Recommendations for Image Prefetch or Film Digitization Strategy Based on an Analysis of an Historic Radiology Image Database

Eliot L. Siegel and Bruce I. Reiner

Picture archiving and communications systems (PACS) utilize short- and long-term storage to provide both rapid retrieval and large storage capacity. Owing to the practical limitations imposed on the size of the much faster short-term storage, it is important to use an effective algorithm in the retrieval of comparison images from long to short-term storage. A strategy must be used to maximize the likelihood that the relevant historic images have been previously retrieved into short-term memory. Data were collected with a database consisting of 754 consecutive examinations and 7,723 associated historic studies. The most frequent number of previous examinations was zero (11% of patients). In 45% of cases, no previous matching examinations had been performed. Two basic strategies of image retrieval were evaluated. The first algorithm retrieved the last n studies in chronological order. The second strategy tested was retrieval based on a defined interval of time. This strategy was found to be less efficient. By using the former strategy, a 91% success rate (defined as successful retrieval of the previous matching exam) was achieved with retrieval of only 30% of the prior exams. The second approach required retrieval of 70% of the prior exams to achieve a 90% success rate for the previous matching exam. However, the data from this latter strategy suggest that examinations are often ordered in clusters. Thus, there was found to be a 72% likelihood that a previous matching exam, if present, would be available on a PACS after only 1 week of operation, and an 80% chance after only 1 month of operation. The data therefore suggest that digitization of film in a new PACS environment might not be necessary owing to the relatively short period of time required to populate the database with historical studies.

Copyright © 1998 by W.B. Saunders Company

KEY WORDS: digital, computer, radiology, imaging, algorithm, workstation, efficiency, archival, storage, database, digitization

ONE OF THE MOST important lessons taught to radiology residents is the inestimable value of reviewing previous examinations. Most

radiologists would be uncomfortable interpreting studies without previous examinations and prior radiology reports. However, despite the major strides that have been made in both archival size and network performance, current picture archiving and communication systems (PACS) do not typically have all previous studies performed on a given patient in short-term storage. This is true for PACS with a distributed architecture and limited local storage and for those using a central architecture, including those systems with storage capacities of more than 256 gigabytes. Ideally, the PACS should be able to retrieve historic images for comparison "on demand" rapidly.¹

Unfortunately, this may not be possible because of insufficient time for retrieval of the pertinent historical studies, slow retrieval and transfer speed of the archival devices, the level of patient volume, network throughput, the capacity of the faster, short-term storage devices, and other factors. Virtually all PACS have multitiered archives that include different types of storage devices.²⁻⁴ Capacities have risen and costs have decreased, but this remains the situation today. It was more pronounced in earlier systems. The huge amounts of image data to be archived require large capacity storage devices which retrieve images more slowly than the rapid short term devices used to support workstations. These faster devices are more expensive per megabyte. The trade-off of data retrieval speed and cost per megabyte also applies to intermediate term storage devices. The time required for retrieval from medium- or long-term storage devices can be slower by more than an order of magnitude than is available when using short-term storage. This large difference in access times and the other problems in image retrieval make it critical to utilize a variety of algorithms to increase the likelihood that current and historic images are available in short-term storage. Algorithms can be used to determine which prior studies to transfer from long- to short-term storage and which studies to purge from short-term storage when it is unlikely that they will be utilized. The need for these strategies is particularly great in facilities in which

From the Department of Diagnostic Radiology, University of Maryland School of Medicine, Baltimore VA Medical Center, Baltimore, MD.

Address reprint requests to Eliot L. Siegel, MD, Department of Diagnostic Radiology, Baltimore VA Medical Center, 10 North Greene Street, Baltimore, MD 21201.

*Copyright © 1998 by W.B. Saunders Company
0897-1889/98/1102-0002\$8.00/0*

there is relatively limited short-term memory in comparison to the volume of studies performed.⁵ Even though a number of PACS use prefetching, there is little in the academic literature about it.

There are a number of "events" that can potentially be used to "trigger" transfers from long- to short-term storage ("prefetches"). These include the scheduling of a patient for an outpatient appointment, the admission of a patient to the hospital as an inpatient, the transfer of a patient within a hospital (such as to the intensive care unit [ICU]), the ordering of a new study, the discharge of a patient from the hospital, or the end of a scheduled clinic appointment. Any one or a combination of them, and such variables as patient diagnosis, previous studies performed, patient age, and patient location could be used to determine which images to transfer to and from short-term storage.

This study was undertaken to evaluate two straightforward prefetching methods for possible implementation and clinical evaluation in our PACS, based on the database acquired from 5 years of operation.

MATERIALS AND METHODS

Consecutive studies performed on both inpatients and outpatients during 1 week were collected prospectively at the Baltimore VA Medical Center before the installation of a PACS. For each of the 754 examinations performed, data concerning all of the previous studies performed on that patient were obtained from either the hospital information system or the patient film jacket and entered into a Microsoft Access (Microsoft Corp, Redmond, WA) database. The radiology information system database had approximately two years of data at that time, and the film jacket listing extended back further to six years. Certain

Table 1. Information in the Historical Exam Database

1. Patient identification
2. Examination performed
3. Date of exam
4. Total number of historical exams
5. Date of oldest exam
6. Date of most recent exam
7. Date of last identical exam
8. Chronological order of last identical exam
9. Date of the second to the last identical study
10. Chronological order of the second last identical study
11. Date of oldest identical study
12. Chronological order of the oldest identical study
13. Most recent related exam
14. Date of most recent related exam
15. Chronological order of most recent related exam

Note: These fifteen data elements were recorded for each of the 754 consecutive radiographic examinations included in the study.

Table 2. Related Examinations

CT Abdomen	CT Head	IVP
Abdominal "plain" images	MRI brain	Renal ultrasound
RUQ ultrasound	CT posterior fossa	Abdominal "plain" images
Renal ultrasound	Skull radiograph	CT abdomen
CT pelvis		Retrograde urethro-cystogram

Note: For each examination, a list of "related" studies was compiled. The table above lists examples of related examinations for a CT of the abdomen, CT of the head, and an intravenous pyelogram (IVP).

special studies such as those performed for disability compensation and asbestos exposure were retained for longer periods. The total number of historic examinations entered was 7,723. Thus, the average number of previous radiology or nuclear medicine examinations for a given current study was 10.2 examinations. The specific information recorded for each examination is listed in Table 1.

The most recent and next to most recent identical or comparable (for example, two view chest and single view of the chest) exam were identified for each current study. Also, a radiologist examined all prior exams to identify the "most recent related exam," if any. For example, a right upper quadrant sonogram was deemed to be a related examination to a computed tomography (CT) scan of the abdomen because of the overlap in the area of the body examined and because a CT study of the abdomen is often performed following an abnormal sonogram. Other examples of the "most recent related examination" for a CT of the abdomen or head and intravenous urogram (IVP) are listed in Table 2.

A consensus of fourteen faculty members addressed which prior reports and which historic study images should be available when interpreting a current study. This was not an utterly simple task. The summary presented in Table 3 with some caveats was used to define three levels of retrieval requirements. Level 1, the simplest requirement, was retrieval of only the most recent identical or comparable study. Level 2 was the retrieval of the two most recent identical or comparable studies. Level 3 was the same as Level 2 with the addition of the most recent "related study." The criteria for any of the three levels were considered satisfied either by providing the identified studies or if the required studies had never been done.

Table 3. Radiologists' Criteria for Successful Retrieval

Level	Criteria
Level 1	The previous matching (same modality and anatomic area) exam was retrieved
Level 2	The most recent two matching exams were retrieved
Level 3	The most recent two matching exams and the most recent "related imaging study" was retrieved

Note: For most examinations, the 14 University of Maryland faculty radiologists felt that only the previous matching examination and the old reports for abnormal studies were required, and only the report when the prior study was read as "normal." For portable chest radiographs including ICU studies, most radiologists required the previous two studies for comparison and the old reports.

Fig 1. Frequency distribution of historic examinations. There was a wide variation in the number of previous historic examinations. Eighty-eight patients (11%) had no previous studies. The largest number of prior examinations in the database was 124 (not illustrated).

Two basic prefetching methods were evaluated. Neither method discriminated on exam type or patient origin. The first was a retrieval of the last n examinations in chronological order. The second method tested was a retrieval of all studies performed within the past n months. In each test the database on 7,273 historic studies was examined for each of the 754 current studies to determine the percentage of successful retrievals that would have occurred for each of the three retrieval levels had the test method been in place.

RESULTS

Conventional radiographs accounted for 73% of the examinations performed. The most commonly performed study was a conventional chest radiograph, comprising 44% of the total number of studies. CT examinations constituted 18% of the total. The patient with the largest number of previous studies had 124 examinations. Approximately one-third of new studies had at least one previous "related exam" with the average location of the most recent "related exam" being the sixth prior study.

The number of patients with a given total

number of historic exams is shown in Fig 1. The most common number of previous examinations was zero; this occurred in only 11% of the patients. It is likely that this number would be higher in most medical facilities outside of the VA setting, as VA patients tend to be older, tend to have more complex disease processes, and tend to have the majority of their imaging studies performed within the VA system.

The first prefetching method evaluated was a retrieval of the last n cases of any type. The results of this strategy are summarized in Table 4 for each of the three levels of historic examination retrieval as previously defined. As the number of previous examinations (of any type) retrieved increases from 0 to 10, the percentage of the patient's entire record of previous examinations retrieved increases from 0 to 57%. Figures 2 and 3 illustrate that an algorithm to prefetch the previous four studies would retrieve, on average, 30% of the archive's prior examinations but would result in success

Table 4. Results of First Prefetching Method

Number Retrieved	Percentage of the Patient's Total Images Retrieved	Level 1 Criteria (%)	Level 2 Criteria (%)	Level 3 Criteria (%)
0	0	45	45	35
1	8.6	75	51	45
2	16	83	71	62
3	23	88	77	68
4	30	91	84	75
5	35	94	88	81
6	40	95	92	85
8	49	97	95	89
10	57	98	97	92

(criteria satisfied) in 91% of cases using the Level 1 criteria, 84% using the Level 2 criteria, and 75% using the Level 3 criteria. To achieve an approximately 90% success rate using the Level 2 criteria for success, approximately 37% of the patient's historic examinations would be need to be retrieved (on average). This method would achieve almost a 95% success rate using the Level 1 criteria and between an 81% and 85% success rate using the Level 3 criteria. To achieve a 90% success rate, 4 cases would have to be retrieved to meet the Level 1 criteria, 6 to meet the Level 2 criteria, and 9 to meet the Level 3 criteria.

The second prefetching method evaluated was retrieval based on a defined interval of time. The average interval from the current examination to the first historic study of any type was 48 days. The average number of days from the current study to the first previous identical examination was 99 days. To achieve a 90% success rate using the Level

1 criteria with the time interval method, all studies obtained during the past 6 months would be needed. To achieve a 90% success rate using the Level 2 criteria, the system would need to retrieve approximately 9 months of previous studies. Retrieval of the past one year of studies would be required to achieve a 90% success rate using the Level 3 criteria. Figure 4 shows the success rate using the time interval retrieval method.

DISCUSSION

Given the current requirements for radiologists and clinicians to maximize productivity, it is important to have both new and historic images rapidly available. In a film-based environment these requirements can be addressed by retrieving old film jackets in anticipation of a patient being seen in clinic, admission to an intensive care unit, or before being interpreted by a radiologist. Film based methods fail if the patient's studies are misplaced, out of the department or unavailable for another reason. The ability to view an imaging study simultaneously in different locations and a loss free image archive are advantages offered by PACS. Unfortunately, there are limitations on network speed and short-term storage space, even in a state-of-the-art digital department, that mandate historic imaging studies be retrieved in advance from long- to short-term memory when prefetching is possible. Retrieval "on demand" of relevant historic studies as new studies are ordered is not always possible owing to the performance limita-

Fig 2. Retrieval efficiency: Percentage of all studies retrieved (to pull the last N studies (n varies from 0-10)). This figure illustrates the percentage of a patient's total number of historic studies retrieved in order to pull the last n studies. For example, retrieval of the last 8 studies would result in retrieval of approximately 50% of the total number of historic studies performed on that patient.

Fig 3. Number of studies pulled versus success rate for Level 1, Level 2, and Level 3 criteria. For each of the criteria (Levels 1, 2, and 3), the likelihood that a retrieval of the last n studies will satisfy that criteria. Even with no studies retrieved, the Level 1 and 2 criteria are satisfied 45% of the time because no matching prior study existed in 45% of the cases.

tions. Slow retrieval rates from an optical jukebox or other long term storage device and decreased network transfer rates, especially at times of heavy system loads, are two types of common “bottle-necks.”

The first prefetching method evaluated was to retrieve the last n consecutive examinations. When only 30% of the average patient’s historic studies

were retrieved (4 studies), the likelihood that a previous matching exam was retrieved (Level 1) was more than 90%, and there was a more than an 83% chance that the previous two matching studies (Level 2) were made available. This is more efficient than the second prefetching method, which required the retrieval of 6 months of prior studies to achieve a 90% likelihood of retrieving the previous

Fig 4. Interval of historic exams retrieved for three criteria. For each of the three levels, the likelihood that a retrieval of studies during the last n weeks or months will satisfy those criteria. Even after 1 week, there is a 72% likelihood that a previous matching examination will be available on the system or that no previous matching exam exists.

matching examination (Level 1); the most recent 6 months would require retrieval of 70.5% of the patients' previous studies. Retrieving only the most recent study by the first method (less than 9% of the patient's image file) resulted in a 75% likelihood that it would match the new examination (or that there was no previous examination).

The first prefetching method of retrieving some number of the most recent historic studies was selected for clinical trial on the operational PACS based on this preliminary evaluation with our database of previous exams. On the one hand, the admittedly simple algorithms evaluated will operate without placing excessive load on the system. On the other hand, further refinements of the "prefetch" algorithm beyond these very basic two approaches would undoubtedly yield more effective retrievals. For example, a prefetch algorithm, which retrieves only the last study of any type and the most recent chest radiograph, achieves an 87% likelihood of retrieving the previous matching study. This is an improvement over the algorithm that retrieves the last two studies of any type (83%).

More intelligent retrieval algorithms that incorporate information such as the patient location, patient diagnosis, and types of previous examinations performed, would be very likely to provide even higher rates of success. A neural network or other predictive tool might be used in developing and refining these algorithms. Special routines to search further into a patient's old studies might be added for pediatric studies or mammograms.

Additionally, new and improved mechanisms to monitor image retrieval patterns should be incorporated into future PACS. These could be used to measure the frequency of image accesses after interpretation by the radiologist to determine the length of time images should remain in short-term storage until they are purged. Current PACS do not take full advantage of these tools to optimize the effectiveness of short-term storage. These intelligent caching strategies could even be applied to local video memory to enable even more rapid image retrieval.

One potential limitation of these data is the fact that the patient population in the VA Maryland Health Care System is not representative of the type of patients encountered in other settings such as a small community hospital. This suggests that it would be useful to perform an analysis of the radiology examination database at other medical centers to determine the optimal prefetching strategies at those institutions.

Should Old Studies be Digitized When a New PACS is Begun?

Digitizing some months of the most recent analog studies has certain appeal to make these studies available on-line as historic studies for comparison purposes. However, it is both very labor intensive and it is difficult to properly identify the patient, study and projection or data set for each image digitized. As a practical result, almost no new PACS begins by digitizing old files for these reasons, to avoid cost and because it is not entirely clear how much would be gained.

Although retrieval of historic studies by using the time interval strategy was found to be less efficient, the data from the second prefetching method, retrieving all recent cases back over some period of time, could be utilized to predict the number of months that a PACS would have to be operational to contain the previous one or two matching examinations and related examinations. The data suggest that after only one week of "filmless operation," there is a 72% likelihood that a previous matching examination will be found on the system, an 80% chance after one month, and a 96% likelihood after one year (or that no previous matching examination was performed). Therefore, a hybrid approach of viewing old images on film and new images on a workstation would only have to be in effect for a short number of weeks until the PACS image database was likely to contain previous one or two matching exams. This is consistent with the experience at most large PACS sites, suggesting that large-scale film digitization is unnecessary during the transition to filmless operation.

REFERENCES

1. Siegel EL, Diaconis JN, Pomerantz S, et al: Making Filmless Radiology Work. *J Digit Imaging* 8:151-155, 1995
2. Hindel R: The impact of image storage organization on the effectiveness of PACS. *J Digit Imaging* 3:205-210, 1990
3. Dwyer SJ III, Templeton AW, Martin NL, et al: The cost of managing digital diagnostic images. *Radiology* 144:313-318, 1982
4. Piqueras J, Carreno JC, Ovelheiro M, et al: Worklists, preloading and archiving strategies: 3 years of clinical experience in the Barcelona PACS. *Med Inf (Lond)* 19:123-128, 1994
5. Wong AW, Huang HK, Arenson RL, et al: Digital Archive System for Radiologic Images. *Radiographics* 14:1119-1126, 1994