Content-Based Image Retrieval in Medical Applications: A Novel Multi-Step Approach

Thomas M. Lehmann[a1], Berthold Wein[b], Jörg Dahmen[c], Jörg Bredno[a], Frank Vogelsang[b], and Michael Kohnen[b]

[a] Institute of Medical Informatics, Medical Faculty
[b] Department of Diagnostic Radiology, Medical Faculty
[c] Institute for Computer Science VI, Faculty of Computer Sciences
Aachen University of Technology (RWTH), D - 52057 Aachen, Germany

ABSTRACT

In the past few years, immense improvement was obtained in the field of content-based image retrieval (CBIR). Nevertheless, existing systems still fail when applied to medical image databases. Simple feature-extraction algorithms that operate on the entire image for characterization of color, texture, or shape cannot be related to the descriptive semantics of medical knowledge that is extracted from images by human experts.

In this paper, we present a novel multi-step approach, which is specially designed for content-based image retrieval in medical applications (IRMA). In contrast to common approaches, the IRMA-concept is based on a conceptual and algorithmic separation of: (a) image categorization using global features, (b) geometry and contrast registration with respect to prototypes within the categories, (c) extraction of local features, (d) category and query dependent local feature selection, (e) index generation resulting in hierarchical multi-scale blob representations, (f) object identification that links a-priori knowledge on image content to the blobs, and (g) image retrieval processed on the abstract blob-level.

The IRMA-concept comprises several benefits when compared to existing CBIR-systems. The image categories enable semantics by prototypes. Furthermore, each image might belong to several categories. A-priori knowledge on both image and query content is adjuncted to indexing. Therefore, the IRMA-concept provides a high amount of content understanding and enables intelligent queries on an abstract level of information. Hence, IRMA promises satisfactory query completion in medical applications.

Keywords: Content-Based Image Retrieval (CBIR), Content-Based Indexing, Medical Image Databases, Image Content, Registration, Classification, Categorization, Picture Archiving and Communication Systems (PACS)

1. INTRODUCTION

The importance of digital image retrieval techniques increases in the emerging fields of medical imaging and picture archiving and communication systems (PACS). Up to now, textual index entries are mandatory to retrieve medical images from hospital archives or other sources [21]. This also holds for digital archives in DICOM-format [8]. Contrary, information contained in medical images differs considerably from that residing in alphanumeric format [17].

The majority of today's content-based image retrieval (CBIR) approaches is intended for browsing variegated databases (QBIC [11], Photobook [12], Blobworld [1]), e.g. collected from the World Wide Web (VisualSeek [14], WebSeek [15]). For thorough collections of techniques we refer to special issues of notable journals, such as IEEE Transactions on Pattern Analysis and Machine Intelligence (vol. 18, no.8, 1996), IEEE Transactions on Knowledge and Data Engineering (vol. 10, no. 6, 1998), Computer Vision and Image Understanding (vol. 75, no. 1-2, 1999), and Image and Vision Computing (vol. 17, no. 7, 1999), or comparative surveys, e.g. published by De Marsicoi et al. [9].

Resulting from the variegation of images, common CBIR-systems have low data-entry costs and, consequently, only a rudimentary understanding of image content. Within such systems there is no distinction between important and unimportant features or between multiple objects in the image. The features used for indexing characterize the entire image

[1] Correspondence: Email: lehmann@computer.org; Tel: +49 241 80 88793; Fax: +49 241 8888 426.

Part of the IS&T/SPIE Conference on Storage and Retrieval for Media Databases 2000
San Jose, California • January 2000
SPIE Vol. 3972 • 0277-786X/00/$15.00

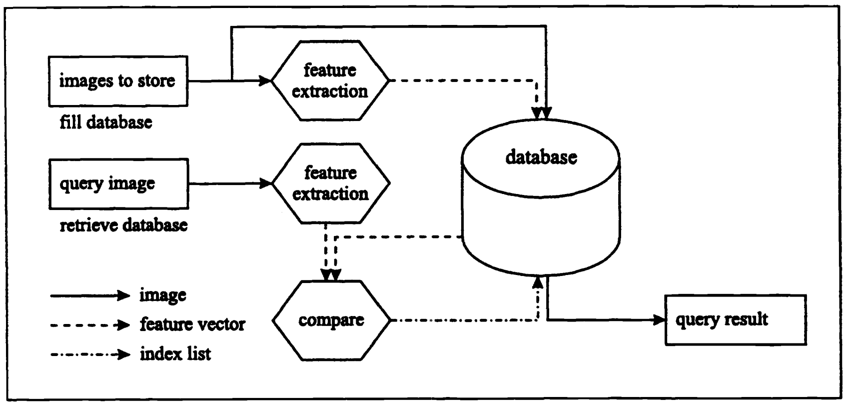

Figure 1: General architecture of CBIR-systems.

rather than unique regions or objects. In contrast, queries of medico-diagnostic relevance include searching for organs, their relative locations, and other distinct features such as morphological appearances. Therefore, common CBIR-systems cannot guarantee a meaningful query completion when used within the medical context [4].

The next section summarizes general properties of CBIR-systems. In contrast, some of the unique challenges in medical images and the resulting constraints to medical image retrieval are stressed in Section 3. In Section 4, we describe our multi-step approach to perform content-based image retrieval in medical applications (IRMA). Advantages as well as limits of the IRMA-concept are discussed in the last section of this paper with respect to system architecture, incorporation of medical knowledge, and modeling of semantic levels.

2. GENERAL SYSTEMS FOR CONTENT-BASED IMAGE RETRIEVAL

Systems for content-based image retrieval have been introduced in the early 1990s [11]. The structure of all the early CBIR-approaches is similar. Images are stored together with features, which usually are extracted automatically from

- color or gray-scale histograms
- texture
- shape

These features are used to describe the entire image and provide a dramatic reduction of data, which is required to perform online retrieval. For that, the user selects an example or draws a sketch from which the system generates the features. Simple classifiers such as nearest neighbor are used to compare the current feature vector with those stored in the database. Note that image retrieval performed using these query-by-example (QBE) methods equals a classification task.

Figure 1 shows the general structure for CBIR-systems [4]. The semantic model, which is used for stepwise description of knowledge that is brought into image processing, consists of only two layers, the raw data layer and the feature layer. Further knowledge on image content is not considered during QBE-processing. Later generations of CBIR-systems add a third layer to the semantic model. For the content-based retrieval engine (CORE), objects and spatial relationships are described by "concepts" within the so called interpretation layer [22]. In Blobworld, this layer is build from ellipsoides (blobs) representing image regions of uniform color or texture on an abstract level of interpretation [1].

3. MEDICAL CONSTRAINTS TO IMAGE RETRIEVAL

Because global color, texture, or shape analyses are insufficient to characterize medical images, the results are rather poor when common CBIR-systems are used to retrieve medical images [10,12]. Furthermore, two or three semantic levels are insufficient to model medical knowledge for image retrieval. In recent reports, some approaches for content-based retrieval designed to support specific medical tasks have been published. KORN et al. describe a system for fast and effective retrieval of tumor shapes in mammogram X-rays [6]. This approach has certain restrictions on both the images (only mammographic

X-rays) and the features (only tumor shapes) that are supported by the system. Likewise, the ASSERT-system operates only on high resolution computed tomographies (HRCTs) of the lung [16]. A physician delineates the pathology bearing region and marks a set of anatomical landmarks when the image is entered into the database. Hence, ASSERT has extremely high data entry costs, which prohibit its application for clinical routine. CHU et al. present a knowledge-based image retrieval system with spatial and temporal constructs [2]. Brain leasions are automatically extracted within 3D data sets from CT and MR. Their representation model consists of an additional knowledge-based fourth layer within the semantic model. This layer provides a mechanism for accessing and processing spatial, evolutionary, and temporal queries. However, those concepts for medical image retrieval are task-specific and not transferrable to other medical applications.

TAGARE et al. point out some of the unique challenges confronting retrieval engines with medical image collections [17]. Medical knowledge arises from anatomic *and* physiologic information, which quite often is obtained by the radiologist simultaneously during the diagnostic process. Hence, regional features are required to support medico-diagnostic queries. However, interpretation of medical images is dependent on both image and query context. Since the context of queries is unknown when images are entered into the database, the database scheme must be generic and flexible. Particularly, the number and kind of features are subject to continuous evolution. Furthermore, medical image interpretation is a complex and poorly understood process. Diagnostic inferences derived from images rest on an incomplete, continuously evolving model of normality. Hence, categorization and registration of medical images is required to support medico-diagnostic queries on a high level of image interpretation.

In medical diagnostics, we distinguish three general applications for automated content-based image retrieval: (a) primitive queries, (b) semantic queries, and (c) browsing:

a) Automatic retrieval of relevant images for follow-up studies within a PACS.
Modern modalities allow submission of textual information about the examination, however, medical staff often does not enter appropriate or sufficient data into the systems. Therefore, text-based retrieval will neither result in complete nor in sufficient information. For example, a physician wants to compare all images of the patient's abdomial region regardless of modality or orientation of the images. Note that this query can be processed successfully by the 3-level CBIR-systems, which have been discussed in the previous section, if they are provided with the prototypes for image categories that are defined for IRMA (see Sec. 4.2). Hence, we name that type of application the *primitive query*.

b) Searching for representative images of known diseases during reporting.
Content-based retrieval by examples of the searched patterns is one of the main activities during diagnostic procedures, especially under unknown circumstances. For example, the radiologist wants to retrieve all images showing fractured tibial bones. Note that this type of query needs processing on the knowledge-based object level (tibial bone with texture of fracture) and could not be performed by those systems discussed in Section 2. This retrieval task is called *semantic query*. However, the support by semantic queries will enhance diagnostic procedures and enforce the welfare of the patient.

c) Scientific and educational studies on pattern obtained from medical imaging.
Time consuming image recalls must be performed to review large amounts of images or studies. Normally, the unspecific primary recall is too large to gain substantial information. Hence, medical knowledge must be adoptable to the system. With pre-selection by a content-based automated retrieval, the query result will be more specific, and larger databases will be processable. It depends on the kind of query whether it must be performed on primitive or semantic levels. However, scanning a database requires mechanisms such as query refinement, where the user is enabled to modify the query successively, and relevance feedback, where the user is enabled to label retrieved images in good ones or bad ones before re-querying. This type of application is called *browsing*.

4. THE MULTI-STEP APPROACH FOR IMAGE RETRIEVAL IN MEDICAL APPLICATIONS

Image retrieval in medical applications (IRMA) requires a system suitable for primitive and semantic queries as well as browsing with no restrictions neither to image category nor to query content. In the following, we present the IRMA-approach for medical CBIR-systems. To enable complex content understanding the IRMA-concept is based on a conceptual and algorithmic separation of seven processing steps:

- categorization (using global features)
- registration (in geometry and contrast)
- feature extraction (using local features)

- feature selection (category and query dependent)
- indexing (hierarchical multi-scale blob representation)
- identification (links a-priori knowledge on image content to blobs)
- retrieval (processed on abstract blob-level)

4.1. Categorization

To enable content-based queries for clinical purposes, the information retrieval system must be familiar with the current image class prior to semantic query processing. Different imaging modalities require different image processing methods. For example, seeking bone fractures in endoscopic images is useless, or ultrasound images must be processed different than skeletal radiographs. According to the deliverable of the USINT working group of the EurIPACS AIM project [20], the IRMA-approach defines three major classes:

- (I) image modality (physical)
- (II) anatomic region (anatomical)
- (III) body orientation (technical)

For example, nine anatomic regions are distinguished within the first level of region code: (1) upper limb, (2) lower limb, (3) pelvis, (4) spine, (5) abdomen, (6) chest, (7) skull, (8) mamma, and (9) total body. These instances build subclasses resulting in hierarchical structured categories (Fig. 2).

Modern DICOM-devices for imaging comprise all information required for IRMA-categorization. Nevertheless, automatic categorization procedures are required for fast archiving of secondary digitized images that were acquired by film-based modalities or non-DICOM digital devices. At this early stage of processing, no information on image content is available. Hence, features characterizing the entire image must be used for image categorization. Inherent properties of color and gray-scale distribution as well as spatial resolution are reflected by *global* measures, such as histograms, cooccurrence matrices, median of gray-scales or colors, and textures [18]. Global features are not necessarily based on simple image value statistics. They might also include specially adapted feature extraction methods and already may include content-based knowledge from prototypes (e.g. shape). For example, ultrasound images of large organs with relatively uniform tissue, such as spleen or liver, present homogeneous image patterns, which are recognizable by Fourier analyses. Microscopic histology images possess unique color signatures and cell textures that are suggested by other researchers to group slices that share similar staining techniques [17].

The automatic categorization of images must not be unique. The IRMA-approach allows each image to be linked to several categories. This can be formulated as one-to-many mapping from images to categories. Furthermore, different images will share the same category, which is captured by many-to-one mapping from images to categories. Hence, data sets and categories are related many-to-many. There are two main methods that are commonly used for many-to-many mapping. One method bases on statistical analyses. A set of samples is used to determine the center of clusters in the feature space. Categories are assigned to cluster centers using minimal distance measures. However, this technique requires a representative set of training data. The other method bases on the definition of exemplary prototypes. The feature vectors are designed such that they best discriminate these prototypes. Because the IRMA-categories are not drawn from statistical data analysis but from modeling the reality of image acquisition, we apply the second strategy. In other words, IRMA-categories are treated as fuzzy sets and fuzzy membership functions are defined for mapping from a category-specific set of global features. The likelihood for each category also is stored in the IRMA-database. Again, medical a-priori knowledge is used for the heuristic optimization of fuzzy categorization that must be rather sensitive than specific.

The number of entries in the global feature vector is extensible. Whenever a new global feature extraction algorithm is incorporated to IRMA, the corresponding vector components are calculated for all images of the IRMA-database. This is done automatically by parallel batch processing on all IRMA-workstations [3].

4.2. Registration

As mentioned above, diagnostic inferences derived from images are deduced from an incomplete but continuously evolving model of normality. In the IRMA-system, this model is represented by a prototype image, which is defined for each category by an expert based on medical a-priori knowledge. The prototype is used for determination of geometric registration parameters, particularly, rotation, scaling, and translation (RST)-movements as well as parameters for contrast adjustment. These parameters are determined to enable further processing for object identification and retrieval. They are

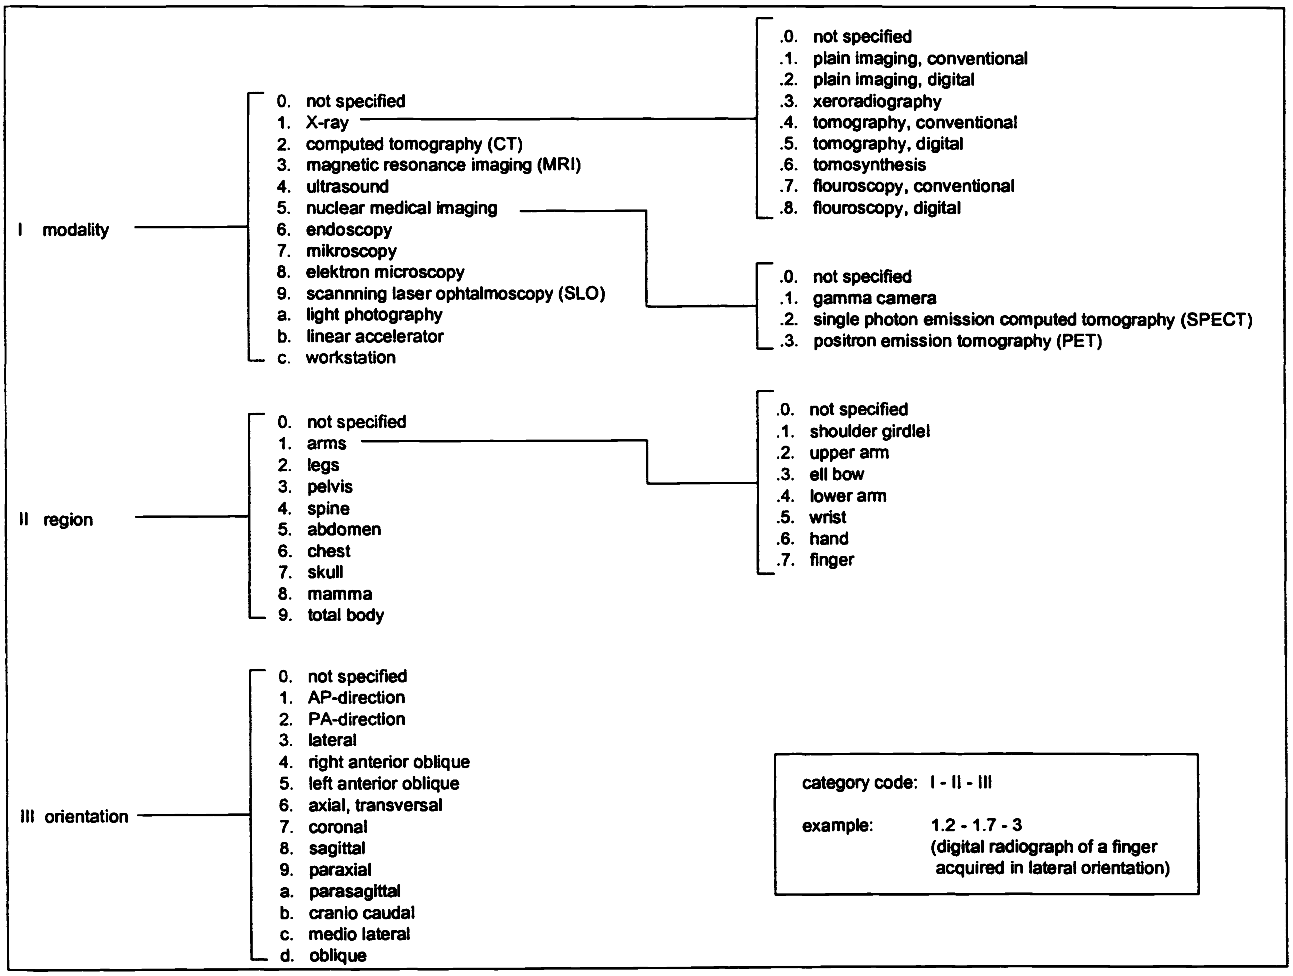

Figure 2: Images are categorized into the three major classes: modality, region, and orientation. The first levels of each of the hierarchically structured classes are given completely while the second level is exemplified only.

stored together with the probabilities for the most likely categories given by the current image. Note that the image is not geometrically transformed at this stage of processing. Since the set of parameters for registration will differ for all likely categories, the interpolation of an image into the registered geometry of each category would require the subsequent algorithms to operate on multiple images. This would result in an unmanageable amount of processing as well as storage costs.

Again, *global* algorithms are used to determine the RST-parameters. The determination of more complex transformation models or warping parameters is redundant with respect to the interpretation of hierarchical blob-structures that are extracted from the registered local features. Symmetric phase only matched filtering or cepstrum processing of Fourier-Mellin power spectra [7] as well as techniques based on Lie derivatives are common approaches and used in IRMA for spatial registration. Contrast registration is done on demand by simple histogram-based algorithms. However, a registered copy of the image may be required for contrast adjustment or specialized feature image extraction methods.

4.3. Feature extraction

A large part of information in medical images is geometric. Hence, the processing of semantic queries drawn from medical routine requires *local* features. The global feature vectors used for categorization are related to the entire image but local features are now connected to pixels or regions resulting in feature images. For example, the Sobel edge detector is one of many algorithms for feature image extraction. The number of local feature images is not limited and extended according to diagnostic requirements.

The feature images strongly extend the data volume. Since they are obtained from an unregistered image they are calculated only once but are suitable for all likely categories (category-free features). Furthermore, category-specific features are used because not all of the feature extractors are applicable to images of either category. These category-specific features include segmentation results obtained by active contours or active shapes [19], which allow the use of a-priori shape knowledge from the categories.

Likewise the global features for categorization, the number of local feature images is extensible freely. Whenever a new local feature extraction algorithm is incorporated to IRMA, the corresponding feature image is calculated for all images within respective categories. Again, this is done automatically by parallel batch processing [3]. Hence, the IRMA-concept regards sufficient flexibility to the database scheme and takes care of evolving processes in medical diagnostics.

4.4. Feature selection

It is important to note that the IRMA-concept allows the de-coupling of feature extraction and selection. This enables the latter task to be retrieval-dependent. A-priori knowledge about the category as well as medical knowledge incorporated to the query is used to select a set of proper feature images. For example, the retrieval of radiographs with respect to bone fractures is done using edge-based feature images (form set). The same radiograph is described by another set of feature images (texture set) if the query concerns bone tumors. However, queries which operate on a certain subset of feature images require processing of the subsequent indexing procedure for all images in the database. This consumes too much time with respect to online queries. Hence, the IRMA-concept comes with a few pre-defined sets combining distinct feature images for each category. Finding these feature sets will be done using statistical feature reduction methods such as Fisher's linear discriminant analysis.

4.5. Indexing

So far, storing an image in the database includes categorization, determination of registration parameters for each category, feature image extraction and selection of some distinct sets in each category. For query processing, this amount of information must be drastically reduced. In general, indexing can be described as the search for an element in the database on the basis of reduced information. For example, reduced information in the setting of a text database is (in ascending order of reduction): the text itself, its abstract, and keywords. Analogously, the indexing of images in IRMA is based on abstractions of the image. Reduced to sets of feature images, the image is segmented into relevant regions. All regions are described by invariant moments resulting in structures called blobs. For example, the blobs equal ellipses if only first and second moments are used for description. Each blob is characterized by a N-dimensional feature vector formed by the medium of the N feature images within the region.

Thereafter, the blob-representation of the image is adjusted with respect to the parameters determined in the registration step. The registration of blob-structures with respect to corresponding image categories avoids those problems that are usually connected to comparisons on blob-structure level. For example in Blobworld, retrieval might fail if a geometric transform distorts size or relations of blobs. In addition, the blob-extraction process is merged with a multi-scale approach. Blob-representations are generated from several resolutions obtained by the segmentation process. This enables the user to retrieve from entire images as well as from regions of interest (ROIs). In summary, this step of processing leads from the pixel domain onto an abstract level of information (i.e. scheme layer, see Section 5).

4.6. Identification

The hierarchical blob-tree assigns blob-structures to dominant regions in different scales of medical images that have been registered with respect to a certain category and prototype. If the local features for blob-generation reflect characteristic and discriminant properties of tissue, certain blobs will correspond to well defined morphological structures in the image, such as organs, bone, or others. Vice versa, medical a-priori knowledge on the content and structure of category prototypes can be used to build a prototype blob-structure. The entities of the prototype blob-structure contain properties of the corresponding image region, such as location, size, shape, and texture, as well as labels from a medical nomenclature. While the characteristic properties are used for blob-identification, the labels are used for semantic queries.

Note that the labeled prototype blob-trees are an applicable data structure to import medical knowledge into the IRMA-system. Knowledge on physiological occurrence of morphological structures within the human body is of significant value to allow the processing of medico-diagnostic queries. Hence, blob-identification provides content understanding at a morphological level. However, blob-identification will not be successful for all blob entities.

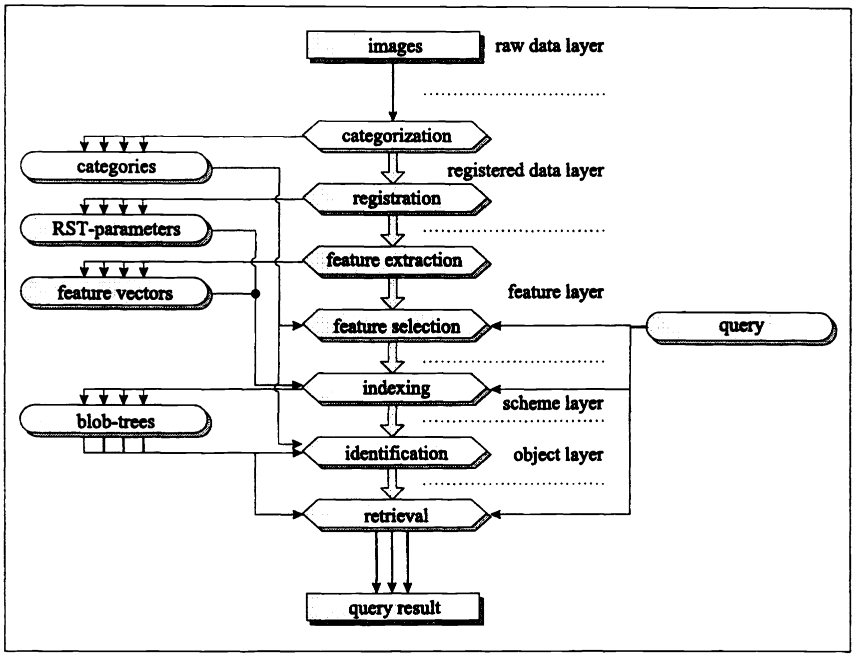

Figure 3: Processing modules and semantic layers of the IRMA-system

4.7. Retrieval

Retrieval is performed by searches in the hierarchical blob-structures. In IRMA, a query is build from the following components:

- a list of possible categories of the recall images
- a QBE-blob-structure on the optimal scale to process the query
- the set of local features that best describe the significant properties for the query (e.g., texture set, form set)

However, not all of these components need to be set for all queries. For a primitive query, the blob-structure is not used as it is searched for a similarity in global features. For semantic queries, the user has to describe the blob-structure of interest. This can either be done using the identification of blobs and, therefore, assigning properties that are queried to morphological structures, or, by providing an example image with a choice of significant blobs for the query. The search is performed by comparison of blob-structures with respect to the blobs of interest and selected features.

The assignment of categories to the images significantly increases the performance of retrieval as only the blob-structures for images of the possible categories have to be compared. Data structures and distance measures that have been described for image retrieval [5] need refining so that only selected blobs and features are relevant in the query. Note that the multi-scale approach for blob-representation efficiently integrates queries restricted to image details (ROIs). Furthermore, blob-identification supports semantic queries, which concern defined organs or other labeled objects in the image. Nevertheless, methods must be provided to extend this choice with qualitative descriptors like "larger", "more intense", or "between these two examples", and combination of these. The formulation of such queries requires intense work on the design of the user interface. However, the current state of development of the IRMA-system is not focused on this work.

5. DISCUSSION

Figure 3 summarizes the multi-step approach of the IRMA-system. The seven processing steps (1) categorization, (2) registration, (3) feature extraction, (4) feature selection, (5) indexing, (6) identification, and (7) retrieval are sequentially combined. These processing steps correspond to five semantic layers for knowledge representation. Likewise other systems, the unprocessed images form the first layer, which is called the *raw data layer*. Categorization and registration within each category are the first level where medical knowledge is incorporated into the IRMA-system. Hence, both steps result in the second semantic layer, which is called the *registered data layer*. While other medical CBIR-systems are restricted to a certain modality or diagnostic procedure [2,6,16], the registered data layer in IRMA allows queries across all kind of medical images regardless of modality, region, or orientation. The *feature layer* is the third level of knowledge-based image processing. The separation of local feature extraction and feature selection is the major advantage when compared to other systems. This separation makes the feature layer query dependent. Note that in contrast to other authors, spatial relationship characteristics are not modeled in the feature layer [2]. The forth layer is obtained from indexing. We use the nomenclature introduced by CHU et al. and call it the *scheme layer* [2]. In the scheme layer, blob-structures represent the entities and the spatial relationships among them. The modeling of spatial relationships is emphasized by the hierarchical blob-concept. Hierarchies not only provide an efficient way to focus on ROIs but also enable the introduction of query-specific knowledge to the processing. The fifth layer results from object identification. Therefore, we call it the *object layer*. The object layer contains detailed knowledge on image content and hence, it is also referred to as knowledge layer [2]. Keep in mind that those names are not used uniquely in the literature.

TAGARE et al. postulate the necessity of several tools for retrieval in medical imaging databases: (1) non-textual indexing, (2) customized scheme, (3) dynamic modules, (4) similarity modules, (5) comparison modules, (6) iconic queries, (7) descriptive language, (8) multi-modality registration, (9) image manipulation [17]. Although all CBIR-systems provide non-textual indexing, the IRMA-approach is specially designed to handle primitive and semantic queries as well as browsing of medical images with respect to medico-diagnostic applications. Hierarchical blob-structures are build from selected local features. However, new feature extraction methods can be added anytime to the IRMA-system. The IRMA-system provides a development platform in which deamons automatically poll job lists from distributed systems and recreate the index of database entries by parallel processing [3]. Hence, the IRMA blob-structure equals a customized and dynamic scheme in the sense of TAGARE et al. We agree with the authors that similarity and comparison modules must be available within the descriptive language. Those concepts fit the IRMA-approach although they have not been implemented yet. Iconic queries cover image examples, sketches or prototypes. They are directly supported within IRMA by the incorporated QBE-concept. Because prototypes are part of the IRMA-system, their importance is even more emphasized by the IRMA-approach when compared to the approach by TAGARE et al. This also holds for registration. In IRMA, registration of images is performed automatically while TAGARE et al. only suggest to provide interactive tools for registration and image manipulation.

6. SUMMARY

We have introduced a multi-step approach for content-based image retrieval in medical applications. In general, the IRMA-concept is related to the Blobworld-project [1]. However, there are several important extensions of the Blobworld-concept especially designed for medical purposes:

- Each of the processing steps uses a more conceptional formulation of a-priori knowledge and it hierarchically regards details from global to local image properties.
- Medical images are categorized to enable content-based processing.
- Each image belongs to several categories with different probability.
- Several blob representations are generated since several sets of feature images exist for each category. The selection depends on both query and image content.
- Each blob-representation is hierarchically ordered on different resolutions (multi-scale) supporting queries on entire images and regions of interest.
- The blob-representations are registered to a prototype. This enables efficient image comparison on blob-level as well as easy identification of image structures.

By these extensions, a-priori knowledge on both image and query content is adjunct to content-based image indexing. Therefore, the IRMA-concept provides a high amount of content understanding and enables highly differentiated queries on an abstract information level. Furthermore, the IRMA-concept fulfills the demands for medical image retrieval systems postulated by TAGARE et al. and therefore, IRMA promises satisfactory query completion [17].

REFERENCES

1. C. Carson, M. Thomas, S. Belongie, J.M. Hellerstein, and J. Malik, "Blobworld: A System for Region-Based Image Indexing and Retrieval," *Procs. 3rd Int. Conf. On Visual Information Systems*, Springer-Verlag, 1999.
2. W. W. Chu, C.-C. Hsu, A. F. Cárdenas, and R. K. Tiara, "Knowledge-Based Image Retrieval with Spatial and Temporal Constructs," *IEEE Trans. on Knowledge and Data Engineering*, **10(6)**, pp. 872-888, 1998.
3. J. Bredno, M. Kohnen, J. Dahmen, F. Vogelsang, B. Wein, and T. M. Lehmann, "Synergetic Impact Obtained by a Distributed Developing Platform for Image Retrieval in Medical Applications (IRMA)," *Proceedings SPIE* **3972(33)**, (this issue).
4. J. Dahmen, T. M. Lehmann, K. Spitzer, and H. Ney, "Image Retrieval für klinische Bilddatenbanken," In: T. M. Lehmann, V. Metzler, K. Spitzer, T. Tolxdorff (Hrsg): *Bildverarbeitung für die Medizin 1998*, Springer-Verlag, Berlin, pp. 442-446, 1998.
5. J. Goldstein, R. Ramakrishnan, U. Shaft, and J.-B. Yu, *Using Constraints to Query R*-Trees*, Technical Report No. 1301, Dept. of Computer Science, University of Wisconsin, Madison, Feb. 1996.
6. P. Korn, N. Sidiropoulos, C. Faloutsos, E. Siegel, and Z. Protopapas, "Fast and Effective Retrieval of Medical Tumor Shapes," *IEEE Trans. on Knowledge and Data Engineering*, **10(6)**, pp. 889-904, 1998.
7. T. M. Lehmann, C. Goerke, A. Kaupp, W. Schmitt, and R. Repges, "A Rotation-Extended Cepstrum Technique and its Application to Medical Images," *Pattern Recognition and Image Analysis* **6(3)**, pp. 592-604, 1996.
8. S. L. Lou, D. R. Hoogstrate, and H. K. Huang, "An Automated PACS Image Acquisition and Recovery Scheme for Image Integrity Based on the DICOM Standard," *Computerized Medical Imaging and Graphics* **21(4)**, pp. 209-218, 1997.
9. M. de Marsicoi, K. Cinque, and S. Levialdi, "Indexing Pictorial Documents by Their Content: A Survey of Current Techniques," *Image and Vision Computing* **15**, pp. 119-141, 1997.
10. M. Nappi, G. Polese, and G. Tortora, "FIRST: Fractal Indexing and Retrieval System for Image Databases," *Image and Vision Computing* **16**, pp. 1019-1031, 1998.
11. W. Niblack, R. Barber, W. Equitz, M. Flickner, P. Yanker, and J. Ashley, "The QBIC-Project: Querying Images by Content Using Color, Texture and Shape," *Proceedings SPIE* **1908**, pp. 173-187, 1993.
12. A. Pentland, R. W. Picard, and S. Scarloff, "Photobook: Content-Based Manipulation of Image Databases," *Proceedings SPIE* **2185**, pp. 34-47, 1994.
13. E. G. M. Petrakis and C. Faloutsos, "Similarity Searching in Medical Image Databases," *IEEE Trans. on Knowledge and Data Engineering* **9(3)**, pp. 435-447, 1997.
14. J. R. Smith and S. F. Chang, "VisualSEEK: A Fully Automated Contend-Based Image Query System," *Proceedings ACM Multimedia '96*, Boston, MA, pp. 87-98, 1996.
15. J. R. Smith and S. F. Chang, "Tools and Techniques for Color Image Retrieval," *Proceedings SPIE* **2670**, pp. 426-437, 1996.
16. C.-R. Shyu, C. E. Brodley, A. C. Kak, A. Kosaka, A. M. Aisen, and L. S. Broderick, "ASSERT: A Physician-in-the-Loop Content-Based Retrieval System for HRCT Image Databases," *Computer Vision and Image Understanding* **75(1/2)**, pp. 111-132, 1999.
17. H. D. Tagare, C. C. Jaffe, and J. Dungan, "Medical Image Databases: A Content-Based Retrieval Approach," *Journal of the American Medical Informatics Association* **4**, pp. 184-198, 1997.
18. F. Weiler, F. Vogelsang, M. Kilbinger, B. Wein, and R. W. Günther, "Automatic Recognition of Image Contents Using Textural Information and a Synergetic Clasifier," *Proceedings SPIE* **3034**, pp. 985-989, 1997.
19. F. Weiler and F. Vogelsang, "Model-based Segmentation of Hand Radiographs," *Proceedings SPIE* **3338-1**, pp. 673-681, 1998.
20. T. Wendler, B. Wein, and R. van den Broek, *Deliverable USINT in the EurIPACS-Project*, Technical Report, AIM, Ref. XIII, European Commission, Brussels, Feb. 1994.
21. S. T. C. Wong and D. A. Tjandra, "A Digital Library for Biomedical Imaging on the Internet," *IEEE Communications Magazine*, pp. 84-91, January 1999.
22. J. K. Wu, A. D. Narasimhalu, B. M. Mehtre, C. P. Lam, and Y. J. Gao, "CORE: A Content-Based Retrieval Engine for Multimedia Information Systems," *Multimedia Systems* **3**, pp. 25-41, 1995.