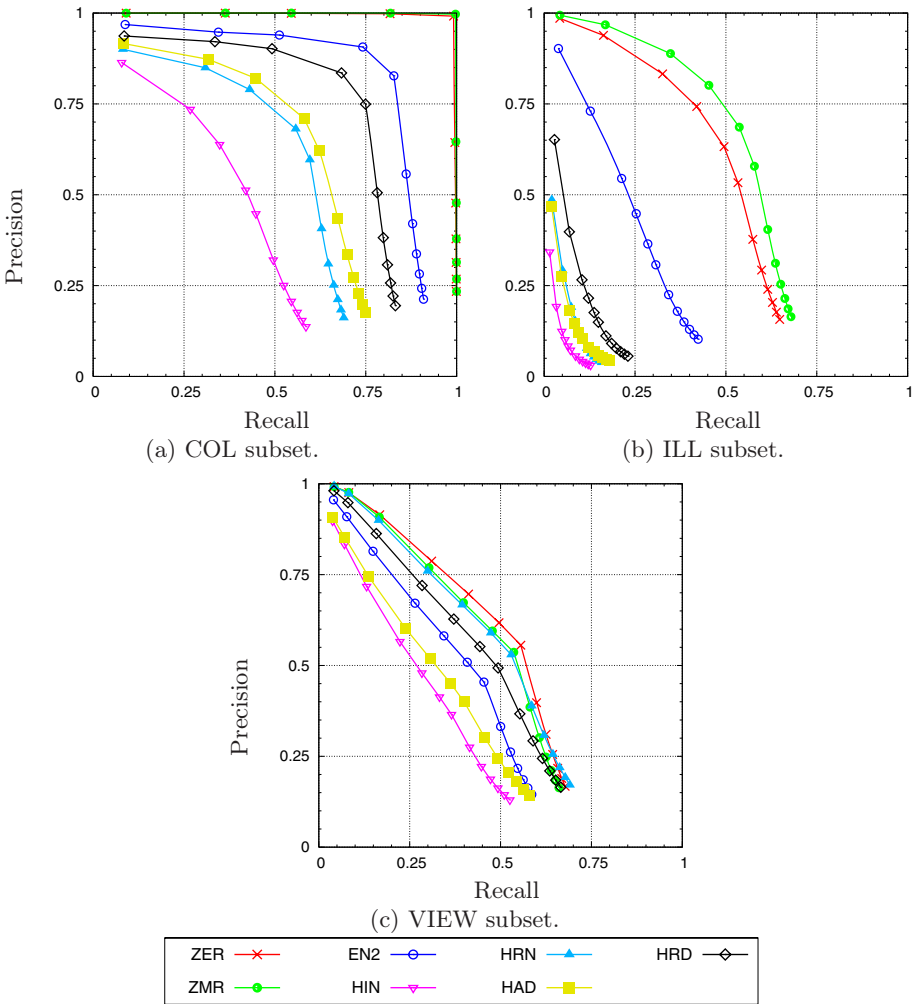


Fig. 7. Precision vs. recall over the ALOI database

Results analysis. Figure 6 presents the achieved precision and recall values for the three different ALOI subsets (COL, ILL, VIEW). The abscissa axis shows the number of images taken into account to calculate the precision and recall values. To be more precise each abscissa value stands for a percentage of retrieved images compared to the *a priori* known number of relevant images, i.e., 100 means 12 images for COL subset, 24 for ILL and 72 for VIEW. This is done in order to be able to compare the results among subsets.

Figure 7 shows the same data than Figure 6 but in this case the recall and precision values are depicted one against the other. This graph allows to see the behavior of the precision value together with the behavior of the recall value.

Results show that shape-based primitives, ZER and ZMR, outperforms the rest of primitives when using this database. In particular these two primitives achieve excellent results when dealing with the COL subset data. It can be seen that their precision and recall graphs shown in Figure 6(a) are almost the ideal graphs, achieving recall and precision values of 1 when retrieving 12 images (that is, a 100% of relevant images per object). On the other hand the performance of these primitives using the other two subsets (ILL and VIEW) are also good (ZMR and ZER are the best results for ILL and VIEW respectively) but not as much as the values achieved when using COL. This can be explained easily taking into account the fact that the variation of the color properties does not change the perception of the captured object's shape, leading to a very good performance. However, when using ILL and VIEW subsets, images' shape changes because of the shadows produced by the variation of the illumination direction, what makes shapes to appear different. A viewpoint variation has the same effect if the object is not "symmetric".

Comparing the results achieved by the multiresolution shape approach (ZMR) against the non-multiresolution (ZER) it can be stated that the results achieved are quite similar except in the ILL subset in which the ZMR performs slightly better.

5 Conclusions and Ongoing Work

The application of a shape primitive to a content-based image retrieval system using a multiresolution approach has been studied over two different databases.

The shape primitive based on Zernike invariants has achieved interesting results using low level order polynomials, diminishing the main drawback of this technique: its demanding computational cost. In the particular case of dealing with color changes the Zernike-based primitive has shown an excellent performance. The use of multiresolution information in the Zernike based primitive has improved the results in some cases but has never given a big increase with the databases tested.

It has been shown that increasing the order of the moments used in the experiments does not guarantee an image with enough detail for discriminating purposes, as can be deduced from the examples presented in Fig. 3.

Using Zernike moments up to order 20 on a database composed of images related to so different topics and themes has not improved the results in comparison to low order moments. It must be noticed that the examples found in the available bibliography apply Zernikes to a set of images in a restricted domain or do not report any results about recall but from a computational cost point of view.

Further research on multiresolution primitives will allow, on one hand, the use of additional wavelet basis to achieve a more compact representation of the multilevel information without losing discriminant capacity. On the other hand, it will contribute to an improvement on the behavior of the Zernike based primitives. To conclude, the use of alternatives for combining both primitives will also be studied in the near future.

Acknowledgments

This work has been partially funded by the the Spanish Ministry of Education and Science (grant TIN2007-67188) and Government of the Community of Madrid (grant S-0505/DPI/0235; GATARVISA).

References

1. del Bimbo, A.: Visual Information Retrieval. Morgan Kaufmann Publishers, San Francisco, California (1999)
2. Venters, C.C., Cooper, M.: A review of content-based image retrieval systems. Technical report, Manchester Visualization Center. Manchester Computing. University of Manchester (2000), <http://www.jtap.ac.uk/>.
3. Marques, O., Furht, B.: Content-based Image and Video Retrieval. Multimedia Systems and Application Series. Kluwer Academic Publishers, Dordrecht (2002)
4. Wu, J.-K., Kankanhalli, M.S., Wu, K.W.J.K., Lim, L.J.-H., Hong, H.D.: Perspectives on Content-Based Multimedia Systems. Springer, Heidelberg (2000)
5. Novotni, M., Klein, R.: 3D zernike descriptors for content based shape retrieval. In: The 8th ACM Symposium on Solid Modeling and Applications (2003)
6. Novotni, M., Klein, R.: Shape retrieval using 3D zernike descriptors. *Computer-Aided Design* 36(11), 1047–1062 (2004)
7. Kim, H.K., Kim, J.D., Sim, D.G., Oh, D.I.: A modified zernike moment shape descriptor invariant to translation, rotation and scale for similarity-based image retrieval. In: International Conference on Multimedia and Expo, ICME, vol. 1, pp. 307–310 (2000)
8. Lin, T.W., Chou, Y.F.: A comparative study of zernike moments. In: Proceedings of the IEEE/WIC International Conference on Web Intelligence (WI' 2003), Halifax, Canada (2003)
9. Hwang, S.K., Kim, W.Y.: A novel approach to the fast computation of Zernike moments. *Pattern Recognition* 39(11), 2065–2076 (2006)
10. Papakostas, G., Boutalis, Y., Karras, D., Mertzios, B.: A new class of Zernike moments for computer vision applications. *Information Sciences* 177(13), 2802–2819 (2007)

11. Wee, C.-Y., Paramesran, R.: On the computational aspects of Zernike moments. *Image and Vision Computing* 25(6), 967–980 (2007)
12. Xin, Y., Pawlak, M., Liao, S.: Accurate computation of zernike moments in polar coordinates. *IEEE Transactions on Image Processing* 16(2), 581–587 (2007)
13. Robles, O.D., Rodríguez, A., Córdoba, M.L.: A study about multiresolution primitives for content-based image retrieval using wavelets. In Hamza, M.H., ed.: *IATED International Conference On Visualization, Imaging, and Image Processing (VIIP 2001)*, Marbella, Spain, IATED, ACTA Press, pp. 506–511 (2001) ISBN 0-88986-309-1
14. Strang, G., Nguyen, T.: *Wavelets and Filter Banks*. Wellesley-Cambridge Press (1997)
15. Starck, J.L., Murtagh, F., Bijaoul, A.: *Image Processing and Data Analysis. The Multiscale Approach*. Cambridge University Press, Cambridge (1998)
16. Rosenfeld, A.: *Multiresolution Image Processing and Analysis*. Springer Series in Information Sciences, vol. 12. Springer, Heidelberg (1984)
17. Marr, D., Hildreth, E.: Theory of edge detection. In: *Proceedings of the Royal Society, London, ser. B, vol. 207*, pp. 187–217 (1980)
18. Daubechies, I.: *Ten Lectures on Wavelets*. vol. 61 of CBMS-NSF Regional Conf. Series in Appl. Math. Society for Industrial and Applied Mathematics, Philadelphia, PA (1992)
19. Mallat, S.G.: A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 11(7), 674–693 (1989)
20. Pastor, L., Rodríguez, A., Ríos, D.: Wavelets for Object Representation and Recognition in Computer Vision. In: Vidaković, B., Müller, P. (eds.) *Bayesian Inference in Wavelet Based Models*. Lectures Notes in Statistics, vol. 141, pp. 267–290. Springer Verlag, New York (1999)
21. Khotanzad, A., Hong, Y.H.: Invariant image recognition by zernike moments. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12(5), 489–497 (1990)
22. Kamila, N.K., Mahapatra, S., Nanda, S.: Invariance image analysis using modified Zernike moments. *Pattern Recognition Letters* 26(6), 747–753 (2005)
23. Zernike, F.: *Beugungstheorie des schneidenverfahrens und seiner verbesserten form, der phasenkontrastmethode* (Diffraction theory of the cut procedure and its improved form, the phase contrast method). *Physica* 1, 689–704 (1934)
24. Rodríguez, A., Robles, O.D., Pastor, L.: New features for Content-Based Image Retrieval using wavelets. In: Muge, F., Pinto, R.C., Piedade, M. (eds.) *V Ibero-american Symposium on Pattern Recognition, SIARP 2000*, Lisbon, Portugal, pp. 517–528 (2000) ISBN 972-97711-1-1
25. MIT Media Lab.: *VisTex*. Web Color image database (1998), <http://vismod.media.mit.edu/vismod/imagery/VisionTexture/vistex.html>
26. Over, P., Ianeva, T., Kraaij, W., Smeaton, A.F.: TRECVID 2006: Search task overview. In: *Proceedings of the TRECVID Workshop, NIST Special Publication* (2006), <http://www-nlpir.nist.gov/projects/tvpubs/tv6.papers/tv6.search.slides-final.pdf>
27. Geusebroek, J.-M., Burghouts, G.J., Smeulders, A.W.M.: The Amsterdam library of object images. *Int. J. Comput. Vision* 61(1), 103–112 (2005)