
Building a medical multimedia database system to integrate clinical information: an application of high-performance computing and communications technology*†

By Henry J. Lowe, M.D.

Bruce G. Buchanan, Ph.D.

Gregory F. Cooper, M.D., Ph.D.

John K. Vries, M.D.

University of Pittsburgh
Pittsburgh, Pennsylvania 15261

The rapid growth of diagnostic-imaging technologies over the past two decades has dramatically increased the amount of nontextual data generated in clinical medicine. The architecture of traditional, text-oriented, clinical information systems has made the integration of digitized clinical images with the patient record problematic. Systems for the classification, retrieval, and integration of clinical images are in their infancy. Recent advances in high-performance computing, imaging, and networking technology now make it technologically and economically feasible to develop an integrated, multimedia, electronic patient record. As part of The National Library of Medicine's Biomedical Applications of High-Performance Computing and Communications program, we plan to develop Image Engine, a prototype microcomputer-based system for the storage, retrieval, integration, and sharing of a wide range of clinically important digital images. Images stored in the Image Engine database will be indexed and organized using the Unified Medical Language System Metathesaurus and will be dynamically linked to data in a text-based, clinical information system. We will evaluate Image Engine by initially implementing it in three clinical domains (oncology, gastroenterology, and clinical pathology) at the University of Pittsburgh Medical Center.

BACKGROUND

Patient care generates a large amount of text-based data. The sheer volume of this clinical data and its increasing importance to health care activities has begun a trend toward the use of computerized clinical information systems [1]. Many of these systems are entirely or largely text oriented. However, text is only

one medium through which clinical information is recorded and communicated.

The rapid growth of diagnostic imaging technologies over the past two decades has dramatically increased the amount of nontextual data generated in clinical medicine. Imaging technologies are essential to the modern practice of clinical medicine [2]. Traditional radiological and nuclear medicine images are now complemented by computerized tomography (CT), magnetic resonance imaging (MRI), positron emission tomography (PET), ultrasonography, and endoscopy-generated image data. Health care providers use such images routinely to make clinical decisions. Furthermore, many medical specialties such as pathology, dermatology, and ophthalmology generate a large number of clinically important images. Though reports on these images are added to the

* Based on a presentation at the Medical Library Association's Ninety-Fourth Annual Meeting postconference symposium, "Building the National Health Information Infrastructure: The Role of High-Performance Computing and Communications," San Antonio, Texas, May 19, 1994.

† This project is supported by Biomedical Applications of High-Performance Computing and Communications (BAA/HPCC) Contract no. NO1-LM-4-3507 from the National Library of Medicine.

patient record, the images themselves are usually difficult for the clinician to access and often impossible to integrate with other relevant clinical data [3]. Systems for the classification, retrieval, and integration of clinical images are in their infancy [4]. Traditional picture archival and communications systems (PACS) [5-7] are generally expensive, monolithic solutions that serve primarily the needs of radiologists and are often not well integrated with the patient record. The new clinical imaging technologies demand innovative medical image database models that can integrate all patient data. Such systems may improve the quality of patient care [8-9], increase the patient's involvement in clinical decision making [10-11], and produce significant new medical knowledge [12].

The architecture of traditional, text-oriented, clinical information systems has made the integration of digitized clinical images with the patient record problematic. This lack of integration leads to a fragmentation of the patient record, which hinders the physician's ability to synthesize all the data relevant to clinical decision making. In part, this partitioning of textual and nontextual patient data reflects the technological heritage of the traditional patient record. Until recently, the technology required to integrate textual and image-based clinical information was prohibitively expensive or nonexistent. Recent advances in high-performance computer, imaging, and networking technology now make it technologically and economically feasible to create a truly integrated, multimedia, electronic patient record linking digitized clinical images from a wide variety of sources with the traditional, text-based, medical record.

BEYOND PACS

Although PACS [13-14] have been an active area of research and development for almost twenty years [15-16], much of that work has been within the domain of radiology and has focused on expensive institutional systems [17]. Some involved with PACS research feel that if the fundamental idea is to prosper, it must expand its domain from radiology and radiologist to a broad range of clinically important images and the majority of physicians practicing both inside and outside the hospital environment [18]. In addition, given the current fiscal outlook for health care delivery, we must seek ways to reduce the implementation costs of integrated clinical image delivery systems.

The integration of clinical images with the textual clinical record has been an area of increasing interest over the past few years. Although a number of integrated PACS systems are described in the literature [19-20], one of the most innovative implementations of integration has been in the Department of Veterans Affairs (VA) Decentralized Hospital Computer Pro-

gram (DHCP) [21-23]. This system, developed at the Washington Information Systems Center of the VA, is a "distributed imaging system that provides image management and communications functionality as an integral part of its existing integrated hospital information system" [24]. The DHCP is an object-oriented system based around the MUMPS VA File Manager database using UNIX/X Windows or Microsoft Windows workstations and multiple image servers. It goes beyond the traditional PACS radiology model by "handling a variety of medical images including cardiology studies, microscopic pathology slides and endoscopic examinations" [25].

There is currently a confluence of developments that makes it technologically and financially feasible to implement institutional systems integrating digitized clinical images from a wide variety of sources with the traditional medical record. These innovations include inexpensive microcomputer workstations with the processing power, memory, and display characteristics necessary to permit real-time decompression and display of high-resolution digital still images and digital video. Image compression technology is now widely and inexpensively available, and international compression standards are emerging [26]. Digital video technology is now an integrated part of many microcomputer operating systems and permits us for the first time to display real-time, on-screen, color-video sequences using widely available and affordable computers. Networking advances increasingly supply the necessary bandwidth for institutional transport of compressed digital images.

Image compression will be an important enabling technology in the development of integrated, multimedia clinical information systems. Even with current advances in affordable computer hardware, the large size of uncompressed images would preclude workable, large-scale systems. For example, one minute of digital video (480 by 640 pixels, 24-bit color at 30 frames per second) requires 1.5 gigabytes of storage, while a single 480 by 640 pixel, 24-bit color still image is approximately one megabyte in size. Clinical information systems working with uncompressed data of this scale would rapidly fill even today's gigabyte storage devices and jam most institutional networks.

IMAGE ENGINE

Our Biomedical Applications of High-Performance Computing and Communications (HPCC) contract from the National Library of Medicine (NLM) will allow us to further develop Image Engine, a prototype microcomputer-based system for the storage, retrieval, integration, and sharing of a wide range of clinically important digital images [27]. Images stored in the Image Engine database will be indexed and or-

ganized using the Unified Medical Language System Metathesaurus and will be dynamically linked to data in a text-based, clinical information system. We will evaluate Image Engine by initially implementing it in three clinical domains at the University of Pittsburgh Medical Center (UPMC): oncology, gastroenterology, and clinical pathology.

The Image Engine System architecture consists of three layers: the Server Layer, the Object Database Layer and the Image Browser Layer. In addition, Image Engine will make use of two network-based services: the Medical ARchival System (MARS) clinical information system and the Probabilistic Indexing (Pindex) server.

THE MARS CLINICAL INFORMATION SYSTEM

MARS is a text-based, clinical information system developed by Dr. John Vries and Russell Yount at UPMC [28]. MARS has been in operation at UPMC for five years. It contains approximately 3.2 million whole-text, word-indexed, clinical records. These document records contain the full text of patient histories and physicals; operative and procedure notes from multiple clinical specialties; discharge summaries; laboratory results; and reports from the pharmacy, microbiology, pathology, and radiology departments. The data stored on MARS includes 80% of all information generated at UPMC and all information produced by the Central Transcription Service for UPMC's hospitals and outpatient clinics. More than 2,500 registered users retrieve an average of 5,000 reports each day, and these numbers are growing steadily. MARS is available twenty-four hours per day. It is anticipated that by 1995, MARS will capture most of the clinical information generated at UPMC.

THE PINDEX INDEXING SYSTEM

We have developed a system called Pindex that takes as input a string of free text and returns an associated list of Medical Subject Heading (MeSH) terms that are each annotated with a probability of relevance. The development of Pindex has been supported in large part by the NLM UMLS project. We are modifying and extending Pindex for the task of indexing images. We plan to use this modified version of Pindex to index medical images based on their free-text descriptions and to assist users retrieving images given free-text input from the user about the type of images that are desired.

Pindex works as follows. A simple parser, P , converts free text input into a set of word phrases, which we denote as S . Pindex has a large table, T , that associates word phrases with MeSH terms. Each association between a phrase and a term has an attached

probability. Let U denote the MeSH terms that are associated with one or more phrases in S . We attach to each term in U the maximum probability it has in association with any phrase in S (as given by table T).

If the free text is a description of an image, then the terms in U can be used to index the image. If the free text is a user query, then we can sort the terms in U in descending order of their attached probabilities. The user can choose terms from this sorted list to construct a search expression for images of interest.

Table T is constructed as follows. A large group of MEDLINE articles is used as the training set. For each MEDLINE article, A_i , we apply parser P to the text in the title and abstract of the article to derive a set of word phrases, S_i . Let U_i denote the main MeSH headings assigned to article A_i in MEDLINE. In A_i , we view each phrase in S_i as co-occurring once with each MeSH term in U_i . We update table T to reflect the phrase-term co-occurrences for article A_i and phrase occurrences in article A_i . After all articles are processed, table T contains the number of occurrences of each phrase that was encountered in the MEDLINE training articles and the number of co-occurrences of each phrase-term pair encountered. From these statistics, we compute the probability of a MeSH term, u , given a phrase, s , as the total number of co-occurrences of s and u divided by the total number of occurrences of s . To avoid misleading probability estimates due to small sample sizes, we require that the value of s be greater than a minimum threshold.

The terms in table T need not be limited to MeSH terms. In general, we can construct the table using any training data that contain free-text descriptions that are annotated with controlled vocabulary terms. Thus, for example, in addition to using MEDLINE training data, one could use training data based on surgical pathology reports in the MARS system for which online SNOMED codes are available. Our overall focus will be on using training data that contain terms within the UMLS Metathesaurus.

THE IMAGE ENGINE SERVER LAYER

The Image Engine Server Layer will consist of a dedicated server computer and a number of gigabyte range hard disks connected to UPMC's high-speed data network. These disks will store digitized, compressed clinical images. Still digital images will be stored in the PICT format and digital video images as Quick-time files. PICT files will be compressed using the International Standards Organization's (ISO) Joint Photographic Experts Group (JPEG) still image compression algorithm [29-30]. Digital video files will be compressed with the proposed ISO Motion Picture Experts Group (MPEG) video compression scheme [31]. Storing images in these widely supported and well-

documented formats will facilitate cross-platform development and image sharing. We plan to support image translation between a number of standard image interchange formats.

One of the issues we will address with image domain experts is defining the optimal compression characteristics for different clinical image types. Images will be automatically compressed as they are added to the server database and decompressed on arrival at a workstation running the Image Engine Browser software. Image compression reduces disk storage requirements and reduces transfer time across the network. While the ISO standard JPEG is lossy at compression ratios above 2:1, it is often capable of compressing a wide range of still image types at compression ratios of 10:1 to 24:1 without detectable loss of image quality. (Lossless image compression means that although the storage requirements of an image are reduced and the image is therefore compressed, the data set for that image remains unchanged. Lossy image compression, on the other hand, means that in the process of compressing an image, some of the original image data set is irreversibly lost.) For example, formal clinical evaluation of JPEG compressed chest x-rays found a compression ratio of 20:1 or less to be acceptable [32]. The optimal degree of compression applied to any image type is highly dependent upon the nature of the image [33] and will have to be determined over time in consultation with image domain experts.

THE IMAGE ENGINE OBJECT DATABASE LAYER

The Object Database Layer uses an object-oriented database model to represent the images stored on the Image Engine Server [34]. This may be viewed as a virtual database, in that it integrates data stored both locally and (by reference) on other systems [35]. This approach has a number of advantages, including avoiding storage replication and problems with data version inconsistencies. This model also allows for future integration with other clinical information systems.

The Image Engine Object Database will contain object-oriented representations of the images stored on the server. This database structure will support multiple independent index files. Each image object record will have a unique identifier. Object property values will be indexed as either text (inverted wordstem index based upon the Porter Algorithm [36]) or data (indexed on full property value). In addition, we are interested in experimenting with data by reference models (virtual data existing in other database systems). Multiple indexes based upon property values will allow for rapid database searching on many independent criteria.

Figure 1 outlines the major elements of the Image Engine Database and its relationship to MARS and Pindex. The image itself is stored in the server layer and has an object representation in the Image Engine Object Database. Each image object has a set of property values. For simplicity, only some of the image objects property slots (e.g., Image Reference, Patient ID) are shown. All image objects in the database will be indexed by multiple property keys in the Image Engine Index. These property keys point to both the Image Object and the image itself. Thus the index can be used to find a specific image or a set of images meeting some specification. Property values can also point to data stored on MARS (i.e., data by reference). For example, the Patient ID property points implicitly to all the MARS documents for that patient. Combinations of properties can define subsets of MARS data; for example, all pathology reports for a given patient ID between date X and date Y. Pindex is shown in this figure in its role as an automatic indexing tool, taking document input from MARS and sending Metathesaurus terms to the Object Database. However, Pindex could also aid data retrieval by taking text from a user query and returning suggested Metathesaurus terms for use in image retrieval.

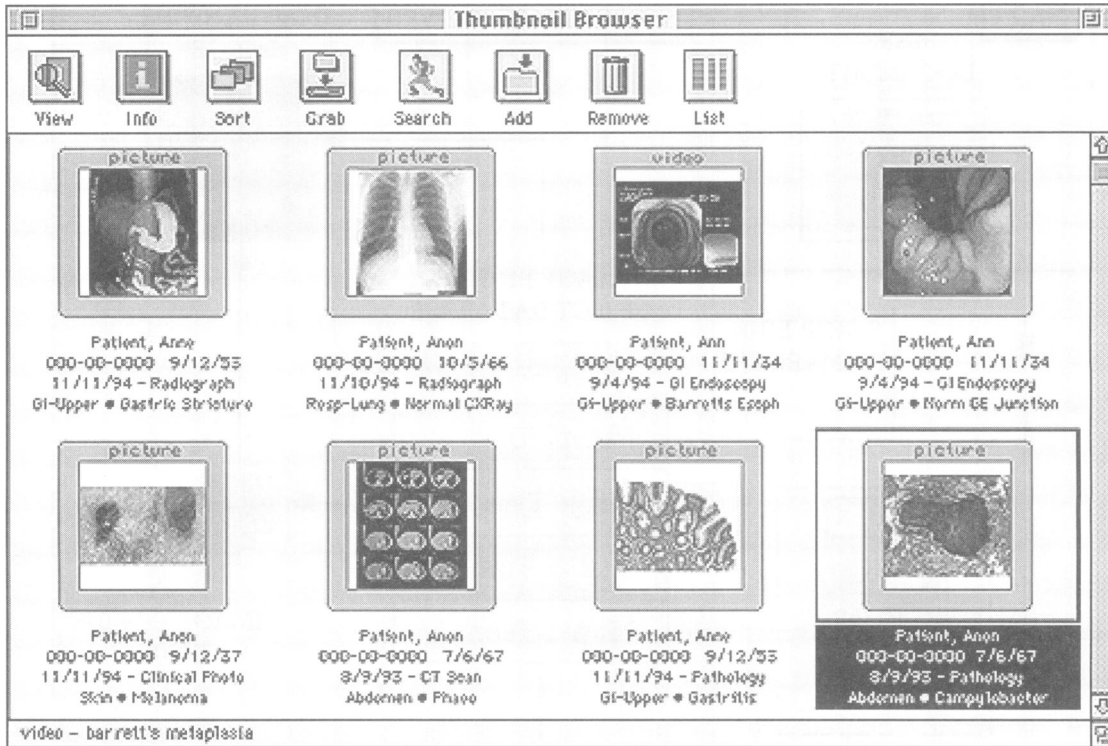
THE IMAGE ENGINE BROWSER LAYER

The Image Engine Browser (Figure 2) is a client application that allows one to interact with the Object Database through an easy-to-use graphical user interface. The Browser communicates with the Object Database over the network and uses a client-server model to request and retrieve image and patient data. Currently the project's workstations are based on Apple Power Macintosh 7100/66 RISC systems with twenty-four megabytes of RAM and twenty-inch, 1,152 by 870 pixel resolution color monitors using accelerated, twenty-four-bit color video cards.

The user may search for image subsets using combinations of image properties, including Metathesaurus terms. Retrieved subset summaries of images can be viewed as either a scrollable list of thumbnail images (100 by 100 pixel scaled, 24-bit color images) or a text list of object identifiers (image name, type, patient ID, and so forth). Images can be selected and viewed at full or scaled size on the computer display in resizable, scrollable windows. Multiple images can be viewed simultaneously. Image information text can be viewed simultaneously with images. Retrieved image sets can be sorted and displayed on a number of criteria.

Digital video images can be displayed on screen and controlled with videotape-like features (e.g., reverse, fast-forward, still frame, and reverse/forward frame capabilities). Users can convert digital video frames to digital still images. Given the large size of

Figure 2
Image Engine Thumbnail Browser



The UPMC data network consists of a basic backbone of three Fiber Digital Device Interface (FDDI) rings using fiber-optic cable with fiber-optic limbs to collapsed Ethernet backbones in each of the constituent hospitals and clinics of the medical center. Fiber-optic cable now extends to most hospital and clinic floors with standard 10BaseT cabling to individual workstations. The current average data transfer rate is ten megabits per second, increasing to a maximum of 100 megabits per second in the near future, with the installation of Level-5 cabling. The current supported routing protocols include TCP/IP, Novell's Internetwork Packet Exchange (IPX), and AppleTalk. Current bridging protocols include Digital Equipment's Local Area Transport (LAT) and Local Area Storage Transport (LAST), and Maintenance Operations Protocol (MOP).

DYNAMIC LINKS TO A TEXT-BASED CLINICAL INFORMATION SYSTEM

The MARS database system supports a scripting language, called ESP, which we will use to implement

dynamic links between clinical images in the Image Engine database and associated patient information in the MARS database. Using ESP, an application on the UPMC data network, can establish a TCP/IP connection to the MARS ESP server and transmit MARS queries written in the ESP language. MARS patient data retrieved by an ESP query is returned over the network to the requesting application in a format defined in the ESP query script. ESP, therefore, allows an application to become a MARS client.

The MARS ESP query language is still under development, and we anticipate that our work on dynamically linking Image Engine objects to MARS may suggest new features and enhancements to the ESP language and server. In particular, we are interested in exploring the use of "shortcut" ESP queries that specify exactly which documents should be retrieved. This will give us the capability of storing sparse ESP queries in Image Object property slots as data by reference. If this approach works, we could expand the Image Object Database to contain image objects that know how to retrieve information about themselves from other clinical information systems without any

user interactions. Such objects could "stack" potential data items from MARS, based upon their past experience with particular users or past sequences of actions leading to data retrieval. For example, such an object could anticipate that if a user views a colonoscopy report and associated images she may also request associated pathology images and reports. Depending upon user preference, these "stacked" images and reports might be automatically retrieved along with the explicitly requested data.

We are eager to ensure that the dynamic links between Image Engine and MARS are implemented in a way that does not preclude Image Engine linking in the future to other important clinical information systems. One method of ensuring this independence may be by defining a communications layer between Image Engine and the other databases.

CLINICAL APPLICATIONS

As part of our HPCC contract, we plan to install and evaluate Image Engine workstations in three clinical environments at UPMC—clinical pathology, gastroenterology, and medical oncology.

In clinical pathology, we will focus on issues involved in digitizing, compressing, indexing, storing, and retrieving both gross and microscopic pathology images. Clinicians in the gastroenterology and oncology test sites will identify pathology specimens (for example, biopsies obtained during gastrointestinal endoscopy or diagnostic oncology procedures) from patients entered into the Image Engine database. These pathology specimens will be digitized, indexed, and added to the database. Image Engine will automatically integrate these pathology images with other digitized images (such as endoscopy and radiology) and the text-based clinical record for that patient.

In gastroenterology, we will work with clinicians specializing in fiber-optic endoscopy of the gastrointestinal and biliary tracts. These domains will involve digital still and video images, pathology images, and radiological imaging studies (including MRI, CT, and ultrasound). For patients entered in the Image Engine database, it will be possible for clinicians to selectively view and manipulate integrated image and textual data.

In medical oncology, we will explore how one manages and integrates the wide range of images (including radiology, MRI, CT, pathology, and clinical photography) that are used in the diagnosis, staging, and treatment of patients with solid tumors. Clinicians will be able to rapidly retrieve both image and textual data for patients entered into the Image Engine database.

Image Engine should be useful in many other image-intensive clinical domains such as dermatology

[37] and ophthalmology [38]. In addition, as the number and variety of images in the database increases, it has the potential to become a valuable educational and research resource. For example, it would be possible for one to retrieve and view a set of pathology, dermatology, ophthalmology, endoscopy, and radiology images for a given disease entity. Alternatively, one could retrieve selected images from a population of patients with certain characteristics.

It has been estimated that medical imaging databases acquire in excess of one terabyte of information per year in a major hospital [39]. Given the large size of even compressed digital images, we have chosen to limit our initial HPCC evaluation domains to a relatively small population of clinicians and patients. Our goal in this project is not to implement a very large-scale, hospitalwide, image database system during this three-year project. Also, Image Engine is not intended to compete with or replace PACS, which we see as continuing to develop as a radiological support system. Instead, we plan to focus on the technical and clinical issues involved in creating a potentially portable system that could be scaled to handle the image storage, retrieval, and sharing needs of clinicians, with an emphasis on integrating a wide range of clinically important images with the text-based patient record using relatively inexpensive high-performance computers and networking technology.

REFERENCES

1. COLLEN MF. A brief historical overview of hospital information system (HIS) evolution in the United States. *Int J Biomed Comput* 1991 Dec;29(3/4):169-89.
2. HUANG HK, ABERLE DR, LUFKIN R, GRANT EG ET AL. Advances in medical imaging. *Ann Intern Med* 1990 Feb 1; 112(3):203-20.
3. MUN SK. Image management and communications (IMAC) in 1991. In: Heshiki A, Mun SK, eds. *The Second International Conference on Image Management and Communications (IMAC) in Patient Care: new technologies for better patient care*, April 10-13, 1991, Kyoto, Japan. Los Alamitos, CA: IEEE Computer Society Press, 1991:6-15.
4. GROSKY WI, MEHROTRA R. *Image database management*. Los Alamitos, CA: IEEE Computer Society Press, 1989:22: 7-8.
5. GUR D, FUHRMAN CR, THAETE FL. Requirements for PACS: users' perspective. *Radiographics* 1993 Mar;13(2):457-60.
6. RATIB O. Current views on the functionalities of PACS. *Int J Biomed Comput* 1992 May;30(3/4):193-9.
7. CHOPLIN RH, BOEHME JM II, MAYNARD CD. Picture archiving and communication systems: an overview. *Radiographics* 1992 Jan;12(1):127-9.
8. GUR D, STRAUB WH, LIEBERMAN RH, GENNARI RC. Clinicians' access to diagnostic imaging information at an academic center: perceived impact on patient management. *AJR Am Roentgenol* 1992 Apr;158(4):893-6.
9. DAYHOFF RE, MALONEY DL. An integrated multidepartmental hospital imaging system: usage of data across spe-

- cialties. In: Frisse ME, ed. Sixteenth Annual Symposium on Computer Applications in Medical Care: supporting collaboration. New York: McGraw-Hill, 1993:30-4.
10. ENDE J, KAZIS L, ASH A, MOSKOWITZ MA. Measuring patients' desire for autonomy: decision making and information-seeking preferences among medical patients. *J Gen Intern Med* 1989 Jan/Feb;4(1):23-30.
 11. BEISECKER AE, BEISECKER TD. Patient information-seeking behaviors when communicating with doctors. *Med Care* 1990 Jan;28(1):19-28.
 12. ZINK S, JAFFE CC. Medical imaging databases: National Institutes of Health workshop. *Invest Radiol* 1993 Apr;28(4):366-72.
 13. GROSKY, op. cit.
 14. GUR, Requirements for PACS.
 15. STECKEL RJ. Daily x-ray rounds in a large teaching hospital using high-resolution closed-circuit television. *Radiology* 1972 Nov;105(2):319-21.
 16. WEBBER MM, WILK S, PIRRUCCELLO R, ALKEN J. Telecommunication of images in the practice of diagnostic radiology. *Radiology* 1973 Oct;109(1):71-4.
 17. HUANG HK. PACS research and development: a review and perspective. In: Huang HK, Ratib O, Bakker A, Witte G, eds. Picture archiving and communications systems (PACS). Berlin: Springer-Verlag, 1991:1-7. NATO ASI series. Series F. Computer and systems sciences;74.
 18. SARANUMMI N, AKISADA M, IRIE G, KIUREE A ET AL. User requirements and standards for PACS: 2d Japan-Nordic PACS Symposium. *Comput Methods Programs Biomed* 1992 May;37(4):237-45.
 19. HUANG HK, TAIRA RK, LOU SL, WONG AW ET AL. Implementation of a large-scale picture archiving and communication system. *Comput Med Imaging Graph* 1993 Jan/Feb;17(1):1-11.
 20. HORII SC, MUN SK, ELLIOTT LP, LEVINE B ET AL. PACS clinical experience at Georgetown University. *Int J Biomed Comput* 1992 May;30(3/4):275-80.
 21. DAYHOFF RE, MALONEY DL, MAJURSKI WJ. VA's Integrated Imaging System on three platforms. In: Frisse ME, ed. Sixteenth Annual Symposium on Computer Applications in Medical Care: supporting collaboration. New York: McGraw-Hill, 1993:791-2.
 22. DAYHOFF RE, MALONEY DL, KENNEY TJ, FLETCHER RD. Providing an integrated clinical data view in a hospital information system that manages multimedia data. In: Clayton PD, ed. Fifteenth Annual Symposium on Computer Applications in Medical Care: assessing the value of medical informatics. New York: McGraw-Hill, 1992:501-5.
 23. DAYHOFF RE, MALONEY DL, KUZMAK PM, SHEPARD BM. Integrating medical images into hospital information systems. *J Digit Imaging* 1991 May;4(2):87-93.
 24. DAYHOFF, Integrating medical images.
 25. *IBID.*
 26. RABBANI M, JONES PW. Image compression techniques for medical diagnostic systems. *J Digit Imaging* 1991 May;4(2):65-78.
 27. LOWE HJ. Image Engine: an object-oriented multimedia database for storing, retrieving and sharing medical images and text. In: Safran C, ed. Seventeenth Annual Symposium on Computer Applications in Medical Care: patient-centered computing. New York: McGraw-Hill, 1994:839-43.
 28. YOUNT RJ, VRIES JK, COUNCILL CD. The Medical Archival System: an information retrieval system based on distributed parallel processing. *Inf Proc Man* 1991;27(4):379-89.
 29. WALLACE GK. The JPEG still picture compression standard. *IEEE Trans Consum Electron* 1992 Feb;38(1):xviii-xxxiv.
 30. KAJIWARA K. JPEG compression for PACS. *Comp Methods Programs Biomed* 1992 May;37(4):343-51.
 31. LE GALL D. MPEG: a video compression standard for multimedia applications. *Commun ACM* 1991 Apr;34(4):46-58.
 32. ISHIGAKI T, SAKUMA S, IKEDA M, ITOH Y ET AL. Clinical evaluation of irreversible image compression: analysis of chest imaging with computed radiography. *Radiology* 1990 Jun;175(3):739-43.
 33. NELSON M. Lossy graphics compression. In: Nelson M. The data compression book: featuring fast, efficient data compression techniques in C. Redwood City, CA: M&T Books, 1991:347-408.
 34. HEATHFIELD H, ARMSTRONG J, KIRKHAM N. Object-oriented design and programming in medical decision support. *Comput Methods Programs Biomed* 1991 Dec;36(4):239-51.
 35. MARGULIES D, MCCALLIE D JR., ELKOWITZ A, RIBITZKY R ET AL. An integrated hospital information system at Children's Hospital. In: Miller RA, ed. Fourteenth Annual Symposium on Computer Applications in Medical Care: standards in medical informatics. Los Alamitos, CA: IEEE Computer Society Press, 1991:699-703.
 36. PORTER MF. An algorithm for suffix stripping. *Program* 1980;14(3):130-7.
 37. STOECKER WV, MOSS RH. Digital imaging in dermatology. *Comp Med Imaging Graph* 1992 May/June;16(3):145-50.
 38. HOLM A, HOLM O. Digitizing 35 mm colour slides for computerized general image handling in ophthalmology. *Acta Ophthalmol* 1991 Oct;69(5):611-7.
 39. KIM Y, PARK HW, HAYNOR DR. Requirements for PACS workstations. In: The Second International Conference on Image Management and Communications (IMAC) in Patient Care: new technologies for better patient care, April 10-13, 1991, Kyoto, Japan. Los Alamitos, CA: IEEE Computer Society Press, 1991:36-41.

Received August 1994; accepted August 1994