

Content-based Image Retrieval in Medical Applications

T. M. Lehmann¹, M. O. Güld¹, C. Thies¹, B. Fischer¹, K. Spitzer¹, D. Keyzers², H. Ney², M. Kohnen³, H. Schubert³, B. B. Wein³

¹Department of Medical Informatics, Aachen University of Technology (RWTH), Aachen, Germany

²Chair of Computer Science VI, Aachen University of Technology (RWTH), Aachen, Germany

³Department of Diagnostic Radiology, Aachen University of Technology (RWTH), Aachen, Germany

Summary

Objectives: To develop a general structure for semantic image analysis that is suitable for content-based image retrieval in medical applications and an architecture for its efficient implementation.

Methods: Stepwise content analysis of medical images results in six layers of information modeling incorporating medical expert knowledge (raw data layer, registered data layer, feature layer, scheme layer, object layer, knowledge layer). A reference database with 10,000 images categorized according to the image modality, orientation, body region, and biological system is used. By means of prototypes in each category, identification of objects and their geometrical or temporal relationships are handled in the object and the knowledge layer, respectively. A distributed system designed with only three core elements is implemented: (i) the central database holds program sources, processing scheme descriptions, images, features, and administrative information about the workstation cluster; (ii) the scheduler balances distributed computing; and (iii) the web server provides graphical user interfaces for data entry and retrieval, which can be easily adapted to a variety of applications for content-based image retrieval in medicine.

Results: Leaving-one-out experiments were distributed by the scheduler and controlled via corresponding job lists offering transparency regarding the viewpoints of a distributed system and the user. The proposed architecture is suitable for content-based image retrieval in medical applications. It improves current picture archiving and communication systems that still rely on alphanumerical descriptions, which are insufficient for image retrieval of high recall and precision.

Keywords

Information storage and retrieval, information management, pattern recognition, computer-assisted image processing

Methods Inf Med 2004; 43: 354–61

1. Introduction

Content-based image retrieval (CBIR) aims at describing the complex object information of digital images by non-textual features, which are applicable for efficient query processing. Color, texture, and shape are used within the first CBIR approaches such as the query by image content (QBIC) system [1]. Besides the raw data layer representing the initial images, QBIC models a feature layer on which the retrieval operates. For the content-based retrieval engine (CORE), objects and spatial relationships are described by “concepts” within the so-called interpretation layer [2]. In Blobworld, this layer is referred to as scheme layer. It is built from ellipsoids (“blobs”) representing local image regions of uniform color or texture on an abstract level of interpretation [3]. However, two or three semantic layers are insufficient to model medical knowledge for image retrieval and consequently, results are rather poor when common CBIR algorithms are applied to medical images [4, 5].

In recent reports, some approaches for content-based retrieval specially designed to support 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 (mammography only) and the features (tumor shapes only), which are supported by the system. Likewise, the automatic search and selection engine with retrieval tools (ASSERT) operates only on high-resolution computed tomography of the lung [7]. A physician delineates the region bearing a pathology 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. Long et al. access a large collection of 17,000 spine radiographs by means of shape analysis, where biomedical categories such as “anterior osteophytes present/not present” are distinguished automatically [8]. The data entry costs are low, but queries are limited to the pre-defined categories. Chu et al. present a knowledge-based image retrieval system with spatial and temporal constructs [9]. Brain lesions are extracted automatically within three-dimensional data sets from computed tomography and magnetic resonance imaging. Their representation model consists of an additional knowledge-based layer within the semantic model. This layer provides a mechanism for accessing and processing spatial, evolutionary, and temporal queries. Nonetheless, those concepts for medical image retrieval are task-specific, i.e. limited to a particular modality, organ, or diagnostic study and, hence, usually not directly transferable to other medical applications.

However, the distinctive characteristics of descriptive text versus self-contained image information in general are well understood and, in particular, their impact on medical image databases intended to support indexing and retrieval is commonly appreciated. Strategies for implementing medical image retrieval systems have been derived, which employ object-oriented iconic queries, semantics by association with prototypes, and a generic scheme [10]. Hence, medical CBIR usually addresses three major problems [11]: (i) which kind of semantics must be modeled by disjunctive

semantic layers, (ii) which kind of features should be used at each stage of abstraction and which distance or similarity measure should be applied, and (iii) how can the data management and computing be organized efficiently?

In this paper, we present an approach for content-based image retrieval in medical applications (IRMA) with particular focus on its semantic layers of information modeling, its hierarchical concept for feature representation and distance computation, and its distributed system architecture for efficient implementation. In addition, we present sample applications to show the general applicability of the concept.

2. The IRMA Approach

2.1 Information Modeling and Steps of Processing

Compared to standard CBIR systems, at least three additional semantic levels of abstraction are needed to cope with the complex medical knowledge that is to be handled by a general system for content-based image retrieval in medical applications. A low level of medical knowledge is determined by the imaging modality including technical parameters, the orientation of the patient position with respect to the imaging system, the examined body region, and the functional system under investigation. Based on prototype images, a mid-level of knowledge is described by regions of interest (ROIs) within the images, and a high level is obtained from information regarding the spatial or temporal relationships of relevant objects. Consequently, IRMA splits the retrieval process into seven consecutive steps (Fig. 1). Each step represents a higher level of image abstraction, reflecting an increasing level of image content understanding [12].

The *categorization* step aims at determining the imaging modality and its orientation as well as the examined body region and functional system for each image entry. For that, a detailed hierarchical coding scheme was developed [13], which exceeds

the complexity of existing tags of the digital imaging and communications in medicine (DICOM) standard, such as (0018/0015) “body part examined” or (0018/5100) “patient position”, but could be consistently integrated to supplement the standard. Automatic categorization is based on a reference database of 10,000 images selected arbitrarily from clinical routine and manually classified by experienced radiologists. This database integrates medical knowledge at a low level of abstraction. The automatic categorization of query-by-example images is performed by combining DICOM header information and global image features, i.e. features describing the entire image. However, categorization in IRMA is not exclusive. Subsequent steps of processing are applied for the most likely categories.

Registration in geometry (rotation, translation, scaling) and contrast generates a set of transformation parameters that is stored for the corresponding image in each of its likely categories. In consent with Tagare et al. [10], registration is based on prototypes which are manually defined for each category, and further incorporate medical expert knowledge into the IRMA system. Note that the transformation is not performed explicitly at this step of processing. Instead, the generated parameters are utilized at higher layers of abstraction.

The *feature extraction* step derives local image descriptions, i.e. a feature value (or a

set of values) is obtained for each pixel. These can be category-free (e.g. resulting from edge detection or regional texture analysis) or category-specific, such as the application of an active shape model that explicitly uses a-priori knowledge derived from the respective category.

Decoupling *feature selection* from feature extraction allows the integration of both the image category and query context into the abstraction process. For instance, the same radiograph might be subject to fracture or cancer examination, resulting in a contour-based or texture-based combination of features, the so-called feature sets, such as the contour set or texture set, respectively. In order to avoid exhaustive computation during query processing, these feature sets are pre-computed for each image in each likely category.

Indexing provides an abstraction of the previously generated and selected image features, resulting in a compact image description. According to the selected feature set, this is done via clustering of similar image parts into regions represented by their second area moment description as ellipses (“blobs”). In contrast to the Blobworld approach [14], this is done at multiple resolutions yielding a multi-scale blob-representation of the image (“blob tree”). Note that hierarchical indexing enables the processing of ROIs, which are marked by the user when issuing a query. In contrast to existing approaches to medical image retrieval, the




Fig. 1 IRMA processing steps (gray, middle) resulting in six semantic layers (dotted lines).

ROI is not determined a-priori. In other words, the incorporation of medical mid-level knowledge becomes possible.

According to Tagare et al. [10], an essential requirement for satisfying medical queries is a high level of image understanding offering object-oriented retrieval. The *identification* step provides linking of medical a-priori knowledge to certain blobs generated during the indexing step. It relies on the prototypes defined for each category, which are labeled locally by medical experts, and the corresponding parameters for geometry and contrast registration. Thus, identification is the fundamental basis to introduce high-level image understanding by analyzing regional or temporal relationships between the blobs.

In IRMA, the *retrieval* itself is processed either on the abstract blob level or referring to identified objects. Note that only the retrieval step requires online computations while all other steps can be performed automatically in batch mode at entry time of an image into the database. This, of course, requires offline computation of all paths generated by the categorization and the feature selection step.

2.2 Feature Representation and Distance Computation

Image categorization is performed by means of *global features*, i.e., a single value or a vector combining a few values is assigned to the entire image. A large number of global features have been proposed in the literature for content-based image retrieval. Besides major components from color and gray scale histograms or the moments of dominant regions, we focus on global measures obtained from (i) frequency, (ii) texture and (iii) structure analysis. Fourier, discrete cosine or wavelet transforms extract suitable characteristics of images [15, 16]. Texture descriptors are obtained from spatial grey-level difference statistics, circular Moran autocorrelation function, entropy and coarseness [17] as well as from the images' fractal dimensions [17, 18]. The structure of an image is an edge-based measure that has been recently

introduced by Zhou and Huang [19]. Based on the Canny edges of an image, statistics on edge length, bifurcation and branching as well as loops are used to characterize the entire image.

The feature values obtained for each image are combined to a feature vector, which is then used for k-nearest-neighbor classification based on e.g. the Euclidean, or Simard's tangent distance [20]. The latter is able to cope with local geometry and contrast differences.

It is important to realize that we do not aim at clustering our feature space in order to find suitable categories but that we apply as many global features as required to distinguish the categories, which are given a-priori. Furthermore, this categorization is not crisp. The most likely categories are tracked through the following steps of processing. Hence, the actual category should be within the most likely ones but does not need to be the likeliest.

Local features are assigned to each image pixel. Regardless of the determined category, local features are extracted uniformly for all images. In other words, all methods for local feature extraction are applied to each image. Besides the gray value itself, anisotropy, polarity, and contrast are determined according to the Blobworld approach [14]. Furthermore, spatial convolution techniques open a great variety of methods. With respect to both image category and query content, suitable feature sets are selected to enable query-adaptive processing without additional computation at the time of query processing.

In contrast to most other CBIR systems, local IRMA features are not used directly to compute a distance measure between two images. Instead, a blob representation is computed first to significantly reduce the amount of data. Following the Blobworld approach [3], dominant image regions are approximated by their best fitting ellipses to which the mean feature vector of the entire region is assigned. While the Blobworld approach applies an expectation maximization clustering technique within the feature space, the partitioning in IRMA is computed in the image domain by means of an edge-preserving region growing algorithm [21]. This ensures connected segments and

a complete image partitioning. In fact, the complete partitioning permits a hierarchical image decomposition to model a multi-scale approach. On the lowest level, each pixel builds its own blob while on the highest level, a single blob represents the entire image. In between, a tree structure of blobs is obtained.

Figure 2 exemplifies the multi-scale abstraction that is modeled by the IRMA system. The image (Fig. 2, left) is partitioned into representations with a decreasing number of regions (Fig. 2, upper row). While the color-coding of regions is initialized randomly, the color of the largest region is maintained during a merge. Dominant regions are represented by their best fitting ellipsoids (Fig. 2, lower row). These ellipsoids form the nodes of the IRMA blob tree representation of images (Fig. 2, right). The strong edges in light blue build the hierarchical tree structure of the graph. Adjacencies within a level of the multi-scale decomposition are coded red. To reduce the total number of nodes to be displayed, stable nodes not repeated in the next level. Instead, the thin blue edges mark adjacencies of nodes crossing the levels of decomposition (e.g., Fig. 2, nodes 5 and 9).

Consequently, the distance or similarity of images or ROIs transforms into a (sub-)graph-matching problem. Several approaches have been suggested as solutions. For instance, similarity flooding is applied to determine corresponding nodes via identical connections to their neighboring nodes using a fix-point iteration [22]. Alternative approaches focus on a direct partial node-to-node mapping, which is solved in polynomial time [23]. More general, one can combine structural matching and the method of solving optimization tasks using Hopfield-style nets to establish a mapping between two graphs that maps similar parts onto each other preserving as much as possible of their overall structural correspondence [24].

2.3 Data Management and Query Processing

Regardless of their semantic layer, all feature extraction and evaluation steps of




Fig. 2 Initial image (left), hierarchical partitioning and corresponding blobs (middle), and resulting graph representation (right)

IRMA are regarded as method transforming features. The *method* joins a program and its parameterization into one. As described in the previous section, IRMA provides a general feature model. It only depends on the method whether it results in value, pixel, or tree data. In particular, value data is numerical values or alphanumeric strings coding global features, pixel data contains images and local features, and tree data holds the hierarchical structured blobs as well as their mean feature vector. Only three types of methods are sufficient to model arbitrary steps of retrieval algorithms: (i) a T:1-method transforms a feature set into a single feature (e.g. generating a matrix for statistical feature reduction by principal component analysis), (ii) a 1:1-method transforms one feature into another (e.g. computing a spatial convolution of an image), and (iii) a 1:T-method transforms one feature into a set of features (e.g. generating multiple representations of one image via small transformations). Queries are modeled as a network of interconnected user-implemented methods with a flexible feature input/output interface. Similar to dataflow process networks, we use directed acyclic graphs composed of methods and control elements [25]. Parameters and information on reference images and prototypes are stored in an experiment description. Since experiments are composed of modules, a structured generation history for each feature allows the automatic identification and re-use of already computed features.

In other words, the manifold structure of information abstraction (Sec. 4.1) is reduced to a small number of mechanisms that can be handled within a distributed client-server architecture (Fig. 3). The cen-

tral components (database, scheduler, web server) run on a server. The client processes (daemons, programs, applications) are used for distributed computing in a cluster [26].

The central relational *database* is used to store administrative information about physical entities, i.e. value, pixel and tree data as well as sources of programs, and logical entities, i.e. methods, networks, and experiment definitions characterizing the algorithms for image processing and retrieval, their parameters, and the feature sets in use. The physical entities are stored as files outside the database and can be hosted by any computer within the distributed system. Using the information about the cluster infrastructure, transparent access to and automatic replication of all physical entities are implemented.

The *scheduler* is a central service that manages the execution of all queries or feature extraction tasks. It uses transmission control protocol (TCP) sockets to communicate with all daemons and the programs running in the cluster. The scheduler has two functional parts. For each invoked query, the process control sub-part creates a data structure to log the progress during the execution of the corresponding experiment. For each node of the network ready to be executed, a job is generated, which includes the method identifier (ID), the IDs of the input features and their locations, and allocated IDs for the output features. The communication subpart assigns jobs to programs running in the cluster. If the program needed by a job is not running, the scheduler selects an appropriate host and

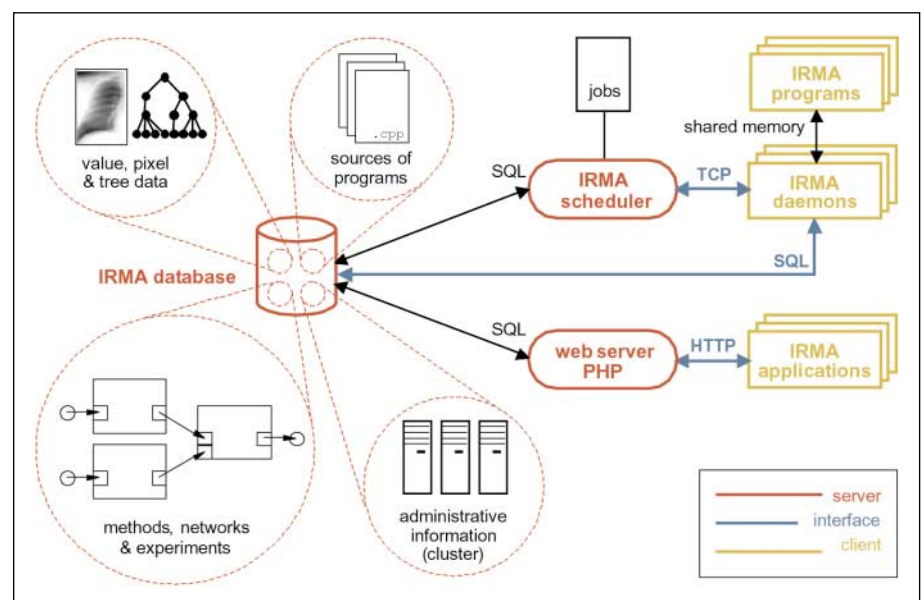


Fig. 3 The IRMA system architecture is composed of three core elements, which are marked in red color. Processes on the clients and the interfaces are displayed in yellow and blue, respectively.

issues the IRMA daemon on this machine to start the program. Additionally, the scheduler can order idle programs to terminate in case this program will not be required within a look-ahead interval. Upon the completion of a job, the scheduler receives a notification from the program and updates the feature information in the database.

Each computer within the system runs an IRMA *daemon*. This background process automatically installs new programs on its host and starts them on demand. The daemon service is also used to inform the scheduler about the current load of its host. Furthermore, it can detect possible abnormal terminations of programs and report them to the scheduler for fault handling. The daemon also performs the automated program transfer: By evaluating the versioning information from the database, the daemon can determine if a program needs to be installed or updated. Due to the heterogeneous cluster infrastructure, programs are transferred as source code. Once the transfer is complete, a make-file is generated and executed, resulting in the program's executable.

The physical executable of a *program* entity consists of a user-implemented subroutine linked to a generic main routine which handles all system communication via the daemon: receiving a job from the IRMA scheduler, loading the parameterization and the input data, starting the subroutine, and storing the resulting output features. Each client hosts a default location where large data objects are stored. If required input data is locally unavailable, a replication is retrieved via the daemon using the same mechanisms as described before for the automated program transfer. Once a program is started on a client, it can compute a sequence of individual data objects with optimized memory access. For instance, communication with the daemon interfacing the IRMA system to the pro-

gram is implemented using shared memory. However, the user providing an IRMA method has to implement only the problem-related algorithm of the feature transformation.

Modular interfaces to support queries within the IRMA system are generated by the *web server* using the hypertext preprocessor PHP. Besides the query by example strategy, two basic mechanisms are provided for medical applications: relevance facts explain to the user why a certain picture has been presented as a query result, and relevance feedback allows the user to correct, adapt, or modify his query by adapting parameters of the algorithm (query refinement). Logging of interactions enables the repetition of complex query refinement processes and allows references to previous query steps. Boolean functions extend the logical linkage of results. Furthermore, the strict separation of information and its visualization makes it simple to add new visualization modules or to change the layout of existing modules.

A standard web browser is used as the graphical user interface (GUI) for all IRMA *applications*. This has numerous advantages. IRMA applications run platform-independent on any computer connected to the Internet and hence, they are not restricted to the IRMA cluster. In addition, physicians are used to handle web interfaces and, therefore, they do not need special instruction to use IRMA applications.

3. Results

Although manual labeling of the IRMA reference database is still in progress, the system was already used for processing basic queries, i.e. queries regarding the category of images. Based on a subset of 1,617 images from six body regions (abdomen, limbs, chest, breast, skull, and spine) ac-

quired with various modalities in several orientations, a statistical classifier based on the tangent distance performed with as low as 8% error rate [27]. Table 1 shows results of other methods on this data, indicating the results to be outstanding considering the difficulty of the task.

In another study, 1,867 chest radiographs were separated automatically into frontal (posteroanterior) and lateral orientation by a correlation measure computed from substantially size-reduced icons of 8×8 pixels. Here, the error rate was below 1% [29].

Currently, the IRMA reference database holds 3,879 images that have been labeled according to the IRMA code (Table 2). Since the frequency of imaged body regions reflects the clinical situation, chest radiographs occur most often. In addition, the sub-region classes are distributed irregularly. Classes with less than ten examples were not considered for leaving-one-out experiments. A downscaled image representation of 16×16 pixels results in 256 features that were used for a 5-nearest-neighbor correlation-based classifier. Regardless of modality, orientation and biological system, automatic classification of the anatomic region is possible with error rates of 11.6% and 15.8% using leaving-one-out for eight body regions and for 26 relevant sub-regions, respectively. Taking into account the large number of categories and the high variability within the categories due to the various imaging modalities, the results obtained are excellent. Also note that the IRMA concept pursues all likely categories for later processing steps: When the three or five most likely sub-regions are considered for each sample, the error rate for 26 sub-regions drops to 8.4% and 5.6%, respectively. For this experiment, the 15 and 25 nearest neighbors were analyzed. Combining the correlation-based classifier with a statistical classifier can further decrease the error rates. If each classifier contributes the 8 and 13 nearest neighbors, error rates yield 6.9% and 4.8%, respectively. Note that all leaving-one-out experiments were performed automatically and controlled by the IRMA scheduler.

Preliminary results were also obtained for the web-based interfaces. The IRMA

Method	Error Rate
co-occurrence matrices	29.0 %
squared images, 1-nearest-neighbor in Euclidian distance	18.1 %
local representation [28]	9.6 %
statistical approach	8.0 %

Table 1
Error rates of different methods on the subset of 1,617 images from the IRMA reference database

code editor is used for manual labeling of reference images (Fig. 4a). The IRMA code can be edited either directly by typing in the code or by selecting the entries from selection boxes. Based on prior selections, the sub-codes offered are adopted properly. All changes are recorded in a history protocol stored in the central database for easy error recovery (Fig. 4b).

4. Discussion

In addition to the requirements originating from content-based retrieval in general [11], the design of a medical image retrieval system requires attention to several other aspects and domain-specific properties [10]. In particular, the IRMA system supports modular design of arbitrary retrieval algorithms. Modularity easily enables the verification of isolated processing steps and allows the re-use of programs for various experiments and applications. All kinds of features (global, local, blob trees) are uniformly accessible, which results from the general and flexible feature model provided by IRMA. This includes a straightforward definition of a set of images (more general: a set of features). The system supports the automatic transfer of new and updated processing components into the pool of retrieval algorithms available to the physicians. In addition, new algorithms can quickly access the image database shortening the cycles between development and testing.

All technical details of the IRMA system are as transparent as possible for all participants. There is no need for the programmer to take care of platform-related communications regarding the query context, feature access, or feature storage. In particular, the storage location is transparent. Furthermore, the programmer is disburdened from data-flow synchronization, i.e. concurrency transparency of processing steps. For the physician, the execution of a query does not require additional technical knowledge about the underlying implementation of the retrieval process, i.e. all steps of the process run fully automatically without further user interaction. In sum-

mary, IRMA provides (i) location and access transparency for data and program sources; (ii) replication transparency for programs in development; (iii) concurrency transparency for job processing and feature extraction; (iv) system transparency at

method implementation time; and (v) job distribution transparency when issuing a query.

According to Tagare et al. [10], the changeability of a medical image database is the single most important aspect of build-

Table 2
Current distribution of image categories within the IRMA reference database

Body Region and Code Sub-region and Code	Images and Frequency	Images and Frequency
0 unspecified	0 0 %	
1 whole body	0 0 %	
10 unspecified		0 0.00 %
11 torso		0 0.00 %
12 extremities		0 0.00 %
2 cranium	179 4.61 %	
20 unspecified		(4) 0.00 %
21 facial cranium		77 1.99 %
22 cranial base		18 0.47 %
23 neuro cranium		80 2.07 %
3 spine	262 6.75 %	
30 unspecified		0 0.00 %
31 cervical spine		106 2.74 %
32 thoracic spine		51 1.32 %
33 lumbar spine		105 2.72 %
34 sacral bone		0 0.00 %
35 coccygeal bone		0 0.00 %
4 upper extremity / arm	678 17.48 %	
40 unspecified		(2) 0.00 %
41 hand		458 11.85 %
42 radio carpal joint		61 1.58 %
43 forearm		23 0.60 %
44 elbow		44 1.14 %
45 upper arm		28 0.72 %
46 shoulder		62 1.60 %
5 chest	1889 48.70 %	
50 unspecified		1873 48.47 %
51 bones		15 0.39 %
52 lung		0 0.00 %
53 hilum		0 0.00 %
54 mediastinum		0 0.00 %
55 heart		0 0.00 %
56 diaphragm		0 0.00 %
6 breast (mamma)	144 3.71 %	144 3.73 %
7 abdomen	169 4.36 %	
70 unspecified		32 0.83 %
71 upper abdomen		40 1.04 %
72 middle abdomen		47 1.22 %
73 lower abdomen		50 1.29 %
8 pelvis	56 1.44 %	
80 unspecified		51 1.32 %
81 sacral bone		(2) 0.00 %
82 iliac bone		(2) 0.00 %
83 pubic bone		(1) 0.00 %
84 small pelvis		0 0.00 %
9 lower extremity / leg	502 12.94 %	
90 unspecified		(4) 0.00 %
91 foot		66 1.71 %
92 ankle joint		62 1.60 %
93 lower leg		26 0.67 %
94 knee		254 6.57 %
95 upper leg		63 1.63 %
96 hip		27 0.70 %
Sum	3,879 100 %	3,864 100 %

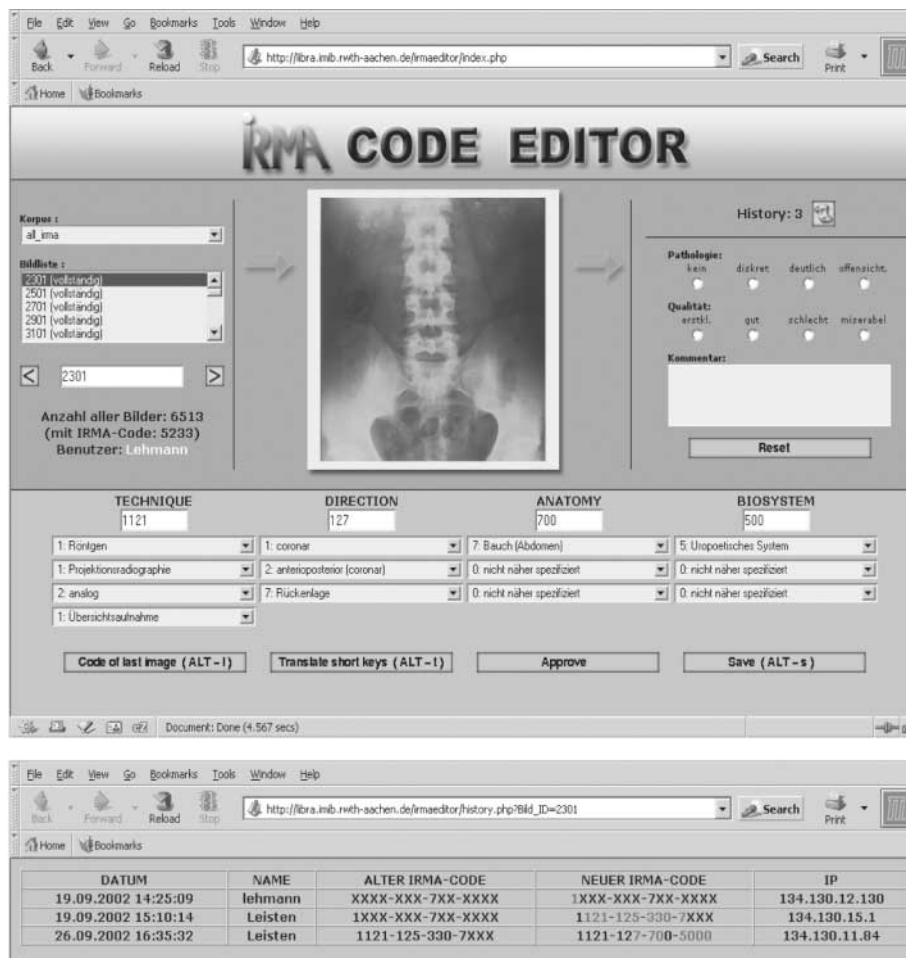


Fig. 4 Web-based interface for manual reference categorization (top) and logging of code modifications within the central database (bottom)

content-based image retrieval systems in medical applications. Based on the transparent mechanisms for distributed computing, data replication, and program installation, IRMA relies on a generic scheme for image content abstraction that is based on a central database. It enables the image and feature database to considerably evolve over the lifetime of the system. Furthermore, image features used in IRMA are computed automatically and are not influenced by knowledge bias arising from a gestalt-driven diagnostic interpretive process. Therefore, IRMA database activity is decoupled from interpretation activity, which meets the second most strongest request of Tagare et al. [10].

5. Conclusion

In contrast to specific applications, a general CBIR approach for medical images is presented and implemented combining a central database with a distributed system architecture suitable for large image databases such as within a picture archiving and communication system. The IRMA system supports rapid prototyping and quick integration of novel image analysis methods. So far, the IRMA system is used to answer primitive queries on the image category level. However, these experiments already have proven the validity and applicability of the IRMA concept. Resulting from system transparency, IRMA is suitable for sophisticated image processing to be performed by a medical user. In other words, CBIR principles are made available for a

variety of medical applications. Hence, IRMA narrows the gap between the semantic imprint of an image and any alpha-numerical description that is always incomplete.

Acknowledgment

The IRMA project has been funded by the German Research Community (DFG, grants Le 1108/4 and Le 1108/6). We are grateful to Tim Dwyer, Department of Computer Science, University of Melbourne, Australia, for providing the WilmaScope 3D graph visualization system (<http://www.wilmascope.org>) used to display the IRMA blob trees. Further information on IRMA is available at <http://irma-project.org>.

References

1. Niblack W, Barber R, Equitz W, Flickner M, Yanker P, Ashley J. The QBIC-project – Querying images by content using color, texture and shape. *Proceedings SPIE* 1993; 1908: 173-87.
2. Wu KJ, Narasimhalu AD, Mehtre BM, Lam CP, Gao YJ. CORE – A content-based retrieval engine for multimedia information systems. *Multimedia Systems* 1995; 3: 25-41.
3. Carson C, Belongie S, Greenspan H, Malik J. Blobworld – Image segmentation using expectation-maximization and its application to image querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2002; 24 (8): 1026-38.
4. Nappi M, Polese G, Tortora G. FIRST – Fractal indexing and retrieval system for image databases. *Image and Vision Computing* 1998; 16: 1019-31.
5. Pentland A, Picard RW, Scarloff S. Photobook – Content-based manipulation of image databases. *Proceedings SPIE* 1994; 2185: 34-47.
6. Korn P, Sidiropoulos N, Faloutsos C, Siegel E, Protopapas Z. Fast and effective retrieval of medical tumor shapes. *IEEE Transactions on Knowledge and Data Engineering* 1998; 10 (6): 889-904.
7. Shyu CR, Brodley CE, Kak AC, Kosaka A, Aisen AM, Broderick LS. ASSERT – A physician-in-the-loop content-based retrieval system for HRCT image databases. *Computer Vision and Image Understanding* 1999; 75 (1/2): 111-32.
8. Long LR, Antania S, Leeb DJ, Krainack DM, Thoma GR. Biomedical information from a national collection of spine x-rays – Film to content-based retrieval. *Proceedings SPIE* 2003; 5033: 70-84.
9. Chu WW, Hsu CC, Cárdenas AF, Tiara RK. Knowledge-based image retrieval with spatial and temporal constructs. *IEEE Transactions on Knowledge and Data Engineering* 1998; 10 (6): 872-88.
10. Tagare HD, Jaffe CC, Dungan J. Medical image databases: A content-based retrieval approach. *Journal of the American Medical Informatics Association* 1997; 4: 184-98.

11. Smeulders AWM, Worring M, Santint S, Gupta A, Jain R. Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2000; 22 (12): 1349-80.
12. Lehmann TM, Wein B, Dahmen J, Bredno J, Vogelsang F, Kohnen M. Content-based image retrieval in medical applications: A novel multi-step approach. *Proceedings SPIE* 2000; 3972: 312-20.
13. Lehmann TM, Schubert H, Keysers D, Kohnen M, Wein BB. The IRMA code for unique classification of medical images. *Proceedings SPIE* 2003; 5033: 440-51.
14. Belongie S, Carson C, Greenspan H, Malik J. Color- and texture-based image segmentation using EM and its application to content-based image retrieval. *Proceedings 6th International Conference on Computer Vision*, 1998; 675-82.
15. Ngo CW, Pong TC, Chin RT. Exploiting image indexing techniques in DCT domain. *IAPR International Workshop on Multimedia Information Analysis and Retrieval*, 1998; 195-206.
16. Wang JZ, Wiederhold G, Firschein O, Wei SX. Content-based image indexing and searching using Daubechies' wavelets. *International Journal of Digital Libraries* 1998; 1 (4): 311-28.
17. Castelli V, Kontoyiannis I, Robinson J, Bergman L, Li CS, Turek J. Progressive search and retrieval in large image archives. *IBM Journal of Research and Development* 1998; 42 (2): 253-68.
18. Chaudhuri BB, Sakar N. Texture segmentation using fractal dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1995; 17: 72-6.
19. Zhou XS, Huang TS. Edge-based structural features for content-based image retrieval. *Pattern Recognition Letters* 2001; 22: 457-68.
20. Simard P, Le Cun Y, Denker J. Efficient pattern recognition using a new transformation distance. In: Hanson S, Cowan J, Giles C (eds). *Advances in Neural Information Processing Systems 5*, Morgan Kaufmann, San Mateo, CA, 1993; 50-8.
21. Thies C, Malik A, Keysers D, Kohnen M, Fischer B, Lehmann TM. Hierarchical feature clustering for content-based retrieval in medical image databases. *Proceedings SPIE* 2003; 5032: 598-608.
22. Melnik S, Garcia-Molina H, Rahm E. Similarity flooding – A versatile graph matching algorithm and its application to schema matching. *Proceedings 18th ICDE Conference*, 2002; 117-28.
23. Depiero F, Trivedi M, Serbin S. Graph matching using a direct classification of node attendance. *Pattern Recognition* 1996; 29 (6): 1031-48.
24. Schädl K, Wysotzki F. Comparing structures using a Hopfield-style neural network. *Applied Intelligence* 1999; 11: 15-30.
25. Lee EA, Parks TM. Dataflow process networks. *Proceedings of the IEEE* 1995; 83(5): 773-99.
26. Güld MO, Thies C, Fischer B, Keysers D, Wein BB, Lehmann TM. A platform for distributed image processing and image retrieval. *Proceedings SPIE* 2003; 5150: 1109-20.
27. Keysers D, Dahmen J, Ney H, Wein BB, Lehmann TM. A statistical framework for model-based image retrieval in medical applications. *Journal of Electronic Imaging*, 2003; 12 (1): 59-68.
28. Paredes R, Keysers D, Lehmann TM, Wein BB, Ney H, Vidal E. Classification of medical images using local representations. *Proceedings BVM 2002*, Springer-Verlag, Berlin, 2002; 171-4.
29. Lehmann TM, Güld MO, Schubert H, Wenning A, Keysers D, Wein BB. Automatic detection of the view position of chest radiographs. *Proceedings SPIE* 2003; 5032: 1275-82.

Correspondence to:

Thomas M. Lehmann
 RWTH Aachen
 Institut für Medizinische Informatik
 D-52027 Aachen, Germany
 E-mail: lehmann@computer.org