

Texture Based Medical Image Indexing and Retrieval: Application to Cardiac Imaging*

Tristan Glatard, Johan Montagnat, Isabelle E. Magnin
CREATIS (CNRS-Inserm), INSA
69621 Villeurbanne Cedex, France
{glatard,johan,magnin}@creatis.insa-lyon.fr

ABSTRACT

Although digital images indexing and querying techniques have extensively been studied for the last years, few systems are dedicated to medical images today while the need for content-based analysis and retrieval tools increases with the growth of digital medical image databases. We analyze medical image properties and we evaluate Gabor-filter based features extraction for medical images indexing and classification. The goal is to perform clinically relevant queries on large image databases that do not require user supervision. We demonstrate on the concrete case of cardiac imaging that these techniques can be used for indexing, retrieval by similarity queries, and to some extent, extracting clinically relevant information out of the images.

Categories and Subject Descriptors

H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing; H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval

General Terms

Management, Experimentation

Keywords

Medical imaging, texture, indexing, retrieval, segmentation

1. INTRODUCTION

1.1 Motivations

Medical images have become a key investigation tool for medical diagnosis and pathology follow-ups. Digital imaging is becoming the standard for all image acquisition devices and, with the generalization of digital acquisition, there is

*This work is supported by the French Ministry for Research ACI-GRID program:
<http://www.creatis.insa-lyon.fr/MEDIGRID/>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MIR'04, October 15–16, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-940-3/04/0010 ...\$5.00.

an increasing need for data storage and retrieval. Medical images represent an enormous amount of data: the annual production of a single average size radiology department represents tens of terabytes of data. Therefore, petabytes of medical images are produced in industrialized countries each year. For these data, there is a need for very long term archiving (the European Union recommends the archiving of all medical data for 20 years and up to 70 years in specific cases).

With the growth of medical databases, new applications devoted to statistical analysis of medical data have emerged such as breast cancer screening, lungs images analysis, or oncology in general. The datasets manipulated for these applications are not images of one patient or coming from one radiology department but rather images showing a particular pathology or specific features. This data set has to be dynamically assembled by automatically selecting relevant images among available databases. Similarly, a physician is often interested by *similar cases* to the one he is studying and the similarity measurement usually involves a medical background. Archiving of tons of medical images is only relevant if adapted query tools exist that allow medical users to browse the data sources. Driven by applications, photographs image indexing and content based retrieval systems have been thoroughly studied these last years. However, there has been far less effort to deploy medical image dedicated systems so far [11].

In the early years of image database assembling, image indexing used to be mostly text-based and manual: images were annotated with keywords and retrieved by using a text-based database management system. In the medical world, the images acquired are usually accompanied by metadata related to the patient, the image acquisition and the radiology department responsible for the acquisition. Indeed, the DICOM (Digital Image and COmmunication in Medicine) format that recently emerged as the standard for medical image storage and communication encompasses textual information fields. Nevertheless, textual information is limited for two main reasons: the large increase of the data volume, which makes manual annotation tiresome and the difficulty to express the image content with keywords which are often inconsistently assigned among different indexers: medical records are complex, difficult to analyze, and rarely available on-line. Except in some specific cases, only the image content itself carries the necessary information for indexing and retrieval.

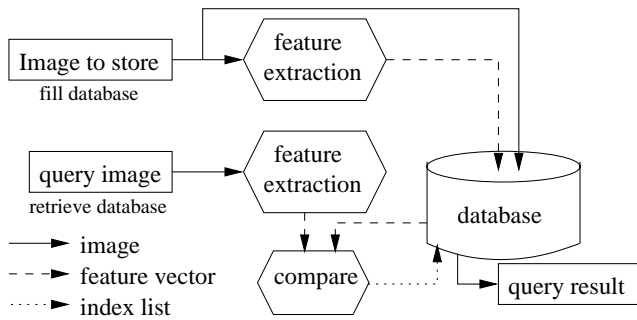


Figure 1: General architecture of CBIR-systems

1.2 Content based image retrieval systems

Content-Based Image Retrieval (CBIR) emerged in the early 1990s mainly to index color photographs. In this approach, images are represented by a vector in a feature space and a similarity measure between images is defined from a distance in the feature space. Figure 1 presents the general architecture of CBIR systems proposed in [14]. Given a query image, such a system first extracts its feature vector and then compares it to those of the images stored in the database. This way, it can rank images of the database according to the distance of their feature vector to the query one. This ranking is the query result.

Many CBIR were developed during the last years, both by commercial firms and academia. The earliest and most famous one is QBIC (Query By Image Content) [7] which was proposed by IBM. Virage is another commercial CBIR which is for example used by CNN [1]. Many systems have also been proposed from the academia like Photobook [19], VisualSEEK [24], Mars (Multimedia Analysis and Retrieval System) [9] - which proposes a solution of relevance feedback -, Candid [12] or Netra [15]. All of those systems only use low level features - mainly color, texture and shape - and do not include any semantic level. Later systems like Blobworld [3] include a segmentation step in the query process to integrate higher level information. In Blobworld, the segmentation step is based on color, position and texture features and leads to a small number of homogeneous regions called blobs. Submitting a request, the user can select the relevant blobs he wants the query to be applied on. Nevertheless, no system offers any interpretation of images yet, which would require a dedicated layer in the system. This loss of information from an image to its representation in a feature space is called the semantic gap [18].

1.3 Medical CBIR

Despite CBIR developments, existing systems hardly address all medical image properties. There exists a large number of medical image acquisition devices among which *computed tomography* scanners (CT), *magnetic resonance imagers* (MRI), *ultrasound* probes (US) and nuclear imagers are the most widely used. They provide images with very different properties in terms of resolution, contrast, and signal to noise ratio. They are highly specialized and they produce images giving very different information on the human body anatomy and physiology. Inside one modality, the tuning of an imager may lead to significantly varying images (*e.g.* an MRI may be used for acquisition of completely different anatomical and functional information).

Although the need is clearly identified, the difficulty to develop medical image indexing and retrieval tools is due to several factors:

- In most cases, medical images are intensity only images carrying less information than color images. In some rare cases however, vectorial images may be produced (*e.g.* tensor MRI). More frequently, multimodality images of a same body may be acquired (*e.g.* MRI and ultrasound images of the same area). However, multiple images are usually not aligned in space and require an additional registration procedure. In the case of soft and deformable structures, complex non-rigid registration procedures are required.
- Medical images are usually low resolution and high noise images. They are difficult to automatically analyze for extracting features. Medical images acquired with different devices, even using the same modality, may have significantly varying properties. Some authors proposed image correction and normalization algorithms to improve image comparison.
- Ideally, medical images should be indexed on medical criteria that are extremely variable depending on the kind of image acquisition considered (imaged body area, clinical context, etc). Moreover, automatic diagnosis in medical images is mostly impossible today except in rare specific cases. Medical images interpretation is often difficult even for trained radiologists.

On the other hand, a large fraction of medical acquisition devices produce 3D data today that provide an additional information not available in 2D images. The use of 3D, and sometimes multisequence, data may enable powerful feature detection in some cases.

As [11] notice, the medical image retrieval must often be processed according to pathology bearing regions which are precisely delimited on the images and could not be automatically detected in the general case. Moreover, low level features like color, texture or shape are not sufficient to describe medical images [14]. As a consequence, medical CBIRs require a high level of content understanding and interpretation of images, which implies their automatic segmentation [18]. Finally, a high level of query completion and accuracy is required by such systems to make them reliable from a clinical point of view [25].

Some specialized CBIR have thus been proposed in medical applications. Chu *et al.* [4] present an image retrieval system dedicated to brain MRI which indexes images mainly on the shape of the ventricular region. Korn *et al* [13] propose a system for fast and effective retrieval of tumor shapes in mammogram X-rays. Comaniciu *et al* [6] describe a system which aims at helping physicians in the diagnosis of lymphoproliferative disorders of blood. Nevertheless, a description of the clinical use of such systems is very rare [18], except for the systems ASSERT [22], which is dedicated to HRCT images of the lung and includes information from physicians - such as anatomical landmarks and pathology bearing regions - and IRMA [14] which proposes a multi-step approach for the classification of images into anatomical areas, modalities and view points.

1.4 Objectives

Digital image processing algorithms have been developed for decades for extracting features. They yield to high level interpretation by providing quantitative informations on medical image contents. They are often highly specialized and not well adapted to general image indexing and retrieval.

In this paper we avoid to address the medical image interpretation problem that is very specific to some particular applications. Instead, we focus on general purpose feature extraction techniques for image indexing. We present local feature extraction measures and we show how classification and un-supervised pre-segmentation can help in indexing medical image databases.

The intend is the automatic indexing of acquired medical data without user supervision for responding to medical users queries. In this context, it is only relevant to compare images of a same modality, and sometimes with the same acquisition settings even inside a same modality: image variability across modalities is far too high and only high level image analysis and interpretation tools could provide information to correlate such different data. Results are shown on magnetic resonance cardiac sequences.

2. MEDICAL IMAGES INDEXING AND RETRIEVAL

CBIR systems have often indexed images on global features extracted from the whole image. However, the need for features related to subcomponents of the image has emerged and bloc-based algorithms have been designed with different granularity levels. Bloc size can range from full image chunks to local areas made of a limited number of pixels.

2.1 Global features

Global features such as image histograms have often been used for color images indexing. They provide an “idea at a glance” of the image content. In medical images, the full image histograms are not very informative as the image are usually intensity images only and the histogram may be very variable for different image resolutions.

Similarity measures have often been used in the medical image domain for comparing images *to a reference* [20, 17]. Similarity measures use a voxel-to-voxel measurement to return a single similarity coefficient to the user. However, they have two drawbacks in the context of medical indexing. First, images require proper registration before a similarity measure can be computed and general purpose fully automatic registration algorithms are not available yet. Second, these measures depend on a reference image and will vary with a reference change.

Therefore, local descriptors are more appropriate for generic medical image indexing. Indeed, bloc-based algorithms have been recently proposed for similarity measurements, trying to make this kind of measure more specific [21].

2.2 Local features

Medical images are often highly textured and voxel based analysis is an interesting path to explore. Moreover, as color and intensity are not as important in medical images as in photographs, texture analysis becomes crucial in medical image retrieval. Defining texture is something quite ambiguous because of the imprecise nature of this concept. According to Smith and Chang [23], texture refers to visual

patterns which have properties of homogeneity and cannot result from the presence of only a single color or intensity. Texture perception plays an important role in the human visual system of recognition and interpretation. Two main approaches can be identified concerning texture analysis: statistical methods and filtering methods.

Among statistical methods, co-occurrence matrices, which were proposed by Haralick in 1973 [8], are still widely used. Those matrices model spatial dependencies between grey levels of an image. Given a distance d and an orientation θ , the (i, j) coefficient of the corresponding matrix is the probability of going from a grey level i to a grey level j with an intersample spacing of d along the axis making an angle θ with the x axis. For instance, with distance $d = 1$ and orientation $\theta = 0$, the unnormalized coefficients are defined in an image I by:

$$P(i, j, 1, 0) = \text{card}\{ (k, l), (m, n) \in (I \times I) / |k - m| = 1, \\ l - n = 0, I(k, l) = i, I(m, n) = j \}$$

where $\text{card}\{\}$ denotes the number of elements in the set and $I(m, n)$ gives the grey level of the pixel (m, n) .

After the computation and normalization of these matrices, a set of 14 second-order statistics that are the textural features of the image is computed out of the coefficients .

Tamura *et al* [26] also proposed a statistical approach to texture extraction based on psychological studies on human perception. They identified 6 meaningful texture features, namely, coarseness, contrast, directionality, line-likeness, regularity and roughness. The autocorrelation function of an image can also be used to quantify the regularity and the coarseness of a texture. This function is defined with the following formula for an image I :

$$\rho(x, y) = \frac{\sum_{u=0}^N \sum_{v=0}^N I(u, v) I(u+x, v+y)}{\sum_{u=0}^N \sum_{v=0}^N I^2(u, v)}$$

If the texture of an image is coarse, its autocorrelation function smoothly decreases and if the texture is fine, the function rapidly decreases. If the texture is regular, the autocorrelation function shows local extrema [27].

Filtering based texture extraction consists in using sets of filters to analyze the image spectrum and characterize it. Among filters used in texture extraction, Gabor ones are widely mentioned in the literature [2, 16, 5]. Gabor filters are used in banks, in which each filter is tuned to a specific orientation and spatial frequency. A two dimensional Gabor filter is a Gaussian-modulated sinusoid. The impulse response of its real (even) version is given by:

$$h(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp \left[-\frac{1}{2} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) \right] \cos(2\pi Fx)$$

where F is the frequency of the filter

In the spatial-frequency domain, this filter is represented by two symmetrical Gaussians:

$$H(u, v) = \exp \{ -2\pi^2 [(u - F)^2 \sigma_x^2 + v^2 \sigma_y^2] \} + \\ \exp \{ -2\pi^2 [(u + F)^2 \sigma_x^2 + v^2 \sigma_y^2] \}$$

To implement a Gabor bank of filters, the parameters F , θ , σ_x , σ_y and the frequency and angular bandwidths B_F and B_θ must be properly set in order to capture textural information well. Most authors recommend to set B_F and B_θ in order to obtain non overlapping filters and to determine σ_x

and σ_y according to the following formulas:

$$\sigma_x = \frac{\sqrt{\ln 2}(2^{B_F} + 1)}{\sqrt{2\pi}F(2^{B_F} - 1)} \quad \text{and} \quad \sigma_y = \frac{\sqrt{\ln 2}}{\sqrt{2\pi}F \tan(\frac{B_\theta}{2})}$$

As [5] mentioned, using center frequency spacing of one octave with frequency bandwidth B_F of one octave and orientation spacing of 30° with angular bandwidth B_θ of 30° is characteristic of the human visual system.

Even if no single best texture model has been identified, Gabor filters are strongly correlated with the human visual system which is *made up of a number of independent detector mechanisms each preceded by a relatively narrow band filter tuned to a different frequency* [5]. Thus, their use for texture features extraction is particularly relevant in image retrieval applications [16] in particular in the medical domain where texture plays a key role in image interpretation. As a consequence, we subsequently use them for texture description.

2.3 Features localization: segmentation

Local image features extraction in medical images will give different image information in areas covered by different tissues. In some medical applications where a tissue of interest covers a large fraction of the image or a prior knowledge on the region of interest is available, extracting features by fixed blocs in the image is sufficient. However in the general case, one would like to identify features for each tissue in the image. This would require prior image segmentation. Medical image segmentation is one of the most challenging problem in medical image analysis and a very active research topic. Therefore, there is no algorithm available in the general case for isolating medical image regions.

Alternatively, texture analysis by blocs gives information on the composition of various areas in an image and therefore provides clue for a rough segmentation of the image. Classification of image voxels by their texture features is a possible pre-indexation stage of the image. Indeed, image retrieval and segmentation are strongly correlated. On one hand, segmentation is the first step to introduce semantic features in a CBIR system because it is required to identify objects or relevant regions in an image. And on the other hand, segmenting an image is equivalent to the classification of its pixels and image retrieval also corresponds to a classification task. That is why we considered the segmentation problem in this work.

2.4 Algorithms

Image indexing. In order to index the images, we used a bank of 42 Gabor filters with an angular spacing of 30° (corresponding to the orientations $0^\circ, 30^\circ, 60^\circ, 90^\circ, 120^\circ$ and 150°) and a frequency spacing of one octave (corresponding to the frequencies $\sqrt{2}, 2\sqrt{2}, \dots, \frac{N}{4}\sqrt{2}$ cycles per image, N being the size of the image). We used angular bandwidths of 30° , frequency bandwidths of 1 octave and determined σ_x and σ_y as presented in 2.2 to obtain non overlapping filters. The spectrum of this bank of Gabor filters is shown in figure 2. We then computed the mean and the standard deviation of the magnitude response of the images, in order to obtain feature vectors as proposed by Manjunath *et al* [16].

Segmentation. Similarly, we propose a texture-based segmentation algorithm which will be applied to an indexation task in the next section. For each pixel of our images,




Figure 2: Spectrum of the used bank of Gabor filters (half-peak supports of the filters)

we considered a 8×8 neighborhood, on which we applied a bank of 16 Gabor filters with an angular spacing of 30° (corresponding to the orientations $0^\circ, 30^\circ, 60^\circ, 90^\circ, 120^\circ$ and 150°) and a frequency spacing of one octave (corresponding to the frequencies $\sqrt{2}$ and $2\sqrt{2}$ cycles per image). Standard deviations and bandwidths were tuned as presented in the preceding section.

The feature vectors obtained were classified using the k-nearest neighbors algorithm [10]. An advantage of this method is that the only input parameter of the classifier is the number of classes. Each class is represented in the feature space by its barycenter that is initially randomly set. Then, each element is assigned to its nearest class (according to the Euclidean distance) and barycenters are iteratively updated. The algorithm stops when no more element moves from one class to another during an iteration.

Given that no assumption is done concerning the nature of the input data of the algorithm, it can be used for general purpose. Consequently, we tested this method on various images. Figure 3 shows the segmentation result on cardiac MRI segmented into three classes corresponding to the blood, the myocardium and the background (top row) and on brain MRI segmented into white matter, grey matter, cerebro-spinal fluid and skull (bottom row).

Image retrieval. The principle of the queries is the computation of a distance between the features of the query image and those of the images in the database. Once all the distances are computed, the algorithm ranks the images of the database from the nearest to the furthest to the query image. We used Euclidean distance to process our queries. Thus, the distance between two feature vectors f_0 and f_1 is given by the following formula:

$$d = \sqrt{\sum_{i=1}^N (f_0(i) - f_1(i))^2}$$

where $f_0(i)$ denotes the i^{th} coordinate of the vector f_0 and N is the dimension of the feature space.

3. APPLICATION TO CARDIAC IMAGING

3.1 Magnetic Resonance Images of the heart

In this section we demonstrate the relevance of a texture based indexing and retrieval system for medical images with




Figure 3: A cardiac (top) and a brain (bottom) MRI slice and the corresponding classified images

a database of cardiac magnetic resonance images. MRI is a modality widely used for cardiac analysis. It can provide 3D images and temporal acquisitions along the cardiac cycle. Figure 4 shows an example of the images used in this experiment. Each 2D image is one slice of the thorax of a patient. The heart appears on the top left area of the center (a zoom on the heart area can be seen on left of figure 3). The darker area next to the heart is a lung. Inside the heart myocardium, the blood filled cavities of the right and left ventricles appear as bright structures.

The MRI device acquiring the images used in this section produced temporal sequences of 3D images: 8 to 10 volumes are acquired along one cardiac cycle. The volumes are stored as sets of cutting slices in the image database. Each volume is made of 10 slices. Each slice is similar to the ones shown in figure 4. The total amount of slices processed in this experiment is 170.

One sequence shows a complete heart beat: contraction phase (systole) followed by a dilation phase (diastole). There are two kind of queries that could be made considering that kind of images: (i) given one slice, find all similar slices (slices corresponding to the same body area *i.e.* with the same vertical position) or (ii) given one slice, find slices corresponding to a given instant in time (in particular the end of systole or the end of diastole).

3.2 Cardiac images indexing and database clustering

The indexes proposed in section 2.4 were computed for each slice and the feature vectors were stored in the image database. We then clustered the database according to those indexes using the k-nearest neighbors algorithm. We split the database into 10 clusters, each of them being assumed to




Figure 4: Feature-based image retrieval : example of query.

gather images with the same value of the vertical coordinate in the initial sequences.

3.3 Cardiac images retrieval

Figure 4 shows the result of a feature-based query performed on the database given a sample image. The sample image is shown on the top left and the “nearest” database images are shown by decreasing order from left to right and top to bottom. t represents the image instant and h the slice vertical position. Notice that the system best retrieves images that belong to the same vertical position in the query image. This result is consistent because the features were computed on the whole image whereas the limited cardiac region is almost the only one which varies with time. Therefore, the overall image aspect has more influence than the temporal variations of the heart.

3.4 Segmentation assisted retrieval

Nevertheless, taking into account medical considerations, the retrieval of images that belong to the same instant that the query image is an important problem to cope with. Therefore, we focused on the retrieval of the end of systole into a 3D+t cardiac MRI sequence. To face this problem,

	μ	σ
Cluster 0	0.74	0.74
Cluster 1	0.97	0.81
Cluster 2	2.96	0.84
Cluster 3	2.91	0.29
Cluster 4	6	0
Cluster 5	4.61	0.57
Cluster 6	4.92	0.62
Cluster 7	6.53	0.50
Cluster 8	8	0
Cluster 9	9	0

Table 1: Mean (μ) and standard deviation (σ) of the vertical coordinate of images among clusters.

we propose the following method that includes the texture-based segmentation algorithm of the cardiac region we presented before. The fundamental hypothesis of the method lies in the fact that the contraction of the myocardium is correlated to the fineness of its texture: the more the myocardium will be contracted, the finer its texture will be. This hypothesis makes sense in an anatomical view because the contraction of the myocardium corresponds to a reduction of its volume and its fibers then lie more closely.

To better capture the myocardium texture, we first consider the cardiac region from the thoracic images, by selecting a region of interest shown in figure 3. The algorithm then segments this region into three classes corresponding to the myocardium, the blood and the background, as presented in section 2.4.

We then considered the myocardium class and computed a parameter ρ corresponding to the evaluation of the coarseness of the myocardium texture. ρ is computed from the Gabor texture features of the barycenter of the myocardium class resulting from the k-nearest neighbors algorithm. It is defined as the quotient between the sum of the means of the magnitude responses to the low frequency Gabor filter and the sum of the means of the magnitude responses to the high frequency Gabor filter:

$$\rho = \frac{\sum_{i=0}^5 \mu_{0,i \times 30}}{\sum_{i=0}^5 \mu_{1,i \times 30}}$$

where $\mu_{0,\theta}$ denotes the mean of the magnitude response to the low frequency Gabor filter at the orientation θ and $\mu_{1,\theta}$ denotes the mean of the magnitude response to the high frequency Gabor filter at the orientation θ .

Figure 5 shows a part of the sequence we indexed using this ρ parameter. Figure 6 presents the evolution of the mean of ρ on the slices of a volume along time. The sequence consistently present a cardiac cycle evolution.

3.5 Results evaluation

Image indexing and database clustering. To quantitatively evaluate the relevance of the feature vectors-based clustering of the database, we computed the mean and standard deviation of the vertical coordinate of the images in each cluster. Table 1 presents the values we obtained. The standard deviations are all inferior to 0.84 (their mean is




Figure 5: Extraction of the cardiac sequence we indexed using the parameter ρ . One can see the evolution of a slice of this 3D+t sequence along time. Systole occurs between instants t=0 and t=5. Then diastole begins.




Figure 6: Evolution of the parameter ρ in a cardiac sequence along time. ρ effectively reaches a minimum around the end of systole instant.

0.44), which indicates a relatively low dispersion among clusters. In clusters 4, 8 and 9, standard deviation is even null, which means that all the images of the concerned clusters have the same vertical coordinate in the initial sequences. Each of the other clusters contains an homogeneous set of images having close vertical coordinates in the initial sequences. This clustering result demonstrates the ability of our system to gather images of the database having the same vertical position in the initial sequences.

Image retrieval. To quantify the accuracy of the retrieval procedure, we compare the known vertical position of the slices (ground truth) with the one of the computed results. We define the *vertical score* of a query as the mean of the vertical distances between the query image and the N first results of the query, balanced with the ranks of the results. Thus, the *vertical score* is defined as follows:

$$\text{score}_v = \frac{\sum_{i=1}^N (N+1-i) |z_{\text{query}} - z_{\text{result}}^i|}{\sum_{i=1}^N (N+1-i)}$$

where z_{query} is the vertical position of the query image in its original sequence and z_{result}^i is the vertical position of the i^{th} result in its original sequence.

Note that result number 0 is always the query image. That is why we did not include it into score_v . We computed this score for the 170 images of our database and with $N=8$. We obtained a mean *vertical score* of 0.23. It means that given a query image in the database, the mean vertical distance to the 8 best results of the query is 0.23 slices. Taking into account that the two original sequences have a vertical dimension of respectively 10 and 7, we can conclude that our system is pertinently able to retrieve images of the database having the same vertical position than the query image using unsupervised texture features only.

Segmentation assisted retrieval. We can notice that the curve of figure 6 shows a minimum around $t=5$, which corresponds to the end of systole instant indeed, as it can be seen in figure 5. Thus, the parameter ρ is consistently indexing cardiac sequences and enables retrieval of the end of systole instant.

4. CONCLUSIONS

Although still preliminary this work demonstrates the relevance of texture-based features for medical images indexing and retrieval. We have analyzed the properties of medical images. Global features are hardly usable for medical image indexing and therefore we have experimented local Gabor filter-based features extraction on cardiac MRIs. In the future, an extension of that kind of filter to 3D and/or multimodal images would be interesting to benefit from all information available in medical images.

Automatic segmentation of medical images would be of great interest for region-based analysis of the images. Although generic and automatic segmentation tools are far from being available, indexing and classification is an interesting method to obtain an unsupervised rough image partitioning suited for assisting local feature extraction in images.

The experiment described in section 3 shows that clinically relevant information can be extracted out of the images without any user supervision nor fine parameter tuning for a specific purpose. Local features extraction on a per-tissue basis enables the study of clinically relevant parameters such as the myocardium contraction evolution along time. It demonstrates that some semantic interpretation of the image can be accomplished through low-level textural information extraction and classification. Texture filters are therefore interesting for accomplishing general purpose medical image indexing although high level query tools will need to add domain-specific knowledge to perform queries responding to specific clinical requests.

5. REFERENCES

- [1] J. Bach, C. Fuller, A. Gupta, A. Hampapur, B. Horowitz, R. Humphrey, and R. Jain. Virage image search engine : An open framework for image management. In *SPIE Storage and Retrieval for Image and Video Databases*, volume 2670, pages 76–87, San Diego/La Jolla, CA, Jan. 1996.
- [2] A. Bovik, M. Clark, and W. Geisler. Multichannel texture analysis using localized spatial filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:55–73, 1990.
- [3] C. Carson, M. Thomas, S. Belongie, J. Hellerstein, and J. Malik. Blobworld : A System for Region-Based Image Indexing and Retrieval. In *Third International Conference on Visual Information Systems (VISUAL '99)*, volume 1614, pages 509–516, Amsterdam, The Netherlands, June 1999.
- [4] W. Chu, C. Hsu, A. Cardenas, and R. Taira. A knowledge-based image retrieval with spatial and temporal constructs. *IEEE Transactions on Knowledge and Data Engineering*, 10(6):872–888, 1998.
- [5] D. Clausi and M. Jernigan. Designing gabor filters for optimal texture separability. *Pattern Recognition*, 33:1835–1849, Jan. 2000.
- [6] D. Comaniciu, P. Meer, D. Foran, and E. Medl. Bimodal system for interactive indexing and retrieval of pathology images. In *Workshop on Applications of Computer Vision*, pages 76–81, Princeton, NJ, Oct. 1998.
- [7] M. Flickner, H. Sawhney, W. Niblack, J. Ashley,

- Q. Huang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D. Petkovic, D. Steele, and P. Yanker. Query by image and video content : The QBIC system. *IEEE Computer*, 28(9):23–32, 1995.
- [8] R. Haralick, K. Shanmugam, and I. Dinstein. Textural features for image classification. *IEEE Transactions Systems on Man and Cybernetics*, 3(6):610–621, 1973.
- [9] T. S. Huang, S. Mehrotra, and K. Ramchandran. Multimedia Analysis and Retrieval System (MARS) project. In *Proceedings of 33rd Annual Clinic on Library Application of Data processing - Digital Image Access and Retrieval*, pages 100–117, Urbana/Champaign, Illinois, 1996.
- [10] A. K. Jain and R. Dubes. *Algorithms for clustering data*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1988.
- [11] A. Kak and C. Pavlopoulou. Content-Based Image Retrieval from Large Medical Databases. In *3D Data Processing, Visualization, Transmission*, Padova, Italy, June 2002.
- [12] P. Kelly and T. Cannon. Query by image example: the candid approach. In *Storage and Retrieval for Image and Video Databases III*, volume 2420, pages 238–248. SPIE, 1995.
- [13] F. Korn, N. Sidiropoulos, C. Faloustos, E. Siegel, and Z. Protopapas. Fast and effective retrieval of medical tumor shapes. *IEEE Transactions on Knowledge and Data Engineering*, 10(6):889–904, 1998.
- [14] T. Lehmann, B. Wein, J. Dahmen, J. Bredno, F. Vogelsang, and M. Kohnen. Content-Based Image Retrieval in Medical Applications : A Novel Multi-Step Approach. In *International Society for Optical Engineering (SPIE)*, volume 3972(32), pages 312–320, Feb. 2000.
- [15] W. Y. Ma and B. S. Manjunath. Netra: a toolbox for navigating large image databases. In *International Conference on Image Processing (ICIP)*, pages 568–571, Oct. 1997.
- [16] B. Manjunath, P. Wu, S. Newsam, and H. Shin. A texture descriptor for browsing and similarity retrieval. *Journal of Signal Processing : Image Communication*, 16(1-2):33–43, 2000.
- [17] J. Montagnat, V. Breton, and I. Magnin. Partitionning medical image databases for content-based queries on a grid. *to appear in Methods of Information in Medicine*, 2004.
- [18] H. Müller, N. Michoux, D. Bandon, and A. Geissbuhler. A review of content-based image retrieval systems in medical applications - clinical benefits and future directions. *International Journal of Medical Informatics*, 73(1):1–23, Feb. 2004.
- [19] A. Pentland, R. Picard, and S. Sclaroff. Photobook : content-based manipulation of image databases. *International Journal of Computer Vision*, 18(3):233–254, June 1996.
- [20] A. Roche, G. Malandain, X. Pennec, and N. Ayache. The Correlation Ratio as a New Similarity Measure for Multimodal Image Registration. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 1496 of LNCS, pages 1115–1124, Cambridge, USA, Oct. 1998. Springer.
- [21] D. Russakoff, C. Tomasi, T. Rohlfing, and C. Maurer, Jr. Image similarity using mutual information of regions. In *ECCV04*, pages Vol III: 596–607, 2004.
- [22] C. Shyu, A. Kak, A. Kosaka, A. Aisen, and L. Broderick. ASSERT : A physician-in-the-loop content-based image retrieval system for HRCT image databases. In *Computer Vision and Image Understanding*, volume 75, pages 111–132, 1999.
- [23] J. Smith and S. Chang. Automated binary texture feature sets for image retrieval. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 4, pages 2239–2242, May 1996.
- [24] J. Smith and S. Chang. Visualeek : A fully automated content-based image query system. In *ACM International Conference on Multimedia '96*, pages 87–98, Boston, MA, USA, Nov. 1996.
- [25] H. D. Tagare, C. C. Jaffe, and J. Duncan. Medical Image Databases: A Content-based Retrieval Approach. *J Am Med Inform Assoc*, 4(3):184–198, 1997.
- [26] H. Tamura, S. Mori, and T. Yamawaki. Texture features corresponding to visual perception. *IEEE Transactions on Systems, Man and Cybernetics*, 6(4):460–473, 1976.
- [27] M. Tuceryan and A. Jain. *The Handbook of Pattern Recognition and Computer Vision (2nd Edition)*, chapter 2.1, pages 207–248. World Scientific Publishing Company, C.H. Chen, L.F. Pau, and P.S.P. Wang Eds., 1998.