

Geographically Distributed Complementary Content-Based Image Retrieval Systems for Biomedical Image Informatics

Sameer K. Antani^a, Thomas M. Deserno^{a,b}, L. Rodney Long^a, George R. Thoma^a

^a National Library of Medicine, National Institutes of Health, Bethesda, MD, USA

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

Abstract

There is a significant increase in the use of medical images in clinical medicine, disease research, and education. While the literature lists several successful systems for content-based image retrieval and image management methods, they have been unable to make significant inroads in routine medical informatics. This can be attributed to the following: (i) the challenging nature of medical images, (ii) need for specialized methods specific to each image type and detail, (iii) lack of advances in image indexing methods, and (iv) lack of a uniform data and resource exchange framework between complementary systems. Most systems tend to focus on varying degrees of the first two items, making them very versatile in a small sampling of the variety of medical images but unable to share their strengths. This paper proposes to overcome these shortcomings by defining a data and resource exchange framework using open standards and software to develop geographically distributed toolkits. As proof-of-concept, we describe the coupling of two complementary geographically separated systems: the IRMA system at Aachen University of Technology in Germany, and the SPIRS system at the U. S. National Library of Medicine in the United States of America.

Keywords: *Medical informatics applications, image information storage and retrieval, Internet services*

Introduction

There has been an explosive growth in the acquisition of medical images for clinical diagnosis, and use in medical research and education [1]. Hospitals have been adopting technology such as Picture Archiving and Communication Systems (PACS) and Hospital Information Systems (HIS) to assist in the digital collection, organization, and storage of patient data. The goal of these systems is to make patient data more accessible; in reality, the amount of data that is entered and stored in these systems have created new challenges in effective information indexing and retrieval. Retrieval of image information from these systems is done using limited text keywords in special fields (e.g., unique patient identifier, fields in the image header). These keywords, however, do not

capture the richness of features depicted in the image itself. It would be beneficial if the images could be retrieved by their visual content to help improve research, education, or medical practice. Content-Based Image Retrieval (CBIR) has received significant attention in the literature as a promising technique to ease the management of large image collections in a variety of domains [2-4]. Recently there has been an increasing interest in applying it to medical image repositories [5]. Rather than limiting queries to textual keywords, users can also query by example image or image feature (e.g., color, texture, or shape computed from a region of interest) to find similar images of the same modality, anatomical region, and disease along with the matching associated text records.

While the literature lists several successful systems for content-based image retrieval and image management methods [2, 5], they have not made significant inroads in routine medical informatics. In addition, although many large imaging databases exist, such as the National Cancer Imaging Archive (NCIA) or the Lung Imaging Database Consortium (LIDC) created under the aegis of the Cancer Imaging Program¹ at the U.S. National Cancer Institute (NCI), these efforts have concentrated on data collection and transmission but have left development of applications to the research community. Lack of CBIR adoption has been attributed partly to the difficulty of integrating current implementations with existing healthcare systems [6]. The following reasons may further explain this anomaly: (i) challenging nature of medical images, (ii) need for specialized methods specific to each image type and detail, (iii) lack of effective image indexing methods, and (iv) lack of a uniform data and resource exchange framework between complementary systems. Most systems tend to focus on varying degrees of first two items, making them very versatile in a small sampling of the variety of medical images, but unable to share their resources or strengths. This requires each project to redevelop what may exist as an advanced implementation, but inaccessible. The lack of suitable image indexing methods is a problem for large image collections. Image comparisons performed linearly are inefficient and too slow for practical use. This paper proposes to overcome these shortcomings by defining a data and resource exchange framework using open

¹ <http://imaging.cancer.gov> (Last accessed: March 27, 2007)

standards and software to enable such specialized systems to act as geographically distributed toolkits. The approach enables communication between two or more geographically separated complementary systems with possibly different architectures and developed on different platforms, and specialized for different image modalities and characteristics. The resulting system provides the user with a rich functionality

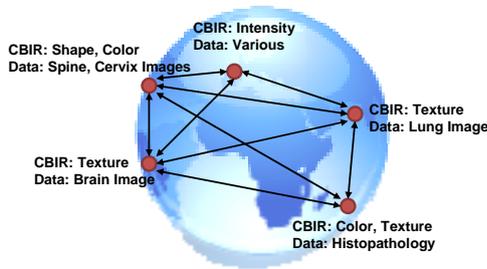


Figure 1 – Concept: Geographically distributed framework of complementary CBIR systems.

operating within a familiar interface on the Web-browser, making it portable and independent of location and underlying user operating systems. Figure 1 illustrates this concept.

In this figure, each circle represents a CBIR system specializing in particular image types, pathologies, or CBIR techniques. Using a standard open protocol any system could act as a client using the services available elsewhere to provide the user with a rich medical image informatics resource.

As proof-of-concept, we describe the coupling of two leading, complementary geographically separated image informatics CBIR systems: The Image Retrieval in Medical Applications (IRMA) project at the Aachen University of Technology (RWTH) [7, 8] and the Spine Pathology & Image Retrieval System (SPIRS) project at the U.S. National Library of Medicine (NLM) [9-11]. IRMA and SPIRS retrieve images based on different approaches: Traditionally, IRMA has focused on image retrieval by computing overall (or global) image characteristics such as color, intensity, and texture. In particular, the image distortion model was developed and proven as a robust distance measure for differences on smaller sub-regions in the image [12]. Such an approach permits queries on a varied image collection and helps identify similar images, e.g., all chest X-rays in the A-P view. The IRMA system lacks the ability to find particular pathology that may be localized in specific regions within the image. In contrast, SPIRS can retrieve images that exhibit pathology that may be localized to a particular region under the assumption that the query to a large image collection containing images of only one type, e.g., vertebral pathology expressed in spine x-ray images in the sagittal plane. SPIRS lacks the ability to select pertinent images from a large varied image collection typical in a hospital PACS system, for example. We believe that combining the strengths of these two complementary technologies of whole image and local feature-based retrieval is unique and valuable for research into the retrieval of images in large repositories that are similar in type as well as pathology.

Background

IRMA Project

The IRMA project² aims to develop and implement high-level methods for CBIR with prototypical application to medico-diagnostic tasks on radiological image archives. Stated goals include support for semantic and formalized queries to the medical image database with support for inter- and intra-individual variance and diseases. Example tasks are the staging of a patient's therapy or the retrieval of images with similar diagnostic findings in large electronic archives. Formal content-based queries also take into account the technical conditions of the examination and the image acquisition modalities. The system classifies radiological images in a general way without restriction to a certain diagnostic problem or question. Pattern recognition and structural analysis methods describe the image content in a feature-based, formal and generalized way. The *mean* image description enables a fast and reliable image comparison and retrieval. The project also includes an automatic classification and indexing process for insertion of new data into the system without manual interaction.

The IRMA project has several interfaces and can be characterized by its features: (i) Automated classification of radiographs based on global features with respect to imaging modality, direction, body region examined and biological system under investigation; (ii) Identification of image features that are relevant for medical diagnosis; these features are derived from a priori classified and registered images; and (iii) Image retrieval on similarity to an a priori selected feature set based on the visual similarity of certain image structures. Current image data consists of radiographs, with future plans to include medical images from other modalities. An IRMA retrieval interface supporting query refinement is shown in Figure 2.



Figure 2- IRMA retrieval interface with query refinement.

² <http://irma-project.org> (Last accessed: March 27, 2007)

SPIRS Project

SPIRS³ provides a Web-based interface for performing image retrieval from a database of digitized spine x-rays using the morphological shape of the vertebral body. Its framework enables interaction with and retrieval of relevant information from large databases of image and patient data using rich hybrid image and text query methods.

A query editor enables users to pose queries by sketching a unique shape, or selecting or modifying an existing shape from the database. It aims to capture query semantics through support of advanced mechanisms like multiple partial shape matching. Additional text fields enable users to supplement visual queries with other relevant data (e.g., anthropometric data, quantitative imaging parameters, patient demographics). These hybrid text-image queries may be annotated with pertinent pathologies by selecting and weighting local features to indicate importance. Query results appear in a customizable window that displays the top matching results and related patient data. SPIRS provides a working proof-of-concept that is capable of accommodating large amounts of imaging data expected in the near future.



Figure 3 – SPIRS interface with (a) crop and (b) detail view

At NLM, the focus of CBIR research has been to develop systems capable of performing a range of queries on large medical multimedia databases comprising various biomedical images and patient health data. Such a database in current use contains digitized spine x-rays and associated person metadata that comes from a large nationwide survey, the Second National Health and Nutrition Examination Survey (NHANES II). NHANES is conducted regularly by the National Center for Health Statistics in the United States. The goals of NHANES include estimating prevalence of selected diseases, monitoring disease trends, and studying the relationship between nutrition and health. The NHANES II collection is considered very valuable to radiologists, bone morphometrists, researchers in osteo-arthritis, and medical educators. Domain experts reviewed a sample of the data and identified 23 key biomedical features that were exhibited in the x-rays. Of these, anterior osteophytes, spondylolisthesis, and disc space narrowing were determined to be frequently occurring and reliably detectable. Each of these key features may be identified by examining the boundary shape of the vertebra. SPIRS may be used to determine what features (e.g., protrusion

on the anterior edge of the cervical vertebra) are consistently associated with a certain symptom (e.g., neck discomfort) or whether a certain feature is a precursor to more serious illnesses (e.g., arthritis). Pathology or medical condition may also be documented in the text of the patient record as a survey response or in the medical diagnoses. The SPIRS Web-interface is shown in Figure 3. In addition to the Web-interface, SPIRS also offers a service to its core shape similarity algorithms and data. Although SPIRS focuses on shape-based queries its framework is extensible to adopt features particular to other biomedical image and data collections, e.g., its core architecture is being extended to include color, texture, and spatial location in uterine cervix images from the National Cancer Institute at the NIH [11]

SPIRS-IRMA Multilevel Distributed CBIR

Medical CBIR systems can be classified by the type of numerical features used to characterize the images: (i) **global** approaches extract a single feature vector from the entire image, (ii) **local** approaches assign the feature vector to distinct regions of interest, and (iii) **structural** approaches additionally cope with spatial and temporal relations between the image objects of interest. In this context the IRMA system is a global feature system with particular structural aspects of some image regions that are also computed. In contrast, SPIRS takes a more local approach to image similarity with some structural knowledge available from the vertebral image labels. In this sense, SPIRS-IRMA jointly may be considered a multilevel distributed CBIR system [11, 13].

In this joint system, IRMA serves as the front-end for the end user, while SPIRS provides specific shape similarity algorithms and supplies formatted responses to structured queries. In this proof-of-concept system, extensive data interchange was minimized through data mirroring. It is conceivable, however, that the data could also be securely shared in a purely service oriented setup. Screenshots from the combined SPIRS-IRMA system are shown in Figure 4.

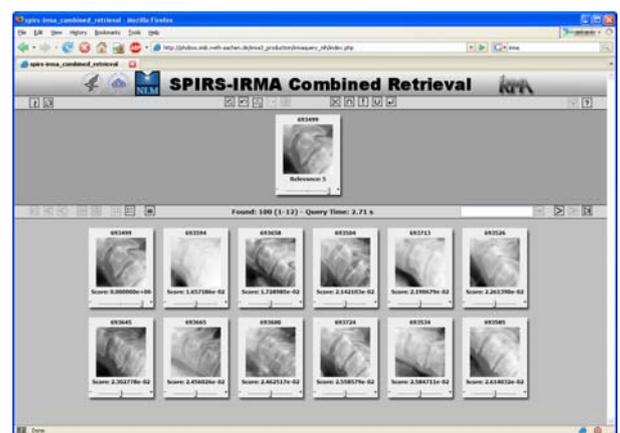


Figure 4 – SPIRS-IRMA interface

The system has plans for development in several phases. SPIRS-IRMA is utilizing only limited services provided by SPIRS with near term goals to use all available features. Future phases will also include permit users to upload their own images, as in the current IRMA system, and segmented or

³ <http://archive.nlm.nih.gov/spirs> (Last accessed: March 27, 2007)

sketched boundary outlines of interest. Next steps include multi-resolution shape queries and support for query logging and relevance feedback.

Data and Resource Exchange Framework

CBIR systems sharing their data and computational resources need to be developed with a distributed architecture as shown in Figure 4. Each CBIR system would require the following components: (i) a gateway that acts as a mediator between client and server-side components, (ii) the indexing and retrieval server, which performs the feature representation and

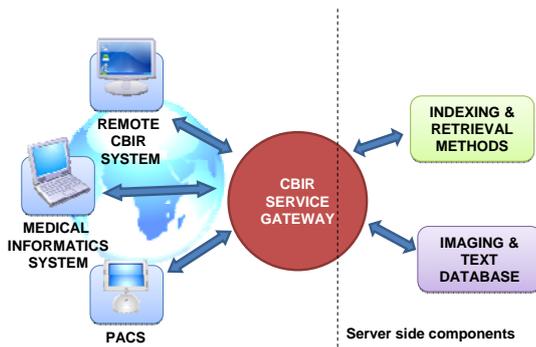


Figure 4 - The distributed architecture of SPIRS.

similarity matching, (iii) the databases containing images and associated text data, and (iv) an open communication standard for sharing of data and commands between the systems. Key components are described using the protocol developed for the SPIRS-IRMA system.

Gateway

The gateway is the entry point for all service requests. It can be implemented in any standard Web interface. In case of the SPIRS-IRMA collaboration, SPIRS has the service gateway implemented as a Java servlet. The gateway acts as a mediator between client requests and server-side components. It manages multiple simultaneous connections (users) as separate sessions and queues requests to the core CBIR engine. It translates query components that require information from the MySQL text database into SQL queries. The gateway is also responsible for formatting responses from server-side components and sending them to the client.

Server-side Components

The core CBIR algorithms that operate on image feature data and text databases comprise the server-side components. While these algorithms rarely need any modification, their implementation is of particular importance. Typically, monolithic CBIR systems have indexing and similarity retrieval components closely wedded to the query and result visualization user interfaces. It is necessary to decouple these core algorithms from their user interfaces and implement them such that the local interfaces also use the gateway for all system communication. SPIRS indexing and retrieval algorithms use a variety of shape similarity metrics embedded within feature indexing techniques for efficient retrieval essential for Web use [11]. All text data is stored in a MySQL database and is

accessible via both the gateway and the CBIR algorithms. In this initial prototype both systems hold all images and shape contours to minimize data exchange. Future plans include allowance for secure image and feature data exchange.

Communication Protocol and Data Exchange Format

A distributed computational framework relies extensively on a robust communications protocol and data exchange format. In case of SPIRS-IRMA, the systems are loosely-coupled over the Internet making it possible to use open communication standards. XML documents styled to a developed DTD are used for data exchange. The entire transaction is divided into three primary events, viz., *querystatus*, *query*, and *queryresult*. Each element in the XML file is designed for a particular event. For example, the *<querystatus>* element is used to determine if a desired service is available and to obtain a list of currently available services.

The *<query>* element is used to make shape queries where each query contains: (i) the query vertebra contour (boundary data points), (ii) partial shape indices and weights (if any) (iii) requested matching method (iv) maximum number of responses requested, and (v) range of similarity result scores. SPIRS responds with the *<queryresult>* element populated with: (i) matching image - vertebra tuples, and (ii) similarity scores. The design of the SPIRS-IRMA collaboration, the DTD, and sample XML data, are discussed in [13].

Discussion

The field of medical informatics has been unable to take advantage of image retrieval methods and systems in spite of their acknowledged importance in the face of growing use of image data in medical practice, research, and education. Image management and pathologically sensitive image content-based retrieval systems are thus increasingly necessary. The challenging nature of images and lack of comprehensive systems have hindered their acceptance. While it is difficult to develop a single comprehensive system, it may be possible to take advantage of the growing research interest and several successful systems with techniques developed for specific image collections. In addition to supporting rich segmentation, validation, indexing, query, retrieval, and visualization methods, inclusion of open interfaces using standard communication protocol, such as that described above, to the projects can enrich individual systems. The proposed approach has the following advantages:

- **Simplicity:** The simple communication interface allows for rapid development of methods for individual systems by expanding the available resources reachable through open communication standards.
- **Extensibility and Flexibility:** By separating the user interface from the core image informatics algorithms, the proposed approach allows systems offering services to continue development of other techniques and add them as they mature. It is no longer necessary to rebuild entire applications. In unusual cases,

the protocol allows removal of some services advertising their unavailability.

- Security: Separating the core algorithms enables selective additional data security, user authentication, and encryption of the communication component (the gateway), where appropriate. The flexibility also allows sharing of just the methods.

We have demonstrated some of these features through the SPIRS-IRMA collaboration. In the IRMA system, it is possible for a user to find images from their database similar to an uploaded image. The query is limited to this extent, however. It requires manual viewing of each resulting image to identify those with pathology similar to that in the query example. While the SPIRS system allows shape queries with multiple parts highlighted, which indicates their relative importance, and includes text fields for further refined responses, it is limited to the spine x-ray database. When completed, users familiar with the IRMA interface can extend their searches beyond finding similar x-rays to include localized searches. Goals include enabling image segmentation services for user-supplied images. The resulting system is a multi-level, pathologically sensitive, geographically distributed system.

Conclusions and Future Work

In addition to formalizing the implementation of all aspects of the SPIRS-IRMA system, we plan to develop and publish a formal specification to enable sharing of image informatics resources among various systems. As a test, we will expand this framework to the CBIR system under development for uterine cervix images that uses color, texture, and spatial location information in generating an image description.

In this article, we have proposed a resource sharing strategy for increasing the impact of traditionally limited medical CBIR systems on medical informatics, and possible applications to clinical medicine, medical research, and education. We propose the use of open standards and a distributed framework that is simple to implement, extensible, flexible, and secure for the development of CBIR systems. We demonstrate its impact through a prototype SPIRS-IRMA combined retrieval system.

Acknowledgements

The authors would like to thank Leif Neve at NLM, William Hsu at University of California at Los Angeles, and Mark Oliver Güld at Aachen University of Technology for their assistance.

This research was supported by the Intramural Research Program of the National Institutes of Health (NIH), National Library of Medicine (NLM), and Lister Hill National Center for Biomedical Communications (LHNCBC).

References

[1] Andriole KP. Addressing the Coming Radiology Crisis: The Society for Computer Applications in Radiology

Transforming the Radiological Interpretation Process (TRIP) Initiative. White Paper (<http://www.siiimweb.org>). November 2005. Last Accessed: March 26, 2007.

- [2] Smeulders A, Worring M, Santini S, Gupta A, Jain R. Content-based image retrieval at the end of the early years. *IEEE Trans Pattern Rec. & Machine Intel.* 2000; 22(12):1349-80.
- [3] Antani S, Kasturi R, Jain R. A survey on the use of pattern recognition methods for abstraction, indexing and retrieval of images and video. *Pattern Rec.* 2002;35(4):945-65.
- [4] Catarci T, Costabile MF, Levialdi S, Batini C. Visual query systems for databases: a survey. *J. Visual Lang. & Comp.* 1997;8(2):215-60.
- [5] Muller H, Michoux N, Bandon D, Geissbuhler A. A review of content-based image retrieval systems in medical applications—clinical benefits and future directions. *Int. J. Med. Info.* 2004; 73(1):1-23.
- [6] Shyu CR, Brodley CE, Kak AC, Kosaka A, Aisen AM, Broderick LS. ASSERT: A physician-in-the-loop content-based retrieval system for HCRT image databases. *Computer Vision & Image Understanding* 1999;75(12):111-32.
- [7] Lehmann TM, Güld MO, Thies C et al. Content-based image retrieval in medical applications. *Methods of Information in Medicine* 2004; 43(4): 354-61
- [8] Thies C, Gueld MO, Fischer B, Lehmann TM. Content-based queries on the CasImage database within the IRMA framework. *Lec. Notes Computer Sci* 2005;3491:781-92.
- [9] Antani S, Long LR, Thoma G. Content-based image retrieval for large biomedical image archives. *Proc. 11th World Cong. on Med Info (MEDINFO) 2004*, pp. 829-33.
- [10] Long LR, Antani SK, Thoma GR. Image informatics at a national research center. *Computerized Med Imaging & Graphics.* 2005;29:171-93.
- [11] Thoma GR, Long LR, Antani S. Biomedical imaging research and development: knowledge from images in the medical enterprise. Lister Hill National Center for Biomedical Communications, NLM, NIH; Sep. 2006. Report No.: LHNCBC-TR-2006-002.
- [12] Keysers D, Dahmen J, Ney H, Wein BB, Lehmann TM. A statistical framework for model-based image retrieval in medical applications. *J Electronic Imaging* 2003; 12(1): 59-68
- [13] Antani S, Lehmann TM, Long LR, Güld MO, Neve L, Thoma G. Interfacing Global and Local CBIR Systems for Medical Image Retrieval. Accepted by Workshop on Image Research for Medicine (Bildverarbeitung für die Medizin). March 2007, Munich, Germany.

Address for correspondence

Sameer Antani, PhD
National Library of Medicine, NIH
8600 Rockville Pike, MS 3824
Bethesda, MD 20894, USA
Email: santani@mail.nih.gov