

Retrieving Similar Segmented Objects Using Motion Descriptors

Konstantinos C. Kartsakalis, Angeliki Skoura, Vasileios Megalooikonomou

Abstract—The fuzzy composition of objects depicted in images acquired through MR imaging or the use of bio-scanners has often been a point of controversy for field experts attempting to effectively delineate between the visualized objects. Modern approaches in medical image segmentation tend to consider fuzziness as a characteristic and inherent feature of the depicted object, instead of an undesirable trait. In this paper, a novel technique for efficient image retrieval in the context of images in which segmented objects are either crisp or fuzzily bounded is presented. Moreover, the proposed method is applied in the case of multiple, even conflicting, segmentations from field experts. Experimental results demonstrate the efficiency of the suggested method in retrieving similar objects from the aforementioned categories while taking into account the fuzzy nature of the depicted data.

Keywords—Fuzzy Object, Fuzzy Image Segmentation, Motion Descriptors, MRI Imaging, Object-Based Image Retrieval.

I. INTRODUCTION

MEDICAL images obtained through Magnetic Resonance (MR) are often difficult to interpret, due to a number of factors like background noise, blurring, poor operator performance or the lack of homogeneity between the materials depicted in the image. Approaches in image segmentation attempt to capture and retain the inherent fuzziness of such images within the context of segmentation [1], utilizing the mathematical framework provided by the theory of fuzzy sets. Objects extracted through the use of the aforementioned methodologies, referred to as "Fuzzy Objects", provide the basis for a new framework of object-based image retrieval that is respectful of their graded composition. Different algorithms implementing such methods tend to either output the object as a graded composition consisting of gray-scale values, or as binary (hard) segmentations [1]-[4], as shown in Fig. 1.

Traditional content-based image retrieval approaches tend to either treat the object depicted as a whole image, operating on features and properties in order to acquire a descriptor, or focus on some of the more or less apparent features of a segmented object or region, like shape, color composition, or texture [5]. More sophisticated approaches include density histograms [6]-[7], color instances [8], correlograms [9], the use of feature points or the construction of specific vectors that contain information from both color and shape of the

object depicted, thus providing similarity measures that are based on both boundary and interior [5]. Further flexibility on measuring boundary similarity is provided by approaches that operate on the contour of the object, like chain code representations [10] or complex zernike moments [11] which also provide scale and translation invariance. In this paper a similar technique that simultaneously draws from both shape and the graded composition of the fuzzy object is employed, based on the traditional Motion Descriptors presented in [12].




Fig. 1 (a) a CT slice of a human knee: each internal structure is considered as an object for segmentation (b) Fuzzy Segmentation depicting pixel degrees of membership within each fuzzy object (c) one binary (hard) segmentation of the same objects (d) multiple binary (hard) manual segmentations of the same object (e) the fuzzy object resulting from combining the manual segmentations shown in (d)

Motion Descriptors presented in [12] are invariant to translations, rotations, scale and reflections of the object under consideration. Furthermore, they have often been used successfully in object color recognition and color image classification. Even more interestingly, H. Fong [13] extended the notion to that of Similarity Descriptors, applying them to grey level images. The author further elaborated on the notion of stability, suggesting that small distortions in the space of objects only result in small distortions in their invariant features, therefore providing a credible measure of similarity between objects.

In this paper, a novel object-based image retrieval framework addressing the issue of efficient retrieval of images

K. C. Kartsakalis is with the Computer Engineering and Informatics Department, University of Patras, Greece (corresponding author to provide phone: 0030-6977822709; fax: 0030-2610-969018; e-mail: kartsaka@ceid.upatras.gr).

A. Skoura and V. Megalooikonomou are with the Computer Engineering and Informatics Department, University of Patras, Greece (e-mail: skoura@ceid.upatras.gr, vasilis@ceid.upatras.gr).

containing fuzzy objects is proposed. Given a query image and a number of automatic or manually created segmentations, the framework scans the pre-existing database of segmentations created offline and retrieves the most relevant to the query, whether binary or fuzzy in nature, as shown in Fig. 2. Based on the distance function calculated using the corresponding Motion Descriptors, similarities in contour, interior or the gradation of intensity values are used as a criterion of matching between images. The most similar matches are then found and reported back to the user.

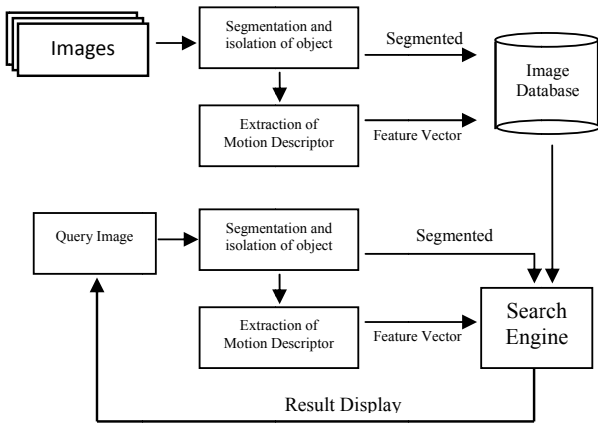


Fig. 2 Framework of proposed image retrieval system for retrieving fuzzy objects according to their segmentation(s)

Motion Descriptors have been used in the past in the context of object recognition and image retrieval of grayscale images, providing a measure of similarity that is both object-oriented and highly robust to small distortions in pixel space and color gradation [12], [13]. The main contribution of this work is therefore the introduction of a framework that operates on images on an object-based level, providing a credible measure of similarity that considers binary and fuzzy segmentations alike.

The rest of this paper is organized as follows: in Section II background knowledge about Motion Descriptors is introduced, along with the pre-processing required for effective application of the various distance measures on the series obtained. Section III describes the retrieval technique in depth. Experimental results validating the technique are then presented in Section IV, and conclusions are given in Section V.

II. BACKGROUND

Motion Descriptors introduced in [12] have been successfully used in the past in order to perform color and invariant object recognition tasks in images. The feature vector for the Motion Descriptor obtained for each object is unique and can be used for shape recognition and similarity measurements. H. Fong extended the notion of Motion Descriptors, defining Similarity Descriptors and applying them to grey level images [13].

The main focus is on a version of Motion Descriptors as

presented in [12], specifically the "Spectral Densities"-type invariants. Let f be a square summable function on the plane, and the Fourier Transform \mathcal{F} being given by:

$$\mathcal{F} = \int_{R^2} f(x) \exp - i(x|\xi) dx$$

By $\langle * | * \rangle$ we denote the scalar product in R^2 . For further application, it is more convenient to explore the nature of the data in terms of polar coordinates: if (r, θ) are the coordinates of point ζ , we shall again denote by $\mathcal{F}(r, \theta)$ the Fourier Transform of f at the point (r, θ) . A mapping from R_+ into R_+ is defined by:

$$D_f(r) = \int_0^{2\pi} |\mathcal{F}(r, \theta)|^2 d\theta$$

D_f is the Motion Descriptor (MD) of f .

Properties: Let f and g be square summable functions. The following properties are valid:

- 1) If M is a motion such that $g(x) = f * M(x)$ for any x in R^2 , then $D_f(r) = D_g(r), r \in R_+$.
- 2) If there exists a reflection R such that $g(x) = f * R(x)$ for any in x in R^2 , then $D_g(r) = D_f(r), r \in R_+$.

The above properties show that Motion Descriptors are invariant under motions and reflections of objects.

The main focus of this work is on a specific family of Motion Descriptors; the family of *Motion Descriptors of Order a* . Let b be an object, \mathcal{B} the Fourier Transform of b , and $a \geq 1$. Let us define the mapping D_b^a from R_+ into R_+ by:

$$D_b(r) = \int_0^{2\pi} |\mathcal{B}(r, \theta)|^a d\theta$$

D_b^a is the Motion Descriptor (MD) of order a of b .

Reference [13] suggests that for any positive number a , the MD of order a of f is *stable*. *Stability* is a property that denotes that similar objects have similar invariant features, without having to be identical or distorted versions of each other.




Fig. 3 Examples of Feature Vectors extracted using Motion Descriptors of 5th order from a mixture of binary, fuzzy on border and textured fuzzy segmentations

Feature Vectors extracted by the use of Motion Descriptors

usually exhibit trends over the radius of integration used in calculating them, as shown in Fig. 3. Whether corresponding to either low or high spatial frequencies used in the calculation of the descriptor over the 2-D Spectrum, the trends can either concern global visual features, such as shape or orientation, or local information such as the fine details, abrupt changes in intensity and the gradation of elements over the object boundaries. Membership grades over the Fuzzy Object are used, in our case, as intensity values.

III. PROPOSED FRAMEWORK

The proposed framework presented in Fig. 2 for content-based image retrieval requires the storage of two kinds of information for each image in the database: segmentation(s) of the image, either manual or automatically generated by a segmentation algorithm, and the corresponding Feature Vector generated by the Motion Descriptor of the segmented object. Therefore, two representations of a segmented object are stored in the database; a binary image depicting the object in black background, and a vector representing the Motion Descriptor. The retrieval of similar segmentations is performed based on similarity measures between feature vectors.

Let N be the total number of images in the database. For each image n of size $w \times h$ pixels, a feature vector D_n of length $L, L \leq \frac{\min(w,h)}{2}$ is stored containing the calculated Motion Descriptor.

Three distinct schemes implementing the retrieval procedure are presented. The first, *Difference Matrix*, calculates the Euclidean Distance between a query vector and all vectors in database, finding the one(s) with the minimum distance. The second, *Low Frequency-Matching*, provides a more cost-effective version of the *Difference Matrix* scheme, only operating on a few elements of each vector instead of the whole length. The last, *Two-Step Matching: Low-Frequency Pre-filtering*, aims at minimizing the number of element-to-element comparisons between vectors while at the same time guaranteeing a small number of false dismissals and a precision comparable to that of the first scheme.

A. Difference Matrix

The retrieval system is either supplied with an unsegmented image, in which case a fuzzy Segmentation Framework such as Fuzzy Connectedness is applied [1], or directly provided a manual or computer generated segmentation. The corresponding Motion Descriptor for the query object q is calculated, which is used as a query feature vector v for retrieval in a database structured as described above. By means of an appropriate distance measure e.g. *Euclidean Distance*, a difference matrix with all distances between query vector v and feature vectors D_n in database is built, and the segmentation images corresponding to the minimum distances calculated are the retrieved results.

Considering the query vector v of size $|v| = L$, and N number of images in the database, the distances between v and feature vector D_n are calculated as:

$$d_n(v, D_n) = \sqrt{\sum_{i=1}^L (v_i - D_{n_i})^2}, \forall n \in N$$

The minimum distance $\min(d_n(v, D_n))$ calculated between query vector v and feature vector D_n corresponds to the most similar image n in database to query object q . Fig. 4 shows the results of a retrieval process.

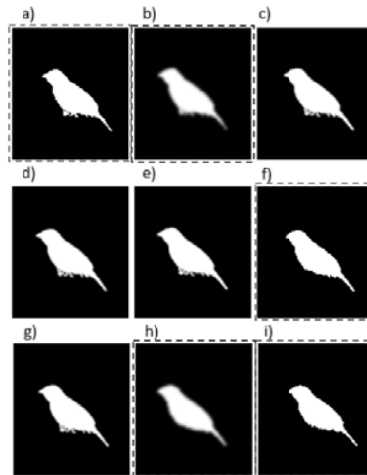


Fig. 4 A retrieval example (a) the query image is a (hard) binary segmentation and is the first retrieved result with $d_a(v, D_a) = 0$ (b), (h) gaussian blurred versions of images a) and f) (f), (i) other (hard) binary segmentations by field specialists (c), (d), (e), (g), (i) fuzzy segmentations of any combination of a), f), i).

Although effective, the above case is considered as naïve, as it calculates all distances between query vector v and all D_n vectors in database, and for each element in the vectors, where $|D_n| = |v| = L$.

B. Low Frequency-Matching

High or even the same precision in retrieval can be achieved by only keeping a small number of the feature vector elements and performing the difference matrix calculations using a truncated version of the original vectors. Utilizing a strategy similar to [14], the first few (4-5) elements of the feature descriptor are kept, which roughly corresponds to only keeping the elements of the Motion Descriptor calculated using *extremely low spatial frequencies* in the *2-D Spectrum*. The above is possible due to the fact that our data have a skewed energy spectrum of the form $O(F^{-b})$. This implies that the first few coefficients contain most of the energy [14], [16], thus providing good estimations of the actual distance between the two sequences.

Due to the fact of the ‘truncation’ causing an underestimation of the distance between the two feature vectors, the presence of false alarms in the retrieved results is expected [14]. Depending on how small the number of elements used is, a few false dismissals are also expected. Experimental results also demonstrate that the order of retrieval is also affected, as compared to the case of using the

whole feature vector for calculations.

The query vector v of size $|v| = L$, and the number of images N in the database are considered again. The distances between v and feature vector D_n are now calculated as:

$$d_n(v, D_n) = \sqrt{\sum_{i=1}^{\Lambda} (v_i - D_{ni})^2}, \forall n \in N, \Lambda < L$$

The “*Low-Frequency Matching*” method significantly improves on the number of comparisons computed, providing robust results that are considerably similar to that of Section III A.

C. Two-Step Matching: Low-Frequency Pre-filtering

The advantages of both matching methods can be combined in a way in which we can enjoy the benefits of a reduced comparison cost, in conjunction with the increased precision and retrieval order of the first method, as shown in Fig. 5.

According to the combined version, a *pre-filtering* retrieval scheme is used which employs the “*Low-Frequency Matching*” method as a first step to quickly filter out non-matching feature vectors. The robustness of the method within a finite number X e.g. 50-60 of the first retrieved results guarantees that an extremely high number of similar segmentations will be supplied to the second step of retrieval, in which all of the elements of the feature vector are used. This results in a time-saving, two-step retrieval scheme which can safely substitute the “*Difference Matrix*” method, at the same time retaining an extremely high or exact order of retrieval and keeping the number of false alarms and dismissals to minimum or even zero, depending on the number of elements used in the first step.




Fig. 5 Two-Step Matching Scheme

The number of N elements used in first-step should be enough to guarantee sufficient successful retrievals within the first X retrieved results. It is thus a point of interest that the statistics of the first step indirectly influence both the precision and the number of false dismissals of the second step.

IV. EXPERIMENTAL RESULTS

Experimental results regarding all three proposed methods are presented here. The “Segmentation Evaluation Database” of Weizmann Institute of Science¹ has been used for

evaluation purposes and in order to quantify the accuracy of the retrieved results. The database was initially compiled by asking human subjects to manually segment gray scale images, with each image being segmented by three different human subjects.

A secondary database extracted from data from MRI depictions of the Corpus Callosum brain fiber [17] has also been used in evaluating the methods.

A. Datasets

The first dataset of the “Segmentation Evaluation Database” contains a large number of binary segmentations, which, for our experimental purposes, had to be extended to also include fuzzy segmentations. On top of the already existing 3 binary segmentations per class, 4 more were added by all possible combinations of the existing ones. Furthermore, we obtain another 6 blurry segmentations by applying intense or light Gaussian blurring on the first 3 images, and 3 additional ones by applying texturing on their light Gaussian blurred version. In total, 95 classes of 16 segmentations per class were employed to apply the three suggested methods.

The second dataset, labeled “Corpus Callosum” [17], contains 19 segmentations depicting the colossal commissure, a wide flatbundle of neural fibers beneath the cortex in the eutherian brain at the longitudinal fissure, in patients suffering from “*Chromosome 22 Deletion Syndrome*” [15]. An extra 11 control (healthy) segmentations are provided. The dataset is therefore grouped into two uneven classes, which are then expanded in a similar manner as with the first dataset, to include blurry versions of the manual segmentations. In total, this dataset contains 120 segmentations, out of which 76 belong to the first class and 44 to the second one.

Prior to the construction of the Motion Descriptor(s), an extra pre-processing step may be required: images in which the black background greatly outcales the depicted object may need to be resized to match the rest of the database. Calculations are best favored by using images with same size and a similar background-to-object ratio, in which the object is large enough to guarantee sufficient high-frequency content. However, since the first few coefficients largely outweigh the significance of the last, we expect the method to be robust in case of a loss of quality in the high-frequency component of an image.

B. Results

Experimental results on both the original and extended forms of the dataset demonstrated that Motion Descriptors can be used as an effective similarity measure.

To evaluate the proposed methods, *Recall* and *Precision* were used. Precision is the fraction of retrieved instances that are relevant to the query, whereas recall is the fraction of relevant instances that are retrieved. A retrieval is considered a “hit” if the retrieved result belongs to the same class as the query object selected. Extra-class similarity, although evident in some cases, is not taken into consideration when calculating the above metrics.

¹ Weizmann Institute of Science – Segmentation evaluation database: http://www.wisdom.weizmann.ac.il/~vision/Seg_Evaluation_DB/

TABLE I (A)
RECALL RESULTS

Order	Measure	Number of Retrievals						
		16	20	25	30	40	50	60
$\alpha = 1$	Recall	98.82%	98.88%	98.88%	99.01%	99.47%	99.80%	100%
	Av. Time	0.0141 s	0.0168 s	0.0146 s	0.0160 s	0.0168 s	0.0164 s	0.0159 s
$\alpha = 2$	Recall	96.91%	97.43%	98.55%	99.14%	99.47%	99.67%	99.74%
	Av. Time	0.0152 s	0.0146 s	0.0155 s	0.0155 s	0.0145 s	0.0145 s	0.0153 s
$\alpha = 3$	Recall	97.17%	98.03%	98.75%	99.08%	99.28%	99.54%	99.61%
	Av. Time	0.0167 s	0.0144 s	0.0143 s	0.0157 s	0.0149 s	0.0144 s	0.0144 s
$\alpha = 4$	Recall	90.86%	92.83%	94.47%	96.32%	98.03%	98.75%	99.21%
	Av. Time	0.0197 s	0.0142 s	0.0143 s	0.0143 s	0.0148 s	0.0161 s	0.0152 s

Recall results for Motion Descriptors of an increasing Order from Weizmann Institute "Segmentation Evaluation" Database

TABLE I (B)
RECALL RESULTS

Order	Measure	Number of Retrievals				
		80	90	100	110	120
$\alpha = 1$	Recall	100%	100%	100%	100%	100%
	Av. Time	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s
$\alpha = 2$	Recall	100%	100%	100%	100%	100%
	Av. Time	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s
$\alpha = 3$	Recall	89,47%	89,47%	89,47%	93,42%	100%
	Av. Time	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s	$2,859 * 10^{-11}$ s
$\alpha = 4$	Recall	94,73%	94,73%	94,73%	94,73%	100%
	Av. Time	$2,852 * 10^{-11}$ s	$2,851 * 10^{-11}$ s	$2,853 * 10^{-11}$ s	$2,854 * 10^{-11}$ s	$2,853 * 10^{-11}$ s

Recall results for Motion Descriptors of an increasing Order from "Corpus Callosum" Database

In Table I, we first present the *Recall* statistics of the first method, that is, the number of "hits" – objects belonging to the same class – within the first 16, 20, 25, 30, 40, 50 or 60 results for the first database, or within the first 80, 90, 100, 110 or 120 (all) results for the second database. The results presented were calculated for Motion Descriptors altering the values of parameter α .

The results of Table I provide solid ground in further

improving on the results of the first method. A large number of experiments has been conducted using Descriptors of different orders, in order to validate the usefulness of the conclusions drawn on the second method of *Low-Frequency Matching*. In Tables II and III we provide results for the second method, when using only a finite number of the very first elements for the calculation of distances between feature vectors.

TABLE II (A)
RECALL RESULTS

#Elements used	Number of Retrieved Results						
	16	20	25	30	40	50	60
1	81.38%	85.20%	88.36%	90.99%	93.68%	94.80%	95.59%
2	93.55%	94.61%	95.92%	97.04%	97.96%	98.42%	98.42%
6	98.62%	98.68%	98.68%	99.01%	99.34%	99.93%	100%
7	98.55%	98.68%	98.75%	99.01%	99.28%	99.87%	100%
10	98.62%	98.88%	98.88%	99.01%	99.47%	99.80%	100%
15	98.82%	98.88%	98.88%	99.01%	99.47%	99.80%	100%

Recall results for Motion Descriptors of 1st Order from Weizmann Institute "Segmentation Evaluation" Database, using only a limited number of elements as in *Low Frequency-Matching* method.

TABLE II (B)
RECALL RESULTS

#Elements used	Number of Retrieved Results					
	76	80	90	100	110	120
1	96.05%	96.05%	96.05%	96.05%	96.05%	100.00%
7	98.68%	98.68%	98.68%	98.68%	98.68%	100.00%
12	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%

Recall results for Motion Descriptors of 1st Order from "Corpus Callosum" Database, using only a limited number of elements as in *Low Frequency-Matching* method.

The above results of Table II concern the use of Motion Descriptors of 1st Order. Motion Descriptors of higher order, although presenting statistics that are less precise, demonstrate similar results in terms of 50-60 retrieved objects for the first database. Further experiments using a higher number of elements in the first step were also conducted, without however providing any improvement in the recall or precision metrics achieved by using only 10-15 elements (as shown in Table II). The above indicates that by keeping only the first X retrievals using only the first few elements we can safely proceed towards a second step that uses the full length of the descriptor and is specifically geared towards eliminating false

dismissals.

Due to the uneven sizes of elements per class, the dissimilarity between the two classes and the high level of efficiency of the *Low-Frequency Matching* method as a first step, results regarding the “Corpus Callosum” database for the case of the *Low-Frequency Prefiltering Two-Step Scheme* do not provide an example of the improvement offered by the second step and are therefore not presented in this paper. For the “Segmentation Evaluation” Database, the above scheme provides results that are either extremely similar or exactly the same with the first method, while at the same time reducing the number of element-to-element comparisons drastically.

TABLE III
PRECISION RESULTS

Order		Number of elements Λ , $\Lambda < L$, used in 1 st step					
		1	3	4	7	15	20
$\alpha=1$	Recall	94.80%	98.28%	98.81%	98.81%	98.81%	98.81%
	Av. Time	$0.1139 * 10^{-3}$ s	$0.1140 * 10^{-3}$ s	$0.1155 * 10^{-3}$ s	$0.1146 * 10^{-3}$ s	$0.1145 * 10^{-3}$ s	$0.1183 * 10^{-3}$ s
$\alpha=2$	Recall	94.73%	96.90%	96.90%	96.90%	96.90%	96.90%
	Av. Time	$0.1189 * 10^{-3}$ s	$0.1236 * 10^{-3}$ s	$0.1233 * 10^{-3}$ s	$0.1213 * 10^{-3}$ s	$0.1385 * 10^{-3}$ s	$0.1304 * 10^{-3}$ s
$\alpha=3$	Recall	94.27%	97.17%	97.17%	97.17%	97.17%	97.17%
	Av. Time	$0.1172 * 10^{-3}$ s	$0.1223 * 10^{-3}$ s	$0.1224 * 10^{-3}$ s	$0.1199 * 10^{-3}$ s	$0.1207 * 10^{-3}$ s	$0.1217 * 10^{-3}$ s
$\alpha=4$	Recall	93.68%	90.85%	90.85%	90.85%	90.85%	90.85%
	Av. Time	$0.1212 * 10^{-3}$ s	$0.1208 * 10^{-3}$ s	$0.1199 * 10^{-3}$ s	$0.1214 * 10^{-3}$ s	$0.1199 * 10^{-3}$ s	$0.1261 * 10^{-3}$ s

Precision results for Motion Descriptors of different Orders using *Low-Frequency Prefiltering Two-Step Scheme* method for the Weizmann Institute “Segmentation Evaluation” Database.

V. CONCLUSION

In this paper, a new framework for retrieving segmentations, either binary or fuzzy, is proposed. The data provided as input or existing in our database can either come from manual segmentations from field experts or through the use of some well-adjusted automatic segmentation framework, such as the Fuzzy Connectedness framework [1]-[4]. Three different alternatives are provided: the first demonstrates the effectiveness of Motion Descriptors when used to provide similarity measurements between different segmentations, performing $N \times \binom{M}{2} - 1$ comparisons where N is the number of images of $M \times M$ size in database and $\binom{M}{2} - 1$ is the size of each descriptor. The second elaborates on properties inherent to the 2-D spectrum, providing a similarity measure that is fast and effective while at the same time keeping the introduction of false dismissals or inaccuracies to a minimum. The method performs $N \times \Lambda$ comparisons where $\Lambda < \frac{M}{2} - 1$ is the number of low-frequency elements kept. The third alternative makes use of a two-step pre-filtering scheme, which quickly filters out irrelevant segmentations using minimum calculations, leaving the second step to operate on a significantly smaller number of input vectors. The third method performs: $N \times \Lambda$ in the first step and another $X \times \binom{M}{2} - 1$ in the second step, where $X < N$, therefore providing a significant improvement in calculation speed. Preliminary experimental results demonstrate the improved speed and efficiency of a method that operates on both boundary and interior of the segmented object.

ACKNOWLEDGMENT

The authors would like to thank Professor K. Berberidis of the Computer Engineering and Informatics Department of University of Patras, for providing constructive comments on previous versions of the paper. This research has been co-financed by the European Union (European Social Fund – ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Frame-work (NSRF) - Research Funding Programs: Thales, Investing in knowledge society through the European Social Fund.

REFERENCES

- [1] J. K. Udupa, S. Samarasekera, “Fuzzy Connectedness and Object Definition: Theory, Algorithms, and Applications in Image Segmentation.” *Graphical Model and Image Processing*, vol. 58, no. 3, pp. 246-261, May 1996.
- [2] K. Ch. Ciesielski, J. K. Udupa, P. K. Saha, and Y. Zhuge, “Iterative Relative Fuzzy Connectedness for Multiple Objects with Multiple Seeds,” *Computer Vision and Image Understanding*, vol. 107, no. 3, pp. 160-182, Sep. 2007.
- [3] J. K. Udupa, P. K. Saha, “Fuzzy Connectedness and Image Segmentation,” *Proceedings of the IEEE*, vol. 91, no. 10, pp. 1649-1669, Oct. 2003.
- [4] J. K. Udupa, P. K. Saha, and R. A. Lotufo, “Relative Fuzzy Connectedness and Object Definition: Theory, Algorithms, and Applications in Image Segmentation”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 11, pp. 1485-1500, Nov. 2002.
- [5] R. Datta, D. Joshi, J. Li and J. Z. Wang, “Image Retrieval: Ideas, Influences, and Trends of the New Age”, *ACM Computing Surveys*, vol. 40, no. 2, pp. 1-60, Sep. 2008.
- [6] M. J. Swain, D. H. Ballard “Color indexing”, *International Journal of Computer Vision*, vol. 7, no. 1, pp. 11-13, Nov. 1991.

- [7] J. L. Hafner, H. S. Sawhney, W. Equitz, M. Flickner, and W. Niblack, "Efficient color histogram indexing for quadratic form distance functions", *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 17, no. 7, pp. 729-736, July 1995.
- [8] M. Stricker, and A. Dimai, "Color indexing with weak spatial constraints", in *Proc. SPIE Storage and Retrieval of Still Image and Video Databases IV*, vol. 2670, 1996, pp. 29-40.
- [9] J. Huang, S. R. Kumar, M. Mitra, W. Zhu, and R. Zabih, "Image indexing using color correlograms", in *Proc. Computer Vision and Pattern Recognition, IEEE Computer Society*, Puerto Rico, 1997, pp. 762-768.
- [10] R. C. Gonzalez and P. Wintz, *Digital Image Processing*, 2nd ed. Boston, MA: Addison-Welsey Longman Publishing Co., 1987.
- [11] S. Li, M. Lee, and C. Pun, "Complex Zernike Moments Features for Shape-Based Image Retrieval," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 39, no. 1, pp. 227-237, Jan. 2009.
- [12] F. Smach, C. Lemaître, J. Gauthier, J. Miteran, and M. Atri, "Generalized Fourier Descriptors with Applications to Object Recognition in SVM Context," *Journal of Mathematical Imaging and Vision*, vol. 30, no. 1, pp. 43-71, Jan. 2008.
- [13] H. Fong, "Pattern Recognition in gray-level images by Fourier analysis", *Pattern Recognition Letters*, vol. 17, no. 14, pp. 1477-1489, Dec. 1996.
- [14] C. Faloutsos, M. Ranganathan, and Y. Manolopoulos, "Fast Subsequence Matching n Time-Series Databases", in *Proc. of ACM SIGMOD international conference on Management of data*, Minneapolis, 1994, pp. 419-429.
- [15] K. M. Antshela, W. R. Katesa, N. Roizena, W. Fremonta and R. J. Shprintzena, "22q11.2 Deletion Syndrome: Genetics, Neuroanatomy and Cognitive / Behavioral Features Keywords", *Child Neuropsychology: A Journal on Normal and Abnormal Development in Childhood and Adolescence*, vol. 11, issue 1, pp. 5-19, Feb. 2005.
- [16] R. Agrawal, C. Faloutsos, and A. Swami, "Efficient similarity search in sequence databases," in *Proc. of 4th International Conference of Foundations of Data Organization and Algorithms (FODO)*, Chicago, 1993, pp. 69-84.
- [17] A Machado, T. Simon, V. Nguyen, D. McDonald-McGinn, E. Zackai, and J. Gee, "Corpus callosum morphology and ventricular size in chromosome 22q11.2 deletion syndrome", *Brain Research*, vol. 1131, pp. 197-210, Feb. 2007.

and three book chapters. His main research interests include medical image analysis, pattern recognition, data mining, data compression, biomedical informatics, and multimedia database systems.

Prof. Megalooikonomou is a member of the ACM, IEEE, SIAM, and SPIE. In 2003 he received a CAREER award from the US National Science Foundation for developing data mining methods for extracting patterns from medical image databases. He regularly serves as a program committee member or referee on several premier conferences and journals in his areas of research.



Kartsakalis Ch. Konstantinos was born in Agrinio, Greece, in 1987. He received his B.E. degree in computer engineering & informatics from the University of Patras, Patras, Greece, in 2010 and an M.Sc degree in computer science & technology from the same university in 2014 respectively. Since 2012 he has been with Intracom Telecom AE, Patras, Greece where he is currently

employed as a Verification Engineer.



Angeliki Skoura received a Diploma (2007), a M.Sc. degree (2009) and a Ph.D. in Computer Engineering and Informatics Department of University of Patras.

Currently, she is a Post-Doctoral Researcher at the University of Patras and also works as a Data Analyst in the private sector. Her research interests lie in the areas of image analysis (image segmentation, feature extraction), data mining (mining of spatiotemporal data, ensemble learning) and decision support systems with applications to, among other, medical informatics and bioinformatics. She holds work experience in designing and developing information systems and she has participated in several research projects.



Vasilis Megalooikonomou (M'95) received a BSc in computer engineering and informatics from the University of Patras, Greece in 1991, and a M.S. and Ph.D. in computer science from the University of Maryland, Baltimore County, USA, in 1995 and 1997, respectively.

He is currently a professor at the Department of Computer Engineering and Informatics at the University of Patras, Greece. Prior to this appointment he has been on the faculty of Johns Hopkins University, Dartmouth College and Temple University. He has co-authored over 150 refereed articles in journals and conference proceedings