Interfacing Global and Local CBIR Systems for Medical Image Retrieval

Sameer K. Antani[1], Thomas M. Deserno[1,2], L. Rodney Long[1],
Mark O. Güld[2], Leif Neve[1], George R. Thoma[1]

[1]U.S. National Library of Medicine, National Institutes of Health, Bethesda, MD, USA
[2]Dept of Medical Informatics, Aachen University of Technology (RWTH), Aachen, Germany
Email: santani@mail.nih.gov

Abstract. Contemporary picture archiving and communication systems are limited in managing large and varied image collections, because content-based image retrieval (CBIR) methods are unavailable. In this paper, an XML-based data and resource exchange framework is defined using open standards and software to enable specialized CBIR systems to act as geographically distributed toolkits. The approach enables communication and collaboration 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 synergy provides the user with a rich functionality operating within a familiar Web browser interface, making the combined system portable and independent of location and underlying user operating systems. We describe the coupling of the Image Retrieval in Medical Applications (IRMA) system and the Spine Pathology and Image Retrieval System (SPIRS) as proof of this concept.

1 Introduction

There has been an explosive growth in the acquisition and use of images in clinical medicine, medical research, and education [1]. In current picture archiving and communication systems (PACS), retrieval of image information is done using limited text keywords in special fields in the image header (e.g., patient identifier). Since these keywords do not capture the richness of features depicted in the image itself, content-based image retrieval (CBIR) has received significant attention in the literature as a promising technique to facilitate improved image management in PACS systems [2,3]. With this approach, rather than limiting queries to textual keywords, users can also provide an 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. In spite of this research interest, the challenging nature and variety of medical image data have contributed to the absence of CBIR in contemporary PACS. CBIR requires specialized methods specific to each image type and content detail. Some systems tend to focus on particular image types, while others that are less specific with respect to particular anatomy tend to concentrate more on image discrimination by overall appearance, and any pathological similarity is only in the gross overall view.

As proof of concept, we describe the collaboration between two leading, complementary geographically distributed CBIR systems. The Image Retrieval for Medical Applications (IRMA) Project[1] undertaken at the Aachen University of Technology (RWTH) [3,4] aims to provide visually rich image management through CBIR techniques applied to medical images using intensity distribution and texture measures taken globally over the entire image. This approach permits queries on a heterogeneous image collection and helps identify images that are similar with respect to global features, e.g., all chest x-rays in the AP (anterior-posterior) view. The IRMA system lacks the ability for finding particular pathology that may be localized in particular regions within the image. In contrast, the Spine Pathology and Image Retrieval System (SPIRS)[2] [5,6,7] at the U.S. National Library of Medicine provides localized vertebral shape-based CBIR methods for pathologically sensitive retrieval of digitized spine x-rays and associated person metadata that come from the Second U.S. National Health and Nutrition Examination Survey (NHANES II). In the SPIRS system, the images in the collection must be homogeneous, i.e., a single type imaging the same anatomy in the same view, e.g., vertebral pathology expressed in spine x-ray images in the sagittal plane. Combining the strengths of these two complementary technologies of whole image and local feature-based retrieval is unique and valuable to find images that are not only similar in overall appearance but also with locally expressed pathology. This is accomplished through use of an XML-based service protocol to access particular internal methods in these systems.

2 System Background

The IRMA Project: IRMA has developed and implemented high-level CBIR methods with application to medico-diagnostic tasks on radiological image archives. Current image data consists of radiographs, with future plans to include medical images from arbitrary modalities. IRMA is a system where methods are treated as black boxes with an input and an output connected as a directed graph [8]. Using this paradigm, new algorithms can be integrated into a Web-based query interface with relative ease. IRMA caches intermediate results to improve efficiency and support its query logging capability. Distributed processing is also enabled by assigning the set of methods particular to a query or user interaction to one of the daemon services running on network-connected computers.

SPIRS: SPIRS provides a Web-based interface, implemented as a Java applet, for performing image retrieval on a database of digitized spine x-rays using the morphological shape of the vertebral body. Its framework of shape indexing and retrieval algorithms communicate with external users through a Java servlet. SPIRS enables CBIR for the large databases of image and patient data using rich hybrid image and text query methods. A shape query editor enables sketching or selecting and/or modifying an existing shape in the database. It also supports advanced mechanisms like multiple partial shape queries. Additionally, text fields enable users to supplement visual queries with other relevant data. A customizable window displays the top matching vertebrae and

[1] `http://irma-project.org`
[2] `http://archive.nlm.nih.gov/spirs`

related text data. SPIRS also offers a service of its core shape similarity algorithms and data with a predefined DTD. It is being extended and generalized to include color, texture, and spatial location in uterine cervix images from the National Cancer Institute [7].

3 Roadmap

The interfacing of the SPIRS and IRMA systems has been divided into the following phases, each phase adding enhancements and features to the combined system. This paper covers Phases 1 and 2 of this development.

Phase 1 - Whole shape matching: In this phase, entire vertebral shapes are used in computing similarity. Additionally, the query shapes are selected from those within the database of pre-segmented vertebrae. Query shape modifications or user sketches are not supported.

Phase 2 - Similarity method selection: During this phase, SPIRS enables support for additional shape similarity algorithms. The user can select a shape similarity method at query time. A list of available methods and server status are provided by SPIRS.

Phase 3 - Partial shape matching: In this phase, the IRMA user will be able to specify multiple non-overlapping partial shapes and provide weights to indicate their relative importance.

Phase 4 - Shape matching by user-sketch: In this phase, the IRMA user will be allowed to alter pre-existing database shapes or sketch new shapes. Additionally, the user may specify a threshold for shape similarity scores to limit the number of results.

Phase 5 - Combined retrieval: At this final phase the IRMA user may securely upload his own spine x-ray image to find similar x-ray images using its global similarity algorithms. Shapes selected within a matching image may then be passed on to SPIRS for local similarity searches.

4 Method

A loosely-coupled Internet-based distributed computing framework, such as the SPIRS-IRMA interface relies extensively on a robust communications protocol and an open standard data exchange format. The elements in the XML file, styled to a predefined DTD, are designed toward particular events, e.g., 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 by IRMA to make shape queries, and the `<queryresult>` element populated by SPIRS provides the results. A fragment of the DTD is shown in Figure 1. The XML data exchange is implemented on top of the standard hypertext transfer protocol (HTTP). This allows easy adoption into any Web server infrastructure without affecting the firewall settings. The SPIRS service is accessed through a pre-defined URL.

IRMA methods for interacting with SPIRS: As noted above, IRMA methods are treated as black boxes. This eases the addition of new SPIRS-IRMA coupling code to provide the communication interface with SPIRS and perform necessary conversion to and from the XML transport format. An additional required conversion is from the

```
<!ELEMENT spirs_irma (querystatus|algorithms|query|queryresult)>

<!ELEMENT algorithmidlist (algorithmid+)>
<!ELEMENT algorithmid (#PCDATA)>
<!ELEMENT query (algorithmid,contour)>
<!ELEMENT algorithmid (#PCDATA)>
<!ELEMENT contour (pointlist)>
<!ELEMENT pointlist (point+)>
<!ELEMENT point (x,y)>
<!ELEMENT x (#PCDATA)>
<!ELEMENT y (#PCDATA)>

 <!ELEMENT queryresult (neighborlist)>
<!ELEMENT neighborlist (neighbor*)>
<!ELEMENT neighbor (imagetag,vertebranum,similarityscore)>
<!ELEMENT imagetag (#PCDATA)>
<!ELEMENT vertebranum (#PCDATA)>
<!ELEMENT similarityscore (#PCDATA)>
```

Fig. 1. Fragment of the SPIRS-IRMA DTD showing query and result elements.

image-vertebra identifiers used in SPIRS to equivalent IRMA image IDs. The response to a query can be visualized using the existing Web-based IRMA interface.

SPIRS service for external queries: SPIRS implements the service gateway as a Java servlet which is the entry point for all service requests and acts as a mediator between client requests and server-side components. It manages multiple simultaneous connections as separate sessions and queues requests to the core SPIRS engine. It also translates query components that require information from the MySQL text database in SPIRS into SQL queries. Finally, the gateway is also responsible for formatting responses from server-side components and the MySQL database into the XML response format and sending them to the client.

Relevance feedback for SPIRS-IRMA: IRMA supports relevance feedback by allowing the user to specify the degree of relevance on results from an initial CBIR query. The set of "relevance specified" images can form a new IRMA query which, however, cannot be directly applied to SPIRS since it uses only a single shape as query. In the modified method: (i) the contour vectors for all elements included in the feedback are collected; (ii) each of these vectors is then used to individually query SPIRS, which breaks down the feedback query into a series of single queries; (iii) the query refinement is performed thereafter as a post-processing step on the set of these single-query results.

5 Results and Conclusion

This paper describes the SPIRS-IRMA collaborative CBIR system, which enables users to pose shape queries local to a particular image region through the interface of a CBIR system that inherently supports only overall image similarity. The communication is done over HTTP using XML tags that conform to a predefined "SPIRS-IRMA interface" DTD. The development of this system is planned in several phases. The first two phases which allow vertebral whole shape queries, with user selection of similarity methods, is currently completed, and work on subsequent phases is in progress. A screen-capture of the SPIRS-IRMA system, which is accessible from the IRMA Website, is shown in Figure 2.

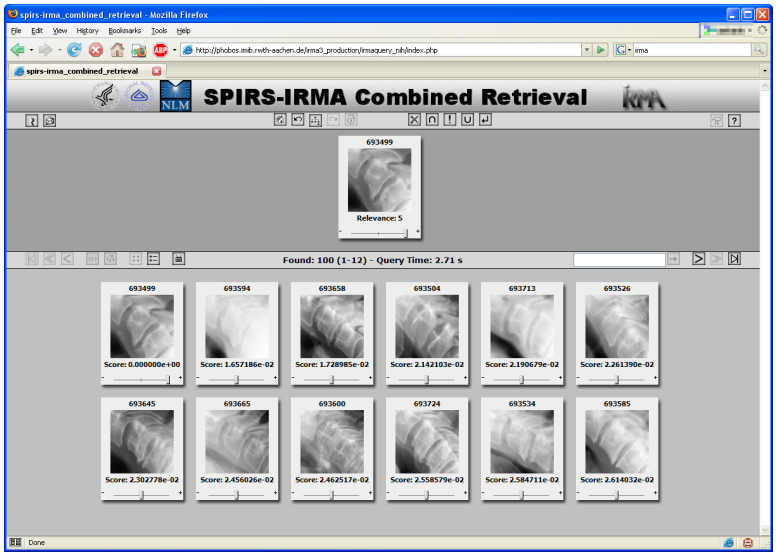

Fig. 2. The SPIRS-IRMA Web interface.

Acknowledgement

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

References

1. Andriole KP, Morin RL. Transforming Medical Imaging: The First SCAR TRIP™Conference. A Position Paper from the SCAR TRIP™Subcommittee of the SCAR Research and Development Committee. J Digital Imaging 2006;19(1):6–16.
2. Müller 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.
3. Lehmann TM, Güld MO, Thies C, Fischer B, Spitzer K, Keysers D, et al. Content-based image retrieval in medical applications. Methods of Info in Med 2004;43(4):354–61.
4. Thies C, Güld MO, Fischer B, Lehmann TM. Content-based queries on the CasImage database within the IRMA framework. Lec Notes in Comp Sci 2005;3491:781–92.
5. Antani S, Long LR, Thoma GR. Content-based image retrieval for large biomedical image archives. In: Proc 11^{th} World Cong Medical Informatics; 2004. p. 829–33.
6. Long LR, Antani SK, Thoma GR. Image informatics at a national research center. Comp Med Imaging & Graphics 2005;29:171–93.
7. Thoma GR, Long LR, Antani SK. Biomedical imaging research and development: knowledge from images in the medical enterprise. Technical Report Lister Hill National Ctr for Biomedical Communications, US National Library of Medicine, NIH 2006;LHNCBC-TR-2006-002.
8. Güld MO, Thies C, Fischer B, Lehmann TM. A generic concept for the implementation of medical image retrieval systems. Int J Med Info 2006;In Press.