SPIRS: A Framework for Content-based Image Retrieval from Large Biomedical Databases

William Hsu[a], L. Rodney Long[b], Sameer Antani[b]

[a] *Medical Imaging Informatics Group, University of California, Los Angeles, CA, USA*
[b] *National Library of Medicine, National Institutes of Health, Bethesda, MD, USA*

Abstract

With the increasing use of medical images in clinical medicine, disease research, and education, the need for methods that effectively archive, query, and retrieve these images by their content is underscored. This paper presents the implementation of a Web-based retrieval system called SPIRS (Spine Pathology & Image Retrieval System) at the U. S. National Library of Medicine that demonstrates recent developments in shape representation and retrieval from a large dataset of 17,000 digitized x-ray images of the spine and associated text records. Users can search these images by providing a sketch of the vertebral outline or selecting an example vertebral image and some relevant text parameters. Pertinent pathology on the image/sketch can be annotated and weighted to indicate importance. This hybrid text-image query yields images containing similar vertebrae along with relevant fields from associated text records, which allows users to examine the pathologies of vertebral abnormalities. Initial experiments with SPIRS have demonstrated the potential for this system, particularly on a large dataset of clinical images.

Keywords: Medical informatics applications, Information storage and retrieval, Content-based image retrieval, Visual access methods, Web-based systems

Introduction

Medical imaging is an increasingly important and versatile tool for acquiring information about a patient and disease. Image features (e.g., shape, color, or texture) have been used to diagnose a variety of conditions. 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 a new challenge for effective information indexing and retrieval. Historically, PACS have limited users to query by certain keywords (e.g., unique patient identifier, fields in the image header). However, these keywords often do not capture the richness of features depicted in the image itself. Over the past two decades, content-based image retrieval (CBIR) systems have been researched to address the problem of indexing and retrieving visual data in a variety of domains [1]. Rather than limiting queries to textual keywords, CBIR users provide a query sketch/image, which is then used to find similar images of the same modality, anatomical region, and disease along with the associated text records.

The goal of this work is to develop a retrieval system that implements recent developments in shape representation, efficient indexing, and similarity matching; supports whole and partial shape matching, which enables a wide range of meaningful queries to be posed; and utilizes a distributed framework, which is customizable to the needs and constraints of healthcare environments. The result, Spine Pathology & Image Retrieval System (SPIRS)[1], provides a Web-based interface for performing image retrieval on a database of digitized spine x-rays using the morphological shape of the vertebral body. A query editor enables users to pose queries by sketching a unique shape or selecting or modifying an existing shape from the database. 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 demonstrate the capability of accommodating large amounts of imaging data expected in the near future.

Background

The problem of retrieving information based on image content has been researched by various groups since the early 1990's resulting in the development of tools such as QBIC, Virage, and Blobworld [2]. CBIR in medicine has been an active area of research [3], but only a small number of proposed systems such as ASSERT [4] and IRMA [5] have been demonstrated in the clinical environment. In addition, although many large image databases exist, such as the National Cancer Imaging

[1] http://archive.nlm.nih.gov/spirs

Archive (NCIA) or the Lung Imaging Database Consortium (LIDC) created under the aegis of the Cancer Imaging Program[2] 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 is attributed partly to the difficulty of integrating current implementations with existing healthcare systems [3]. SPIRS addresses this issue by (i) utilizing open standards to communicate among components, which can be extended to support data encryption to meet privacy regulations, and (ii) using modular open source software components. Much work has been done in the past on visual querying paradigms, as reviewed in [6]. Building upon these paradigms, SPIRS combines visual and text queries to provide users with greater flexibility in retrieving relevant results.

Biomedical Database

At the U.S. National Library of Medicine (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 metadata from a large nationwide survey, the National Health and Nutrition Examination Survey (NHANES), 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 Lister Hill National Center for Biomedical Communications, a research division of the NLM, maintains data from the second survey, NHANES II, which was collected between 1976 and 1980 and featured over 20,000 participants. Each participant's record includes 2,000 textual data points such as health questionnaire answers, anthropometric information, and results from a physical exam. This textual data is stored in a relational database (e.g., MySQL). Supplementing the textual data is a collection of 17,000 cervical and lumbar spine x-rays that were taken from patients aged 25 – 74. These images were originally on film and subsequently digitized using a 146 dpi scanner resulting in 140 Gigabytes of data. The collection is considered valuable to radiologists, bone morphometrists and researchers interested in osteoarthritis, and medical educators. Domain experts reviewed a sample of the data and identified 23 key biomedical features exhibited in the x-rays. Of these, anterior osteophytes, subluxation/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. However, identifying images that exhibit particular types and severities of these pathologies is extremely tedious for such a large collection. SPIRS and related tools [7], developed to analyze the images and populate the database, automate many tasks and enable retrieval of the images using sketches of the vertebral boundary with the desired feature. Furthermore, 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). Illnesses may be documented in the text of the patient record as a survey response or in the medical diagnoses.

[2] http://imaging.cancer.gov/

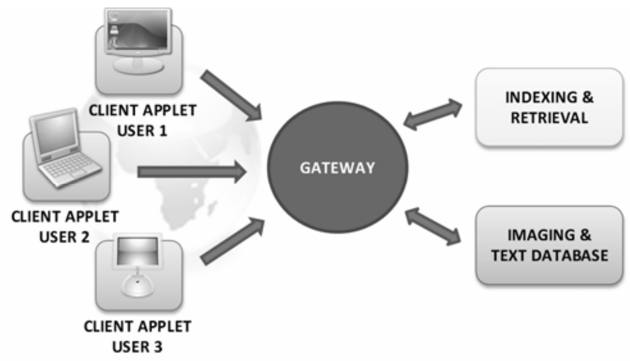

Figure 1 - The distributed architecture of SPIRS.

System Framework

The system's distributed architecture, shown in Figure 1, consists of four components: (i) the client applet, which provides a front-end for users to pose queries and interact with results, (ii) a gateway that acts as an mediator between client and server-side components, (iii) the indexing and retrieval server, which performs the feature representation and similarity matching, and (iv) the databases containing images and associated text data.

Indexing and Retrieval

In our continuing research, we have explored and implemented various algorithms to extract, index, and retrieve shapes of vertebral bodies from x-ray images. Extraction of vertebral shapes is accomplished using manual and semi-automated algorithms such as active contour segmentation, active shape modeling, and hierarchical segmentation [8]. These algorithms have been implemented with the intention of applying them to a variety of image features; therefore, the system can be extended to a wider variety of images and text data. The resulting vertebral shape boundaries are then treated as closed polygons and represented in a variety of forms, such as Polygon Approximation, Fourier Descriptors, or geometric shape properties that help uniquely identify and characterize them [9]. Traditionally, these representations are stored as a feature vector and compared individually with the query. However, for large image archives, this linear comparison approach would be impractical. SPIRS therefore uses a coordinate tree to efficiently index shapes and optimize retrieval time. Embedded with the indexing process are appropriate distance measures which identify matching shapes [10].

Client Applet

The client, shown in Figure 2, is a Java applet that runs in a Web-browser and provides the user with the necessary functionality to query for particular shapes of interest and interact with the results. The applet (i) is a thin client that provides the

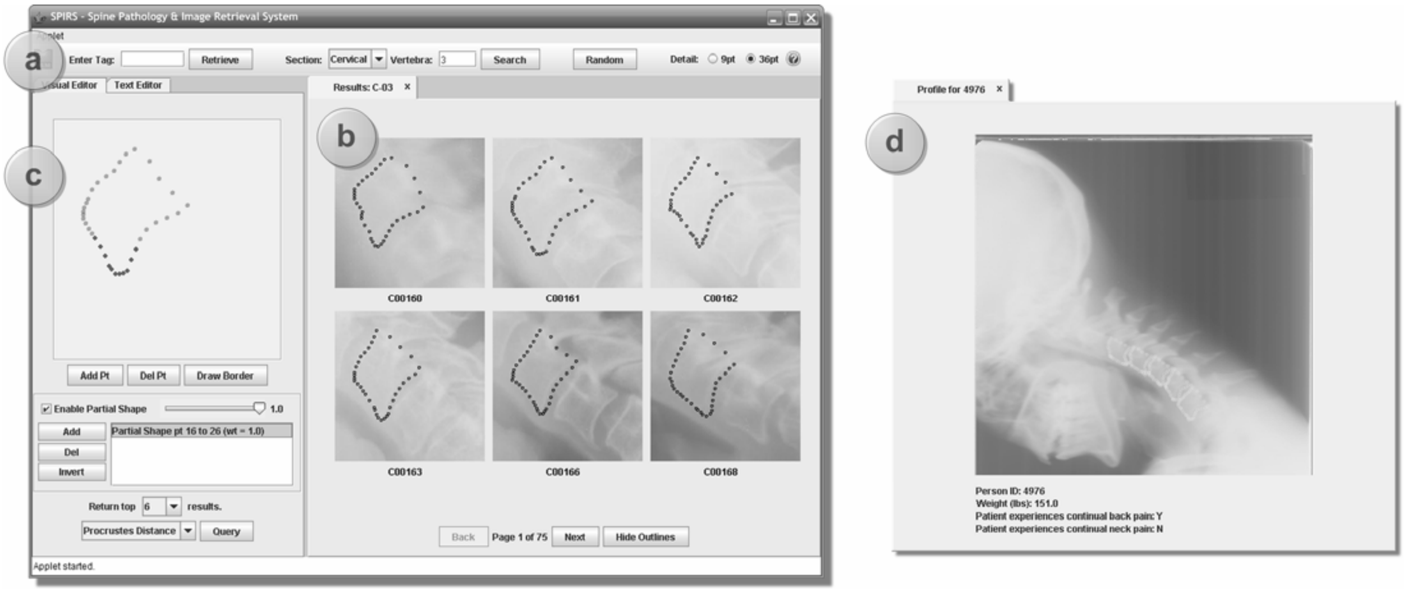

Figure 2 - A screen capture of the SPIRS client applet (left) and a cropped screen capture of the overall results view tab (right).

look-and-feel of a standalone application but uses minimal computing resources, (ii) does not require installation of files onto the local drive, (iii) is easily modifiable without affecting core-CBIR functionality, and (iv) is platform independent.

Query Toolbar

The query toolbar (Fig. 2a) supports query-by-example by allowing users to find vertebral shapes or spine x-ray images that match descriptors stored in the text database, such as image tag or vertebra type and number. A query vertebral shape can be selected from the returned x-ray images in the Overall View or from the selection of vertebral image crops in the Cropped View shown in Fig. 2b.

Overall View

The Overall View (Fig. 2d) displays the entire patient x-ray and allows the user to examine and select any of the segmented vertebral shapes. In addition to the unprocessed version of the scanned image, two forms of enhanced images are also available for improved visualization of subtle detail, viz., one processed with unsharp masking to enhance edges, and the other an adaptive histogram-equalized image which improves contrast. The user may use the mouse hover over and view the vertebral outlines as image overlays. A mouse click on the desired shape selects and transfers it to the query editor.

Cropped View

The Cropped View (Figure 2b) allows a user to inspect multiple vertebral shapes at once. Each displayed vertebra is cropped from the original patient x-ray using the boundary information associated with its shape. The crops are then normalized so that each vertebra is facing in the same direction, which simplifies the comparison between various vertebrae.

Query Editor

The query editor (Figure 2c) enables the user to pose queries using query-by-sketch, select text fields, and supply information pertaining to health history, anthropometric data, quantitative imaging parameters, and demographic data. The visual query component of the editor is a canvas where the user may choose to draw an entirely new shape or edit points that have been imported from an existing shape found using the query toolbar. The query editor also supports multiple partial shape query specification. This feature provides the functionality to select parts of the vertebral shape enabling the algorithms to focus on these boundary intervals, which may exhibit significant pathology. Partial shape querying is shown in the figure as the highlighted interval along the vertebral boundary. Finally, the user can configure the query execution parameters by selecting the retrieval algorithm and limiting the number of results.

Results View

The Results View offers two displays. The Cropped View (Figure 2b) shows returned matching vertebrae while the Overall View (Figure 2d) appears when a user selects one of the resulting vertebral shapes. In the Overall View, the entire x-ray of the patient is shown with the matching vertebra highlighted. Relevant text fields from the text database are displayed alongside the image and may be customized by the user. A unique feature of this system is the ability to use resulting shapes as queries. This can be considered as a form of relevance feedback through the use of iterative querying. More advanced relevance feedback methods have been explored [12] as standalone MATLAB programs and are under development for use with the Web-based system.

Gateway

SPIRS utilizes standard Web communication interfaces to enable interaction with its components. The gateway is a Java servlet that acts as a mediator between client applets and server-side components. It manages multiple simultaneous connections (users) as separate session and queues requests to the core CBIR engine. The servlet translates text components into SQL query syntax for interacting with a MySQL database. In addition, the core architecture of the gateway makes interaction with other geographically distributed CBIR systems possible. An example is the interoperability between IRMA [5],

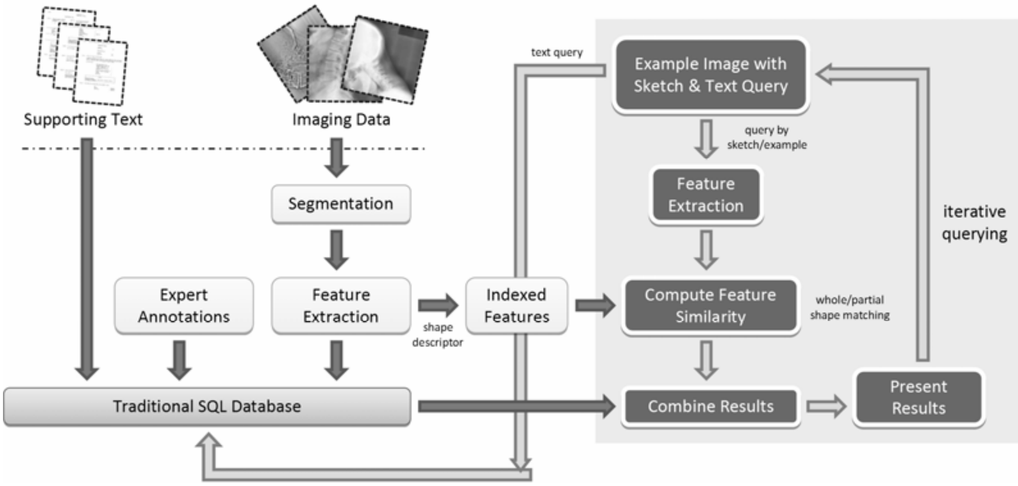

Figure 3 - The query handling process (shaded region) in the context of the overall retrieval system.

which is a system located in Germany for CBIR of medical images using overall image intensity features, and SPIRS. While IRMA's interface provides shape query support for subtle localized pathology, it utilizes the SPIRS server located at the NIH in the USA for computing shape-similarity of spine x-rays.

Query Handling

Although much CBIR research has focused on overcoming technical challenges such as developing accurate and efficient retrieval algorithms, in order for widespread CBIR adoption and use, systems need to provide users with the ability to exploit the capabilities of these algorithms in managing large biomedical multimedia databases. With this goal, we believe that SPIRS allows users to pose meaningful queries that have both clinical and research applicability.

Partial Shape Matching

A problem with many current CBIR approaches for medical images is that they operate on the image as a whole. Whole shape matching methods have been demonstrated to be more effective in identifying shapes with gross similarity because these algorithms operate over the entire shape. They are ineffective, however, in matching subtle shape characteristics, which may have critical pathology. It has been observed that pathologies of interest on a vertebral outline are often localized along a short interval on the boundary. For example, osteophytes are only expressed along the "corners" on the vertebral boundary as seen in the sagittal view. Partial shape matching (PSM) [7, 11] has been implemented in SPIRS allowing users to sketch or identify only the local interval of interest on the vertebral boundary. For instance, consider a user interested in identifying patients in the NHANES II data with a claw anterior osteophyte, which occurs when a spur of triangular shape rises from the vertebral rim and curves toward the adjacent disk. In the query editor, the user would highlight the subset of points that comprise the spur, weight these points greater than the other points in the shape, and execute the PSM query. Initial experiments have shown that PSM can be used to automatically classify the severity of a pathology depicted in the query shape such as whether the claw osteophyte is *slight*, *moderate*, or *severe* depending on criteria defined by domain experts [13].

Query Execution

Figure 3 illustrates the process of query execution and how it integrates with the overall CBIR system. When the user formulates a hybrid text-image query, the query is separated into its visual and textual components. Relevant features and annotations are extracted from the visual component and matched against the indexed features using specific similarity computation methods. Parameters from the textual component are executed against the database containing survey and patient data. The results of these individual searches are then combined before the results are presented to the user. For example, a user who is researching the correlation between certain shape features and the degenerative spine disease spondylosis may be interested in the intersection of results between a visual query depicting, for example, a traction spur and patients who have been diagnosed with spondylosis.

Interaction with Results

The standard method of reviewing the outcome of a query is through the Results View, which displays the whole patient x-ray with the matching vertebral shape highlighted and relevant text descriptions displayed alongside the image. The retrieval algorithms have been evaluated to 68% relevance (precision and recall) when querying for specific osteophyte type (claw or traction) or severity [13]. An additional 22% performance improvement is observed through use of relevance feedback methods [12, 13]. The shortcomings in the performance are linked to (i) erroneous determination of the query semantics and (ii) limitations of the shape matching algorithm. To overcome these problems, SPIRS allows users to select a matching result as the basis of a new query. Traditionally, user interactivity has helped in minimizing similar problems with text retrieval, and user feedback has often been analyzed and employed to improve retrieval relevance. By allowing users the option to iteratively query and refine results, SPIRS implements a basic form of relevance feedback, which will be enhanced with a novel advanced weighted hierarchical feedback method using short-term memory [12] and other approaches as they are migrated from laboratory prototype routines.

Discussion

Image management and pathologically sensitive content-based image retrieval systems are increasingly necessary to interact with the growing volume of biomedical imaging data. In spite of their acknowledged importance, shortcomings in current approaches have prevented their widespread acceptance into medical research, practice, and education. We believe that a biomedical CBIR system should be easily accessible, extensible, and capable of supporting a rich set of segmentation, validation, indexing, query, retrieval, and visualization methods developed using open software and standards. It is difficult for individual systems to support unique requirements of different biomedical images. This can be addressed by providing the capability to combine this system with others (possibly geographically distant) that have complementary features. The SPIRS framework is capable of interacting with and retrieving relevant information from large databases of image and patient data using hybrid image and text query methods. It implements novel shape representation and similarity matching embedded with an index tree that allows efficient retrieval. It aims to capture query semantics through support of advanced mechanisms like multiple partial shape matching and iterative querying that provides simple yet effective relevance feedback to the system. SPIRS is built using open standards and is simultaneously developed as a service, which enables its integration with other complementary information retrieval systems. 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 [13].

Future Work

Future goals for the project include (i) data collection, assimilation, and validation, (ii) system feature enhancements, (iii) improved retrieval quality by learning from user feedback, and (iv) improved user interaction and visualization. Extracting and validating vertebral shapes is an ongoing process which will (at regular intervals) add to the 8,000 shapes indexed currently. In addition, three board certified radiologists are validating the segmented shapes and identifying relevant pathology on them and on the spine. Planned feature enhancements include integration of a generalized shape segmentation toolbox, which is a currently a standalone application, incorporation of additional similarity, relevance feedback, and visualization algorithms, and development of a formalized XML specification for integrating other local and global information systems (or resources) such as WebMIRS and IRMA with core SPIRS services. A multi-user comprehensive qualitative study of the interface is also planned.

Acknowledgements

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). The authors thank Dr. Alex Bui at UCLA for his valuable feedback.

References

[1] Smeulders A, Worring M, Santini S, Gupta A, Jain R. Content-based image retrieval at the end of the early years. IEEE Trans PAMI. 2000; 22(12):1349-80.

[2] 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.

[3] 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.

[4] 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. Comp Vis & Im Understanding 1999; 75(12):111-32.

[5] Thies C, Güld MO, Fischer B, Lehmann TM. Content-based queries on the CasImage database within the IRMA framework. Lec Notes Comp Sci 2005;3491:781-92.

[6] Catarci T, Costabile MF, Levialdi S, Batini C. Visual query systems for databases: a survey. J Vis Lang & Comp. 1997;8(2):215-60.

[7] Antani S, Long LR, Thoma G. Content-based image retrieval for large biomedical image archives. Proc. 11th World Cong Med Info 2004:829-33.

[8] Long LR, Antani SK, Thoma GR. Image informatics at a national research center. Comp Med Im Graphics 2005;29:171-93.

[9] Antani SK, Natarajan M, Long JL, Long LR, Thoma GR. Developing a comprehensive system for content-based image retrieval of image and text from a national survey. Proc. SPIE Med Imaging; 2005:152-61.

[10] Qian X. Shape indexing and its optimization in medical image databases. PhD thesis, Yale University; 2005.

[11] Shou J, Antani S, Long LR, Thoma GR. Evaluating partial shape queries for pathology-based retrieval of vertebra. In: Proc. 8th World Multiconf Sys, Cyb & Info.; 2004:155-60.

[12] Xu X, Lee DJ, Antani SK, and Long LR. Relevance feedback for spine x-ray retrieval. Proc. 18th Int Symp CBMS; 2005 pp. 197-202.

[13] Thoma GR, Long LR, Antani S. Biomedical imaging research and development: knowledge from images in the medical enterprise. Tech Report. LHNCBC, NLM, NIH; Sep. 2006. Report No.: LHNCBC-TR-2006-002.

Address for correspondence

William Hsu
UCLA Medical Imaging Informatics Group
924 Westwood Blvd, Ste 420

Los Angeles, CA 90024, USA
Email: willhsu@mii.ucla.edu