Content-Based Retrieval in Picture Archiving and Communication Systems

Essam A. El-Kwae, Haifeng Xu, and Mansur R. Kabuka

A COntent-Based Retrieval Architecture (COBRA) for picture archiving and communication systems (PACS) is introduced. COBRA improves the diagnosis, research, and training capabilities of PACS systems by adding retrieval by content features to those systems. COBRA is an open architecture based on widely used health care and technology standards. In addition to regular PACS components, COBRA includes additional components to handle representation, storage, and content-based similarity retrieval. Within COBRA, an anatomy classification algorithm is introduced to automatically classify PACS studies based on their anatomy. Such a classification allows the use of different segmentation and image-processing algorithms for different anatomies. COBRA uses primitive retrieval criteria such as color, texture, shape, and more complex criteria including object-based spatial relations and regions of interest. A prototype content-based retrieval system for MR brain images was developed to illustrate the concepts introduced in COBRA.

KEY WORDS: content-based image retrieval, medical image databases, medical information system, picture archiving and communication systems, information retrieval.

1. INTRODUCTION

PICTURE ARCHIVING and communication systems (PACS) are widely used for the acquisition, storage, communication, and display of vast amounts of medical images and text files in a digital radiology department. Compared to traditional film environments, PACS systems reduce cost and turnaround time and provide access to patient studies regardless of location, which allows technologies like telemedicine[1] to be used to improve patient care.

The radiology environment has various multimedia components including images of various modalities, regions of interest, audio dictations, graphic annotations, and text reports.[2] Current PACS systems can only handle queries based on a few keys, such as patient name or hospital ID number. They lack the means to classify and index multimedia files by their information content. Thus, much of the rich, useful patient information stored in PACS has not been used for broader medical practice, research, and teaching.[3]

Content-based image retrieval may be classified into 3 levels that range from the highly concrete to the very abstract.[4] Level-1 comprises retrieval by primitive features such as texture, color, and shape as with the QBIC system[5] and MIT PhotoBook.[6] Level-2 comprises retrieval by derived attributes involving some degree of logical inference about the identity of the objects depicted in the image.[7] Level-3 comprises retrieval by abstract attributes and possibly subjective reasoning about the scenes depicted. The International Standardization Organization (ISO) has begun to clarify the scope of a multimedia content-description interface standard, known as MPEG-7.[8] Whereas Level-1 retrieval may or may not require image segmentation into constituent objects, Level-2 retrieval depends on image segmentation. In spite of the amount of medical image segmentation research in the past years, a general algorithm for segmenting medical images has not yet been developed. However, segmentation has been somewhat successful on specific anatomies, although in some cases, interactive user assistance is required, eg, automatic segmentation of brain MR images,[9] skin lesions,[10] labeling of MR brain images,[11] chest CT images,[12] and extraction of right and left ventricular chambers from cardiac cine MR images.[13]

Several architectures for content-based retrieval in PACS environments have been proposed in recent years. The active index was introduced to prefetch images and multimedia data and to facilitate similarity retrieval.[14] A technique that uses wavelet feature vectors to measure the presence of structures of variable size and orientation and to provide a multiresolution analysis of object contours was proposed.[15] An image indexing method based on the Karhunen-Loève transform has been used to develop a content-based search engine for

From the Department of Computer Science, University of North Carolina, Charlotte, NC, and the Departments of Radiology and Electrical and Computer Engineering, Center for Medical Imaging and Medical Informatics, University of Miami, Miami, FL.

Address reprint requests to Mansur R. Kabuka, PhD, Department of Electrical and Computer Engineering, University of Miami, 1251 Memorial Dr, MCA Rm 406, Miami, FL 33146.

0897-1889/00/1302-0001$10.00/0
doi:10.1053/di.2000.5419

tomographic image databases.[7] A prototype multimedia database system, WebMIRS, to provide WWW access to biomedical databases, is being developed.[16] Super-performance computers were used for the discovery, selection, and optimization of medically useful image feature sets via genetic algorithms and simulated annealing methods.[17] The I2C (Image Indexing by Content) is an object-oriented architecture for the indexing, storage, and retrieval of medical images by content.[18,19] Wong et al[3] developed the content-based image retrieval (CBIR) system architecture as an additional function module of the PACS in which segmentation and extraction of medical images are done interactively. An indexing scheme for content-based retrieval based on the color content of a dominant object was developed in.[20]

Current techniques have several drawbacks that can be summarized as follows:

- The lack of support for new health care and technology standards that allow systems from different vendors and on different platforms to interoperate.
- Content-based retrieval is based on either Level-1 or Level-2 criteria, but not both.
- Study classification and description generation are usually performed manually, which requires user intervention for each study entered into the system. Manual processing is time-consuming and may not be feasible for large medical image databases.
- Adding content-based retrieval capabilities to a PACS system requires redesigning the system rather than upgrading it.

The advent and recent popularity of open systems and emerging standards have created many new areas of interest in the field of medical imaging from a technical standpoint.[21] Digital Imaging and Communications in Medicine (DICOM) is a standard that addresses the issue of vendor-independent data formats and data transfers for digital medical images.[22] The Health Level Seven (HL7) Standard,[23] an American National Standards Institute (ANSI) approved standard for electronic data exchange in health care, enables disparate computer applications to exchange key sets of clinical and administrative information. In addition, the Common Object Request Broker Architecture (CORBA)[24,25] is a specification and an architecture that provides implementation/platform independent access to distributed objects. COBRA uses all three standards to provide PACS systems with vendor, platform, and language independence. This allows the COBRA components to be reused, to be easily maintained and to interoperate with existing systems.

This report introduces a model called COBRA for extending PACS capabilities to include content-based image retrieval. COBRA addresses the issues discussed above as follows:

- COBRA is based on state-of-the-art health care and technological standards. Thus, COBRA components can interact with other standard-based systems regardless of the vendor, platform, or programming language used.
- Content-based retrieval is based on both Level-1 and Level-2 criteria. Efficient indexing schemes are used at both levels.
- An automatic study classification algorithm is included, which eliminates most of the manual processing required for image classification. User intervention for study classification is minimized, and the algorithm dynamically "learns" from previous errors to improve its performance with time.
- COBRA may be incorporated by upgrading an existing PACS system without having to redesign it from scratch.
- In addition to text-based queries, COBRA uses three types of content-based retrieval, query by example, query by region of interest (ROI), and query by sketching.

A prototype was developed to illustrate the concepts introduced in COBRA. The anatomy selected was MR brain images. Images were segmented using an automatic segmentation algorithm.[9] A simple anatomy tree for MR images was built. The study classification algorithm was implemented and tested on a database of more than 22,000 studies. The communication between prototype components was defined as a set of interface definition language (IDL) modules using the CORBA middleware standard to ensure the internal interoperability between COBRA components and the external interoperability to components of other systems. The graphical user interfaces (GUIs) for the COBRA query interface were designed and implemented.

The rest of this article is organized as follows. In section 2, a detailed description of COBRA is given including the main components, the study classification algorithm, and the indexing structure. Details of the prototype are given in section 3, followed by conclusions and future directions in section 4.

2. THE COBRA-PACS SYSTEM DESCRIPTION

COBRA expands the range of capabilities of existing PACS systems by providing content-based querying capabilities. The PACS infrastructure is comprised of several cooperating components. The basic functions of a PACS fall under three major headings, data input, data management, and data output, as in many information management systems.[3] A generic PACS architecture is shown in Fig 1, based on studying several existing PACS architectures.[2,3,18,19,21,26,27]

The database population engine (DPE) receives patient demographic data from the hospital information system (HIS), the radiology information system (RIS), and the image data in DICOM format from various sources (such as scanners and film digitizers). Within the DPE, the database storage engine (DSE) extracts the patient, study, series, and image information and stores it into a physical database. Databases may be distributed and/or replicated. The transfer of image files may be done according to different schemes including prefetching and direct image selection. The database server is responsible for hiding and managing the particular aspects of each physical database so that client applications can view the separate databases as a single database.[28] The radiology workstation (RW) includes components that interact directly with PACS users. The query interface (QI) provides a user-friendly GUI that allows PACS users to retrieve data based on simple, text-based criteria such as patient name, acquisition number, study date, or modality. The report generation (RG) module allows users to dictate their reports to the PACS system. Dictations may be transcribed automatically by voice recognition, or manually by a transcriptionist. The image processing engine (IPE) performs simple image processing tasks such as zooming, panning, window and level setting, and cine mode display. The communication system ties together all sources of information of various types in such a way to provide client applications with the ability to view the image data.[21] The database retrieval engine (DRE) searches the database for data satisfying the query criteria. In traditional PACS systems, the DRE is relatively simple because it deals only with simple text-based queries. In COBRA, the DRE will have added responsibilities, as will be seen.

An example of the above architecture is the PACS system (the Digital Radiology Infrastructure, DRI) developed at the Center of Medical Imaging and Medical Informatics, University of Miami[27] (Fig 2). Mapping the DRI components to the generic architecture reveals that the external inter-

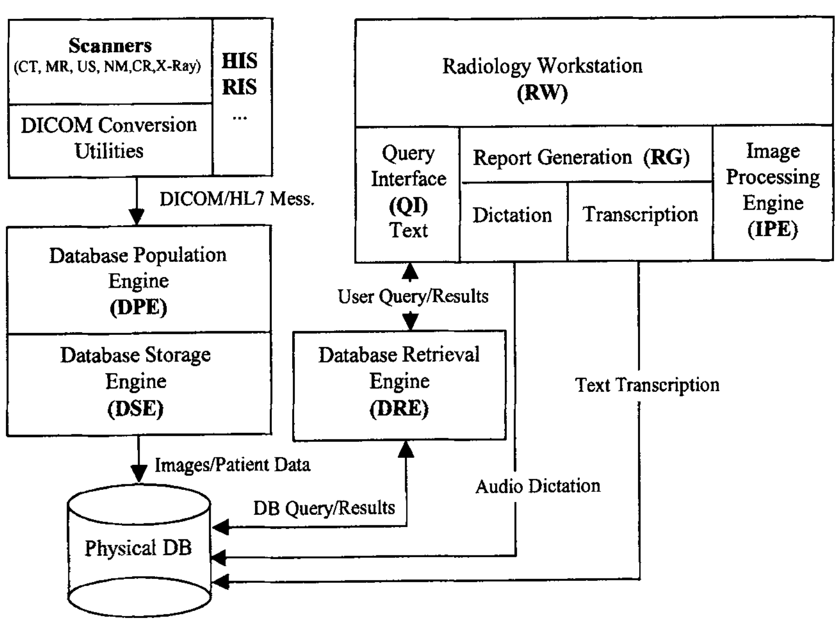

Fig 1. A generic PACS architecture.

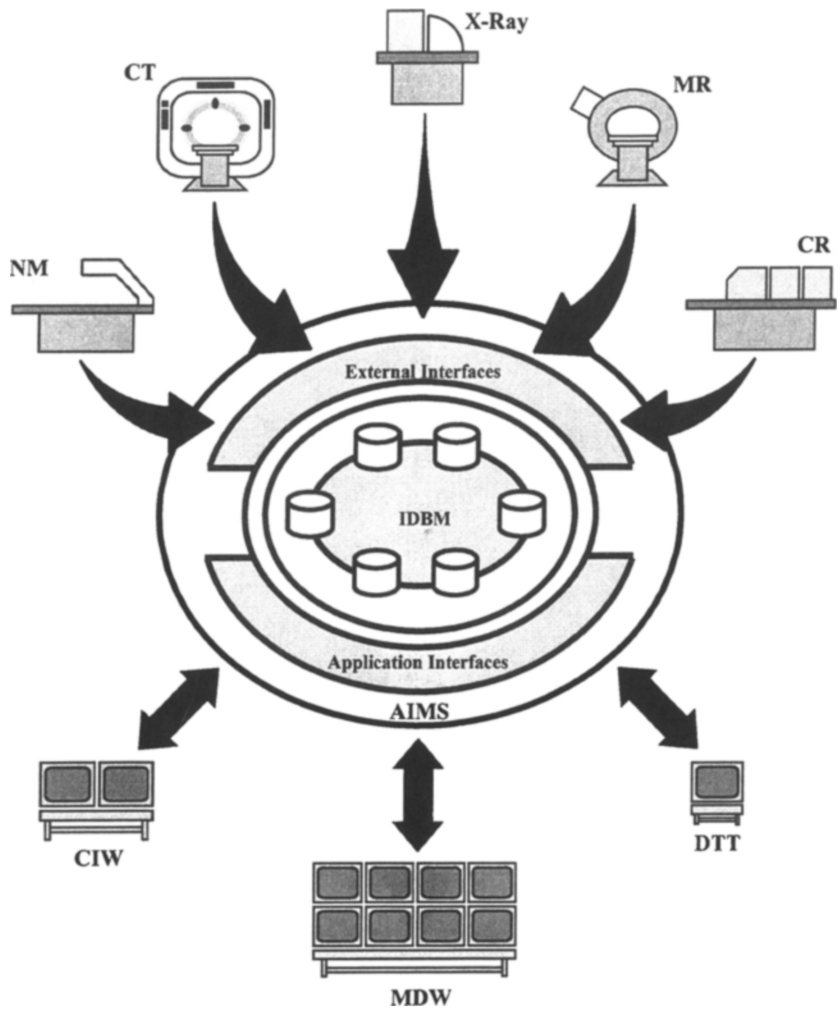

Fig 2. The DRI architecture.

faces are equivalent to the DPE, and the advanced image management system (AIMS) represents the DRE. In addition, the dictation transcription tools (DTT), the multimedia diagnostic workstation (MDW), and the clinician imaging workstation (CIW) constitute the RW of the generic PACS architecture.

The block diagram of COBRA is shown in Fig 3. In addition to the core of the traditional PACS architecture (Fig 1), some components have enhanced roles. These components include the DRE, the QI, and the IPE. New components are added to support the added functionality, such as the descriptor generator engine (DGE) that generates content-based descriptions and the logical database that stores those descriptions.

The basic content-based functionality works as follows. At the time a new study is stored into the PACS database, the content-based description of this study is generated and stored into the logical database using the flowchart shown in Fig 4. A table lookup is used to select the best segmentation algorithm for this anatomy. Segmentation allows the spatial relations between objects to be extracted such as directional relations (eg, north, south, east, west) and topological relations (eg, disjoint, overlap, covers). Spatial relations are important to support object-based high-level queries. In addition, all study, series, image, and patient information pertinent to content-based retrieval are extracted from the DICOM files and saved to the logical database. Examples of these fields are patient age, gender, modality, and orientation. Other primitive content-based descriptors may be ex-

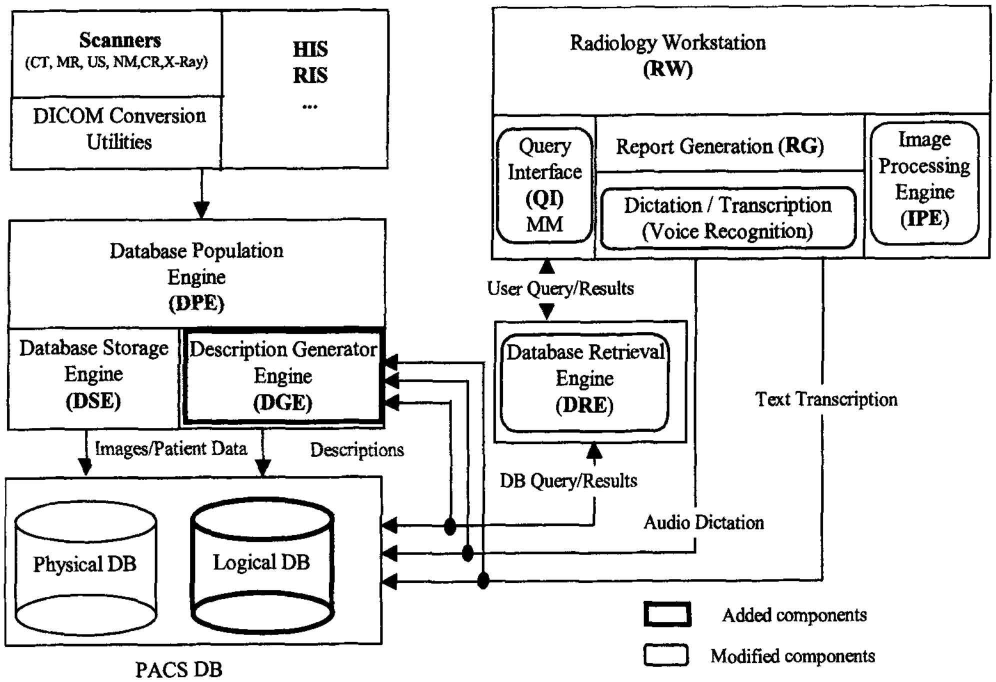

Fig 3. The COBRA system design.

tracted from the given images such as color, texture, and shape.[5,20] COBRA is configurable as to which primitive descriptors to use and how they are represented. Finally, if a user selects one or more regions of interest (ROI) from images in the study, the DGE generates similar descriptions for those regions. When a content-based query is submitted, the query description is generated and matched against other descriptions in the logical database. Because exact matches generally are not feasible for content-based retrieval, query results are ranked based on their degree of similarity to the submitted query. The ranked list of results is returned to the user, who may then refine the query and resubmit it for more specific results.

Several components of the traditional PACS architecture are enhanced in COBRA. The dictation and transcription modules are merged into a single module, and the IPE is extended to handle the extraction of primitive content information from images including color, texture, and shape information. In addition, the QI is extended to allow content-based query creation using textual attributes (eg, patient, study, series, or image information), or visual information (eg, color, texture, shape, or spatial relation).

The DRE is enhanced to allow content-based similarity retrieval based on the submitted query.

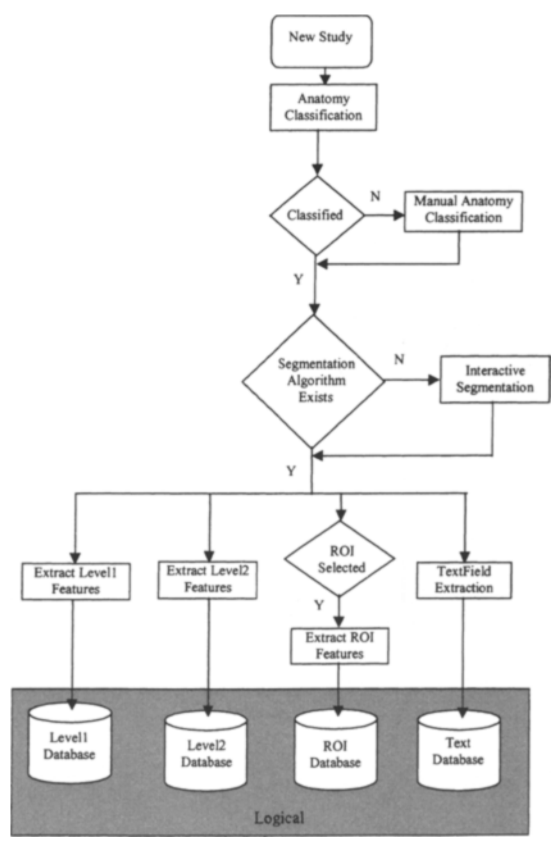

Fig 4. Study description generation.

The DRE translates a user's query into a set of queries understood by the different indices storing content-based information.

Image storage and communication in COBRA is based on the DICOM standard to ensure vendor-independent data format and data transfer of medical images.[22] In DICOM, real-world entities such as patients, studies, and images of various modalities are represented by information object definitions (IOD) composed of groups of attributes termed modules. The textual DICOM IODs include various fields relevant to content-based retrieval such as patient age, gender, and study modality.

The high-level interoperability of COBRA is maintained by adopting the HL7 standard[23] electronic message exchange with HIS and RIS systems. Message formats prescribed in the HL7 encoding rules consist of data fields that are of variable length and separated by a field separator character. Rules describe how the various data types are encoded within a field and when an individual field may be repeated. The data fields are combined into logical groupings called segments. Segments can be defined as "required" or "optional" and may be permitted to repeat.

As will be shown in the prototype, all COBRA interfaces are implemented using the CORBA standard[24,25] to provide implementation/platform independent access to distributed objects. CORBA is based on the client-server architecture. Clients issue requests for services. The object request broker (ORB) provides the interoperability and portability layer between different applications. The ORB is responsible for delivering the request to the object implementation in the format to which it can respond. The interface seen by the client is completely independent of where the object is located, what programming language is used to implement it, or any other aspect that is not reflected in the object's interface. This interface is defined using IDL.

Medical image segmentation, whether automatic or semiautomatic, has been somewhat successful when applied to specific anatomies.[9-13] In COBRA, the DGE determines the anatomy of the current study and selects an appropriate segmentation algorithm. As new segmentation algorithms are developed or current algorithms mature, the DGE is updated without major changes to COBRA.

Anatomy determination by the DGE has to be accurate for the correct segmentation algorithm to be instantiated. In the DICOM standard, the anatomy of an image may be transferred using the BODY_PART_EXAMINED tag in the series IOD (using specific names such as SKULL, CSPINE, ABDOMEN, PELVIS) or using the STUDY_DESCRIPTION tag in the study IOD. The STUDY_DESCRIPTION is a free text field that may not adhere to any standard convention. The anatomy classification algorithm (Fig 5) extracts the anatomy from the STUDY_DESCRIPTION field. The input to this algorithm is an anatomy tree similar to the one in Fig 7. The first level of the tree represents the main body part, and the second level represents a specific technique or body subpart. Different trees may be used for different imaging modalities. The algorithm uses the edit distance between the anatomy to be classified and the anatomy classes in the tree to classify the anatomy to a specific class. The edit distance is defined as the number of changes required to transform one string to another. If a similar anatomy cannot be found within the maximum specified edit distance, the input anatomy is declared ambiguous and the user is prompted to enter the correct classification. The user input is used to update the anatomy tree for future reference. Thus, each node in the tree will behave as a thesaurus for all the anatomies classified as belonging to this node.

Indexing addresses the issue of how the information should be organized so that queries can be resolved efficiently and relevant portions of the data can be quickly extracted. In a medical image database consisting of thousands of images, it is not practical to compute the similarity of a query image with every database image. An efficient indexing mechanism is required to filter irrelevant images as quickly as possible and return only the relevant ones. To improve the performance of COBRA, different indexing mechanisms are used based on the level of retrieval. For Level-1 retrieval, the R-Tree[29] and its variant, the R^*-Tree,[30] are used. R-Trees, proposed by Guttamn,[29] are widely used for spatial and multidimensional databases. The R-Tree is a hierarchical data structure derived from the B-Tree. Arbitrary geometric objects can be handled by an R-Tree by representing each object by its minimum bounding rectangle (MBR). The MBR is the smallest axis parallel hyper-rectangle that encloses the object. The R-Tree and its variants have been widely used for Level-1 retrieval.[5]

Primitive features such as color, texture, and

Input: Anatomy attribute of the study to be classified (A)
Anatomy classes represented by a multilevel anatomy tree (T)
The maximum edit distance to be tolerated (D)
List of word delimiters (d) (e.g. space, backslash, dash, slash)
Output: The location of A in T
Procedure:
Step 1. Parse A into a set of n words $\{ w_i \mid i = 1 \text{ to } n \}$ based on the delimiters in d.
Step 2. Create a set (C) of candidate anatomies including all items in the first level of T
Step 3. Create an empty result list (L)
Step 4. While C is not empty do
4.a. Calculate the edit distance between each of the w_i and each item in C
4.b. For every edit distance $\leq D$,
4.c.1. Copy the corresponding anatomy to L
4.c.2. Remove the parent of this anatomy from L
4.c.4. Remove the corresponding word from w_i
4.c. If leaf level has not been reached
C = all items in the next level of T
else C = empty list
Step 5. If L is empty, return *UNCLASSIFIED*
else if L has more than one item, return *AMBIGUOUS*
else return the only element in L

Fig 5. The anatomy classification algorithm.

shape are extracted from an image. The set of features is represented by a feature vector that is inserted into an R-Tree. To retrieve an image, the feature vector representing the query image is created and the R-Tree is searched for a set of nearest neighbors, which represent similar images.

For Level-2 retrieval, the Two Signature Multi-Level Signature File (2SMLSF)[31] is used. The signature file is a filtering mechanism that eliminates most irrelevant images to the given query. Signature files have been widely employed in information retrieval of both formatted and unfor-

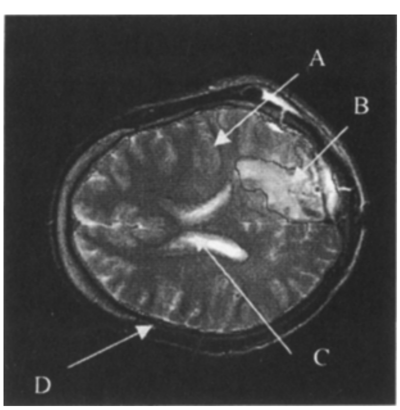

(a) Example Image with

objects:

Cerebrum region(A), Abnormal signal intensity(B), Ventricle(C), Extraventricle(D)

Object Signatures:
(A): 001 000 110 010
(B): 010 001 100 010
(C): 001 000 110 010
(D): 001 010 110 000

Image Signature: 011 011 110 010

(b) Object and image signatures

Queries	**Signature**	**Result**
Cerebrum region	001 000 110 010	Match
Aorta	000 010 100 101	No Match
Ab. Signal and Extraventricle	011 011 110 010	Match
Left Lung	010 010 100 000	False drop

(c) Sample queries and matching results

Fig 6. Signature generation and comparison based on superimposed coding.

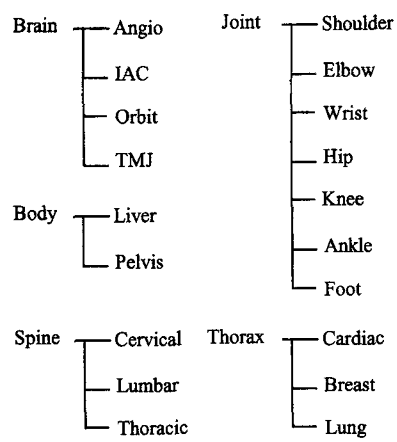

Fig 7. A sample MR anatomy classification tree.

matted data.[32] Recently, signature file techniques were applied to image databases.[31,33] Signatures may be obtained in a number of ways, the most common of which is superimposed coding in which each object (or object pair) in an image is hashed into a word signature. An image signature is generated by superimposing (*OR*ing) all its individual signatures. To resolve a query, the query signature is generated and matched (*AND*ed) against image signatures. Figure 6a is an example showing the generation of an image signature from object signatures. The image has four objects: A, B, C, and D. Figure 6b shows example signatures assigned to each object. These signatures are *Or*ed together to create the image signature. Figure 6c shows some queries and the results of matching their signatures to the image signature.

After images are segmented, the individual objects are extracted and labeled. The image is then encoded into a signature that is inserted into the signature tree. Queries are resolved by matching each query signature against the signature tree.

3. THE COBRA-PACS PROTOTYPE

A prototype was built to demonstrate the concepts introduced in COBRA. Initially, a simple anatomy tree for classifying MR studies was developed (Fig 7). The anatomy classification algorithm (Fig 5) was tested on a database of about 22,000 MR studies. The anatomy field of each study in the database was extracted, and repeated anatomies were discarded. It was found that there are 1,321 textually unique anatomies. A prototype classification system based on the anatomy tree in Fig 7 was built as shown in Figs 8 and 9. The prototype was built using the C++ language and was tested on the Windows NT operating system on a Pentium

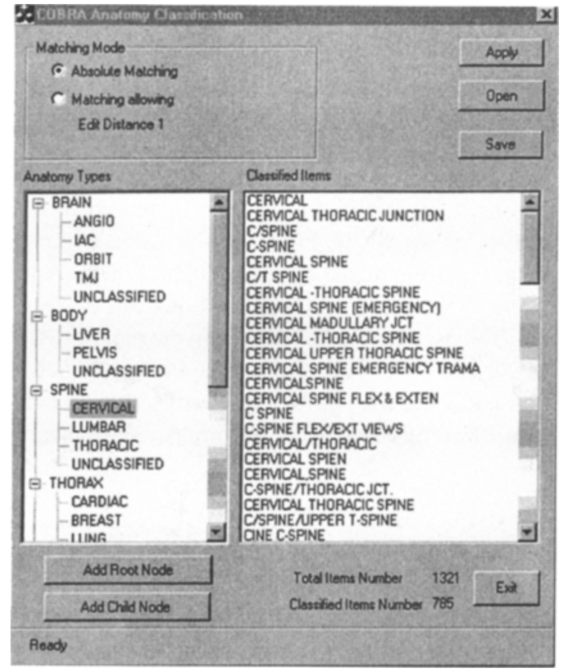

Fig 8. Anatomy classification prototype (absolute matching).

PC. Before classification, the user selects the maximum tolerable edit distance between the query anatomy and the database anatomies. In this prototype, edit distances of *zero* (exact matches) and *one*

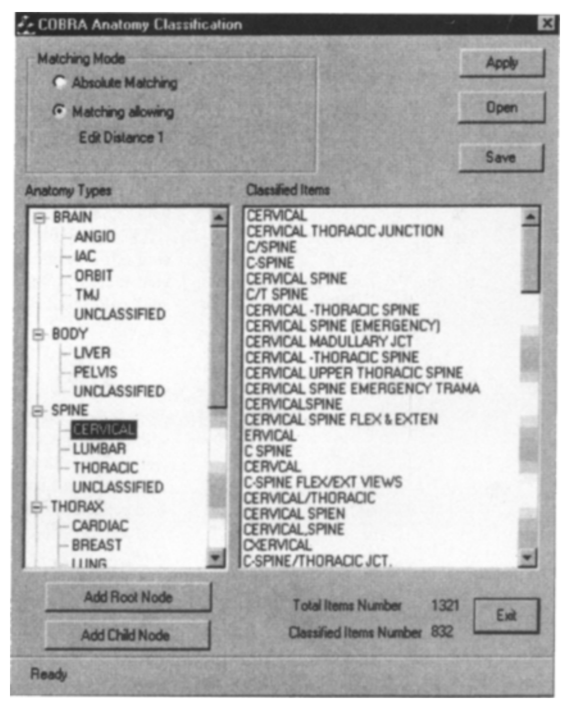

Fig 9. Anatomy classification prototype (edit distance = 1).

(a single spelling error) were used. For a zero edit distance, 785 (59.4%) of the 1,321 anatomies were classified as shown in Fig 8. For example, anatomy items such as "Cervical," "C/Spine," "C-Spine," "Cervical Spine," and "C Spine" were all classified to the cervical spine node. For an edit distance of one, the classified anatomies increased to 832 (63%), as shown in Fig 9. In this case, common single spelling mistakes were caught. For example, using "Cervcal" or "Cxervical" instead of cervical. The user may select a node in the anatomy tree to view the list of anatomies that were classified to this node and may correct any errors in classification. For unclassified anatomies, the user is prompted to classify them manually, to one of the anatomies existing in the tree or by adding new nodes to the anatomy tree.

The second step in the prototype was to design CORBA middleware interfaces for the different components. The interfaces were designed using the IDL. A sample subset of the interface is shown in Fig 10. The main module is called COBRA, and it is composed of three smaller modules: Data, Client, and DBService. The Data module defines the basic patient and DICOM image information. The Client module defines the different query operations such as query by textual attributes, by sample image, and by region of interest. The

```
module COBRA {
        typedef sequence<char> PBYTE;
        interface ExtInvalidDataException {
                setErrorInfo(in string error_info);
                getErrorInfo(out string error_info);
        };

        module Data {
                struct PatientInfo {
                        Date   studyDate;
                        Time   studyTime;
                        string  patientName;
                        string  modality;
                        string  patientID;
                        Date   birthDate;
                        boolean sex;
                };
                interface ExtDicom {
                        void setPatientInfo(in PatientInfo pinfo) raises(ExtInvalidDataException);
                        void getPatientInfo(out PatientInfo pinfo);
                        void setImageData(in PBYTE image_data) raises(ExtInvalidDataException);
                        void getImageData(out PBYTE image_data);
                };
        };
        module Client {
                interface ExtQuery{
                        void setSqlQuery(in string sql_string);
                        void getSqlQuery(out string sql_string);
                        void setSampleImage(in Data::ExtDicom image)
                                raises(ExtInvalidDataException);
                        void getSampleImage(out Data::ExtDicom image);
                        void setROI(in PBYTE roi_mask) raises(ExtInvalidDataException);
                        void getROI(out PBYTE roi_mask);
                };
        };

        module DBService {
                interface ExtModalityImageServer {
                        void setModality(in MODALITY modality);
                        void getModaility(out MODALITY modality);
                        void setReturnNumber(in int num) raises(ExtInvalidDataException);
                        void getReturnNumber(out int num);
                        void setQuery(in Client::ExtQuery query);
                        void getQuery(out Client::ExtQuery query);
                        void search(out Data::ResultSet) raises(ExtDBException);
                };
        };
};
```

Fig 10. A subset of the IDL interface.

DBService module defines the database interface that searches the database in order to respond to the client requests.

Third, the COBRA prototype was applied to the segmentation and content-based retrieval of brain MR images using an automatic segmentation technique developed by Tsai et al.[9] In this technique, the cerebrum region is extracted using a single threshold, which is computed directly from the image histogram, and a sequence of morphological operations. The cerebrospinal fluid (CSF) regions are detected from T_2-weighted images by adaptive thresholding, and the ventricular and extraventricular regions were identified. From the proton density (PD) images, the brain matter is further classified into gray and white matter using a low-level knowledge-based segmentation rule. Finally, a check for any abnormal signal intensities is made. This includes detection of any lesions and abnormal ventricles.

The user interface for the query design process is shown in Fig 11. Three different types of content-based queries are allowed: (1) retrieval by example, in which the user submits an image and requests a list of similar images, (2) retrieval by ROI, in which the user selects a region in an image and requests images with similar regions, and (3) retrieval by sketching, in which the user draws a sketch of the requested image using a set of icons representing the expected objects. The layout of the object icons in the sketch implicitly specifies the spatial relations in the requested image. In addition to content-based queries, the user may combine other regular textual queries using information such as patient name, ID, gender, age, modality, and study date.

The user may select the maximum number of images to be returned. For example, the query shown in Fig 11 requests the retrieval of up to 50 MR brain images similar to the example given where similarity is based on color histograms. When the user performs the query, the DRE performs the search and returns the results. The user interface for the query results is shown in Fig 12. Images are displayed in pages in the order of similarity to the submitted query. Results may be displayed as shown, or numerically, in which a list of images and their textual information is displayed together with the estimated similarity values.

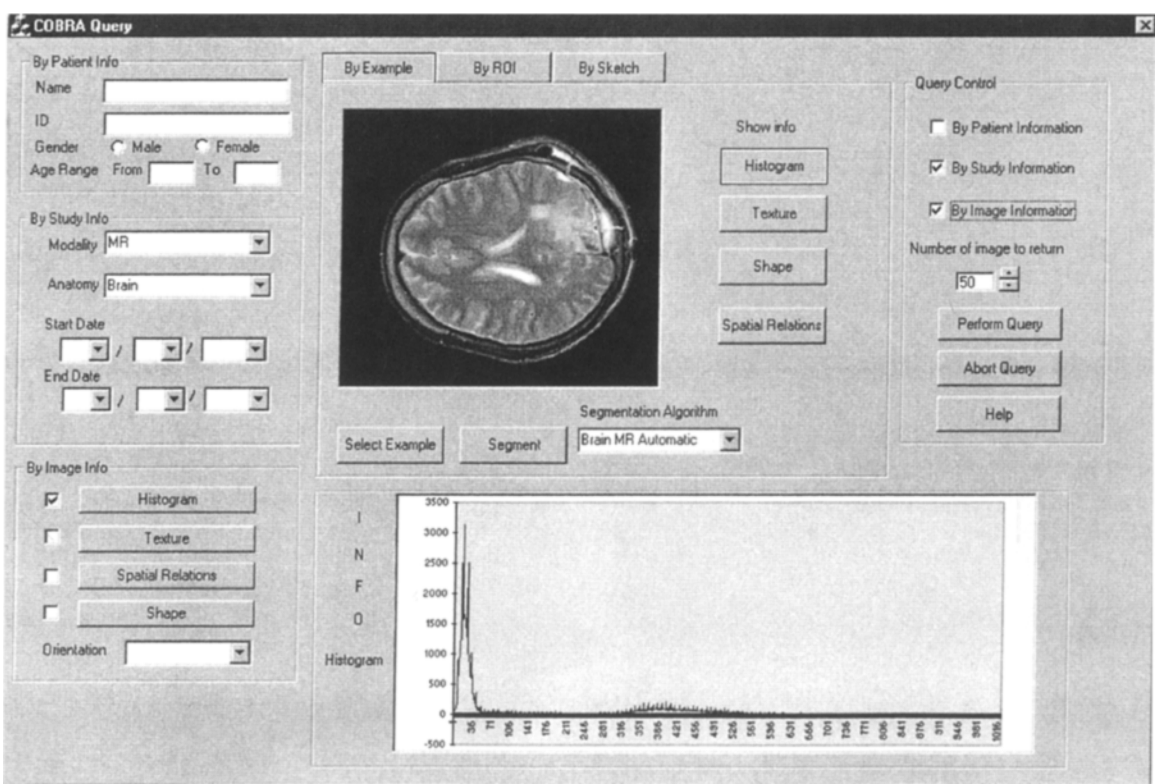

Fig 11. The COBRA prototype query formulation screen.

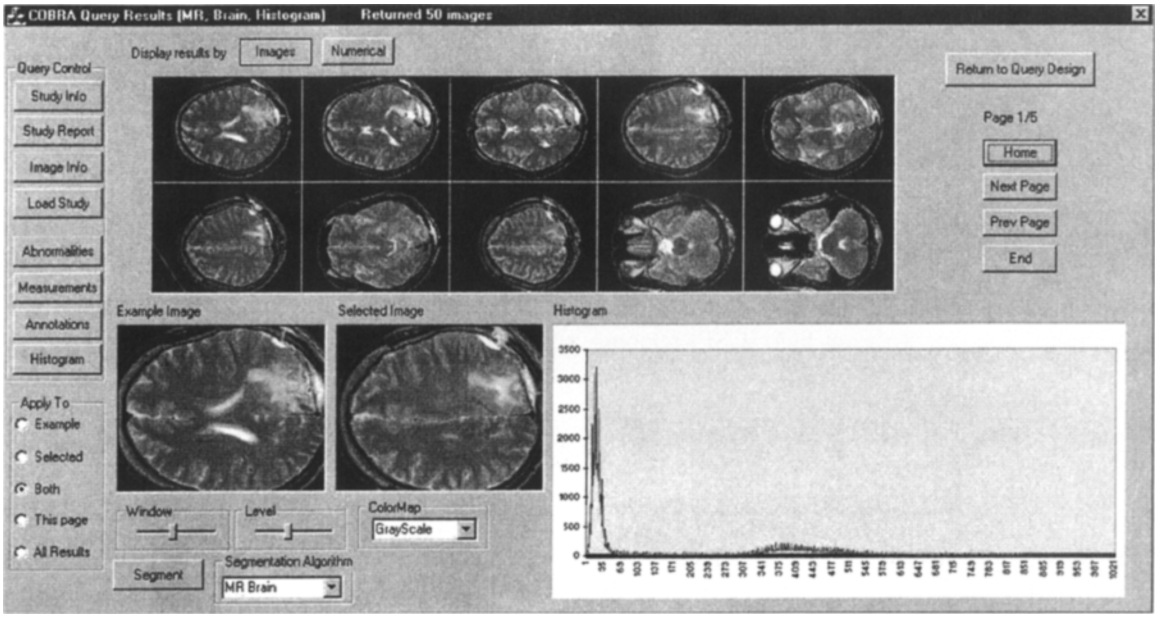

Fig 12. The COBRA prototype query results screen.

The user has a great deal of control with respect to exploration of the results. Various operations may be applied to the example image, the selected image, both images, the displayed page of images, or the entire set of results. The user may select an image and display it side by side with the query image, display the histogram of the two images to explore the areas of similarity, segment one or both images, and/or change the color map or the window and level of the images. In addition, information about the study, including the selected image, may be displayed. The user may view the report or listen to the dictation associated with that study, or may even elect to retrieve the entire study for further exploration. Several other functions may help the examination of results, such as viewing abnormalities or annotations and performing measurements on the displayed images. At any time, the user may return to the query design screen to refine and resubmit the query.

4. CONCLUSION

The COBRA system for content-based image retrieval in PACS systems has been introduced. The architecture may be used to extend existing PACS systems into content-based retrieval systems. COBRA uses two image retrieval levels: Level-1, retrieval by primitive features such as texture, color, and shape; and Level-2, retrieval by derived attributes involving some degree of logical inference about the identity of the objects depicted in the image. The architecture is based on widely used health care standards (DICOM and HL7) and technology standards (CORBA), which makes it feasible for different systems to interoperate regardless of vendor, platform, operating system, or programming language used. COBRA uses efficient indexing techniques to improve retrieval performance. The indexing scheme selected depends on the level of retrieval required. An anatomy classification algorithm is introduced to automatically classify a study given its anatomy and an anatomy tree. Based on the classified anatomy, a suitable segmentation algorithm is selected. The logical description of each image is generated based on Level-1 features, Level-2 features, ROI selection, and textual attributes. The anatomy classification algorithm was tested on a database of MR studies using a simple MR anatomy tree. A COBRA prototype was designed using automatic segmentation of brain MR images. The prototype was used to illustrate the different techniques for querying a COBRA-based PACS system. Future work involves implementing various anatomy-based segmentation algorithms, creating different anatomy trees for different modalities and fully implementing the proposed architecture.

REFERENCES

1. Orphanoudakis SC, Kaldoudi E, Tsiknakis M: Technological advances in teleradiology. Eur J Radiol 22:205-217, 1996

2. Sauer F, Kabuka MR: Multimedia technology in the radiology department. Proceedings of the Second ACM International Conference on Multimedia, ACM Multimedia '94, San Francisco, CA, October 1994

3. Wong S, Hoo K Jr, Huang H: Empower PACS with content-based queries and 3-D image visualization. SPIE 2711: 569-579, 1996

4. Eakins JP: Automatic Image Content Retrieval—Are we going anywhere? Presented at Third International Conference on Electronic Library and Visual Information Research, De Montfort University, Milton Keynes, May 1996

5. Flickner M, Sawhney H, Niblack W, et al: Query by image and video content: The QBIC system. IEEE Comput 28:23-32, September 1995

6. Pentland A, Picard RW, Sclroff S: Photobook: Content-based manipulation of image databases. SPIE Storage Retrieval Image Video Databases II 2185:34-37, February 6-10, 1994

7. El-Kwae EA, Kabuka MR: A robust framework for content-based retrieval by spatial similarity in image databases. ACM Transactions Info Systems, vol 17, April 1999

8. Carlson B: MPEG-7 for multimedia searching: Room for video content recognition? Adv Imaging, March 1997, pp 21-22

9. Tsai C, Manjunath BS, Jagadeesan R: Automated segmentation of brain MR images. Pattern Recognition 28:1825-1837, 1995

10. Lee T, Ng V, McLean D, et al: A multi-stage segmentation method for images of skin lesions. Proceedings of the IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing, Victoria, Canada, May 17-19, 1995

11. Li XH, Bhide S, Kabuka MR: Labeling of MR brain images using Boolean neural network. IEEE Trans Med Imaging 15:628-638, 1996

12. Brown MS, Mcnittgray MF, Mankovich NJ, et al: Method for segmenting chest CT image data using an anatomical model—Preliminary results. IEEE Trans Med Imaging 16:828-839, 1997

13. Goshtasby A, Turner DA: Segmentation of cardiac cine MR images for extraction of right and left ventricular chambers. IEEE Trans Med Imaging 15:56-64, 1995

14. Chang SK: Active index for content-based medical image retrieval. Comput Med Imaging Graphics 20:219-229, 1996

15. Stark HG, Laufkotter GP: Image indexing and content based access to database of medical images with wavelets. Proc SPIE 2569:790-799, 1995

16. Long LR, Pillemer SR, Lawrence RC, et al: WebMIRS: Web-based medical information retrieval system. Proc SPIE 3312:392-403, 1997

17. Wetzel AW: Computational aspects of pathology image classification and retrieval. J Supercomputing 11:279-293, 1997

18. Orphanoudakis SC, Chronaki C, Kostomanolakis S: I2C: A system for the indexing, storage, and retrieval of medical images by content. Med Informatics 19:109-122, 1994

19. Orphanoudakis SC, Chronaki CE, Vamvaka D: I^2Cnet: Content-based similarity search in geographically distributed repositories of medical images. J Comput Med Images Graphics 20:193-207, 1996

20. Ng V, Cheung D, Fu A: Medical image retrieval by color content. IEEE International Conference on Systems, Man, and Cybernetics, October 1995

21. Shebert GM: A communication system for a digitally based radiology environment. MSc Thesis, University of Miami, Coral Gables, FL, May 1995

22. American College of Radiology (ACR), National Electrical Manufacturers Association (NEMA): Digital Imaging and Communications in Medicine (DICOM): Version 3.0. Draft Standard, ACR-NEMA Committee, Working Group VI, Washington, DC, 1993

23. HL7 Organization: Health Level 7 (HL7) Version 2.3, April 1997

24. Object Management Group: The Common Object Request Broker: Architecture and Specification, OMG, revision 2.2, July 1998

25. Orfali R, Harkey D, Edwards J: The Essential Distributed Objects Survival Guide. New York, NY, Wiley & Sons, 1996

26. Bucci G, Cagnoni S, De Dominicis R: Integrating Content-based retrieval in a medical image reference database. Comput Med Imaging Graphics 20:231-241, 1996

27. Kabuka MR: The Digital Radiology Infrastructure (DRI). Internal Report, the Center of Medical Imaging and Medical Informatics, University of Miami

28. Ligier Y, Ratib O, Girard C, et al: Distributed file management for remote clinical image viewing stations. Proc SPIE 2711:475-482, 1996

29. Guttman A: *R*-Trees: A dynamic index structure for spatial searching. ACM-SIGMOD International Conference on Management of Data, June 1984, pp 45-57

30. Beckmann N, Kriegel HP, Schneider R, et al: The *R**-tree: An efficient and robust access method for points and rectangles. Proceedings of the 1990 ACM SIGMOD International Conference on Management of Data, Atlantic City, NJ, May 23-25, 1990, pp 322-331

31. El-Kwae EA, Kabuka MR: Efficient content-based indexing of large image databases. (in press)

32. Lee DL, Kim YM, Patel G: Efficient signature file methods for text retrieval. IEEE Trans Knowledge Data Engineering 7:423-435, 1995

33. Lee SY, Shan MK: Access methods of image databases. Int J Pattern Recognition Artificial Intelligence 4:27-44, 1990