Voice-Activated Retrieval of Mammography Reference Images

Henry A. Swett, Pradeep G. Mutalik, Vladimir P. Neklesa, Laura Horvath, Carol Lee, Joan Richter, Irena Tocino, and Paul R. Fisher

We undertook this project to integrate context sensitive computer-based educational and decision making aids into the film interpretation and reporting process, and to determine the clinical utility of this method as a guide for further system development. An image database of 347 digital mammography images was assembled and image features were coded. An interface was developed to a computerized speech recognition radiology reporting system which was modified to translate reported findings into database search terms. These observations were used to formulate database search strategies which not only retrieved similar cases from the image database, but also other cases that were related to the index case in different ways. The search results were organized into image sets intended to address common questions that arise during image interpretation. An evaluation of the clinical utility of this method was performed as a guide for further system development. We found that voice dictation of prototypical mammographic cases resulted in automatic retrieval of reference images. The retrieved images were organized into sets matching findings, diagnostic hypotheses, diagnosis, spectrum of findings or diagnoses, closest match to dictated case, or user specified parameters. Two mammographers graded the clinical utility of each form of system output. We concluded that case specific and problem specific image sets may be automatically generated from spoken case dictation. A potentially large number of retrieved images may be divided into subsets which anticipate common clinical problems. This automatic method of context sensitive image retrieval may provide a "continuous" form of education integrated into routine case interpretation.

KEY WORDS: Computer assisted decision making, mammography, computer aided diagnosis, radiology reporting systems, decision support systems

DECISION MAKING in radiology depends on visual pattern recognition and association of recognized features with remembered prior experiences. Traditionally, to develop and maintain competence radiologists have depended on a background of formal radiological training, continuing education through journal and textbook review, educational courses, and clinical experience. Computers have shown variable success at automatically recognizing abnormalities in a few specific areas such as pulmonary nodules, pneumothoraces, breast masses, and calcifications. We have been particularly interested in developing computer-based tools to reinforce the radiologist's pattern identification and classification functions and to integrate access to such tools into the natural process of image interpretation and reporting. The ICON system employed an expert system to "discuss" the significance of findings that were observed by the radiologist.[1] The IMAGE/ICON system coupled this kind of case discussion with display of images evoked by the image features under consideration.[2] This article presents the application of context sensitive image retrieval to problem solving in mammography in a computerized decision support system (DSS) called MAMMO/ICON. We have automated feature input by using a voice recognition reporting system to show the concept of making computer-based decision support an integral part of the image reporting process. We discuss the way that this system organizes retrieved images into "axes" of clinical relevance and present the relative clinical utility of these different image groupings as determined by two mammographers during a functional system evaluation.

MATERIALS AND METHODS

Database

We assembled an image database of mammography images consisting of 347 images from 188 cases. There were 64 examples of breast cancer of varying histologies, 6 non epithelial malignancies, 6 inflammatory lesions, and 112 benign entities including fibroadenomas, lipomas, cysts, benign calcifications, scars, and ruptured implants (Table 1). All malignancies were pathologically proven. Histological proof was not available for some benign entities. They were included in the database based on clinical follow-up or if they were considered benign by at least two radiologists.

The images were digitized with a laser film digitizer (Konica

From the Department of Diagnostic Radiology, Yale University School of Medicine, New Haven, CT; the Department of Diagnostic Radiology, St. Raphael's Hospital, New Haven, CT; and the Department of Diagnostic Radiology, Alleghany University Hospitals, MCP, Philadelphia, PA.

This research was supported in part by SBIR Grant 1R43CA58060.01 from NIH.

Address reprint requests to Henry A. Swett, MD, Department of Diagnostic Imaging, Yale University School of Medicine, PO Box 208042, New Haven, CT 06520-8042; or, 820 N. Chelan Ave, Wenatchee, WA 98807.

0897-1889/98/1102-0004$8.00/0

Table 1. MAMMO/ICON Image Database Composition

Histology	No. of Cases
Carcinoma (Various Histologies)	64
Malignancy (Non Epithelial)	6
Benign Entities	118
Fibroadenomas	31
Lipoma	3
Cysts	15
Benign Calcifications	8
Scar	8
Ruptured Implant	1
Abscess	1
Intraductal Papilloma	4
Radiation Mastitis	2
Sclerosing Adenosis	9
Other	36
Total	188

KFDR-S, 250 dpi 12 bits/pixel; 100 μm pixels; Wayne, NJ). Most cases included at least two orthogonal projections. A clinical summary was abstracted and recorded for each case as well as a synopsis of pathological findings when available. They were converted to PICT format. Most images were interpolated down to 640×8 bit resolution (64 dpi) if the radiologists determined that findings were faithfully displayed at this resolution. $64 \times 64 \times 4$ bit image token were also stored. Images were displayed on a 1024×768 gray scale monitor.

Image Description Scheme

We defined a detailed image description lexicon which included finding name, scalar values to quantify the character of each finding (ie, mass density), conspicuity, and typicality. Limited representation of relationships of findings was represented.

Software Architecture

System software was written in the C computer language. This included dedicated database management tools, image retrieval handlers, user interface, and communication with the report dictation system.

Reporting System

We employed a commercially available speech recognition radiology reporting system (VoiceRAD, Kurzweil Applied Intelligence Inc, Waltham, MA) to prepare standard written reports. This system translates "trigger phrases" into complete sentences which are concatenated into a formal radiology report. Each finding trigger was passed to the Macintosh-based (Apple Computer Inc, Coppertino, CA) DSS where it was translated into descriptors in the image description language using lookup tables. To initiate an analysis by the DSS, the radiologist proposes a diagnostic hypothesis that might explain the observed findings and then says the trigger phrase: "Analyze."

Image Clustering Method

A group of related findings, such as breast masses, may be thought of as existing in a three-dimensional image space. If one has a series of images of different carcinomas, for example, they might be grouped by sharpness of border definition. This represents one pathway through the "feature space" of mammography images. On the other hand, a series of masses with varying causes and similar border sharpness may be organized by diagnosis, another axis through the array of images. Similarly, masses may be grouped by density, density and diagnosis, density, diagnosis, and border characteristics, and so on. This conceptual multidimensional arrangement of image findings is the basis for MAMMO/ICONS' image "axes" representation (Fig 1).

The initial global set of retrieved images is subdivided into image axes (Fig 2) as follows:

1. *Hypothesis Axis.* Images that share features with the index case and are caused by the hypothesized diagnosis are included in this array (Fig 3). The examples most like the index case are presented first and become progressively dissimilar along the axis. This paradigm is maintained for each axis.
2. *Differential Diagnosis Axis.* The selection criteria are relaxed for this set of images (Fig 4). To appear on this array, images must share major features with the index case, but may have any cause. This array provides a visual differential diagnosis for the index case.
3. *Spectrum Axis.* For this array, all images must share a common origin to the proposed diagnosis but do not have to share similar features. For example, if a radiologist wishes to review the spectrum of findings that may be seen in lymphoma of the breast, they will be presented on this axis.
4. *Malignant Axis.* All malignancies, regardless of type, are displayed on this axis beginning with those most similar to the index case and becoming progressively dissimilar moving from left to right on the axis.
5. *User Defined Axis.* This axis allows findings to be added or subtracted from the index finding set, or a new set of findings and hypotheses supplied. Here, the radiologist may view examples of any specific finding or finding group.
6. *Best Match Axis.* All images in the initial image set are placed on this axis in order of similarity to the index case. The first images presented are most like the index case (Fig 5).

This system bases its primary search on major mammography findings (masses, and calcifications). If a mass and/or calcifications are present, the system uses morphological descriptors such as mass density, border characteristics, or calcification morphology as its primary entrance into the database. The search routines try to match all features first. After primary features have been sought, the system searches for secondary features such as skin thickening, nipple retraction, and so on. For each axis, given the relevant constraints, the system endeavors to match, in order: (1) all findings, (2) mass characteristics, (3) calcification characteristics, and (4) secondary findings. For each set of matches, the system then relaxes the indicated constraint. For example, if the index case is a high density sharply marginated mass, masses with increased density and sharp margins will be displayed first with progressive relaxation of margin sharpness followed by progressive relaxation of density characteristics.

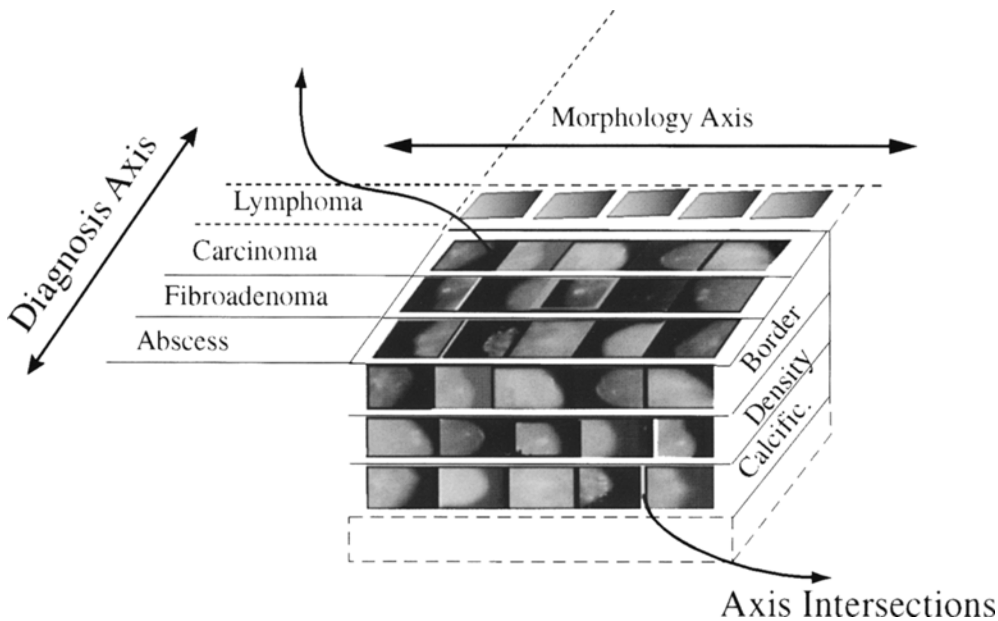

Fig 1. Axis distribution of image findings. Image findings may be distributed in space in many ways including diagnosis, radiographic findings, finding variations, clinical findings, and demographics. All images display. The image tokens represent all images retrieved in response to index case before they have been distributed into image arrays. This array of images may be large and is first displayed in random order. The full sized images are displayed by clicking any image token.

Evaluation of System Function and Clinical Utility

Before further database expansion and system development, a functional evaluation of the system was performed. Our goal was to verify that the system was retrieving images appropriately and setting up the image arrays as specified. We also wanted to find out which forms of image clustering were preferred by mammographers. We refer to this as clinical utility. We did not attempt to measure clinical effectiveness or accuracy at this stage of system development. To measure correct recall of images and proper placement in image arrays, we conducted a structured simulated clinical encounter with the system. Eight prototypical cases were submitted to MAMMO/ICON in compiled form and evaluated by two mammographers. Evaluators examined the image content of each axis and recorded their opinions about the global utility of each array as well as the individual images in each array. To gauge global utility, evaluators were asked to determine if the specified image set provided information that was potentially useful in establishing a diagnosis. This subjective determination was recorded on a 10 point scale (1 = no clinical value; 10 = high clinical value). To determine if the DSS was placing images in the correct arrays and sequencing them correctly, the evaluators inspected each individual image and indicated if the image belonged in the array and if it was placed in logical sequence in relation to the preceding and following images.

Eight prototypical index cases were used as the initial entry into the MAMMO/ICON system in this evaluation. The major findings represented are as follows:

1. Stellate mass
2. High density mass
3. Low density mass
4. Combined density mass
5. Radiolucent mass
6. Multiple masses
7. Clustered calcifications
8. Scattered calcifications.

Additionally, the evaluators submitted multiple variations of each case through "user defined" permutations.

RESULTS

The results of the functional evaluation are summarized in Table 2.

Differential Diagnosis Axis

The number of images retrieved ranged from 2 to 28, depending on the index case and the presence or absence of relevant images in the test database. Clinical utility ranged from 7 to 9 (average = 7.6). Four cases from all arrays were considered inappropriate for display on this axis because the evaluators felt they would not oridinarily be considered in the differential diagnosis of the lesion in question. Three cases of incorrectly coded findings were identified. One image was considered incorrectly sequenced in the array. This could occur if an important image was placed at the end (right end of axis) of the array, or if an irrelevant image was displayed early in the array.

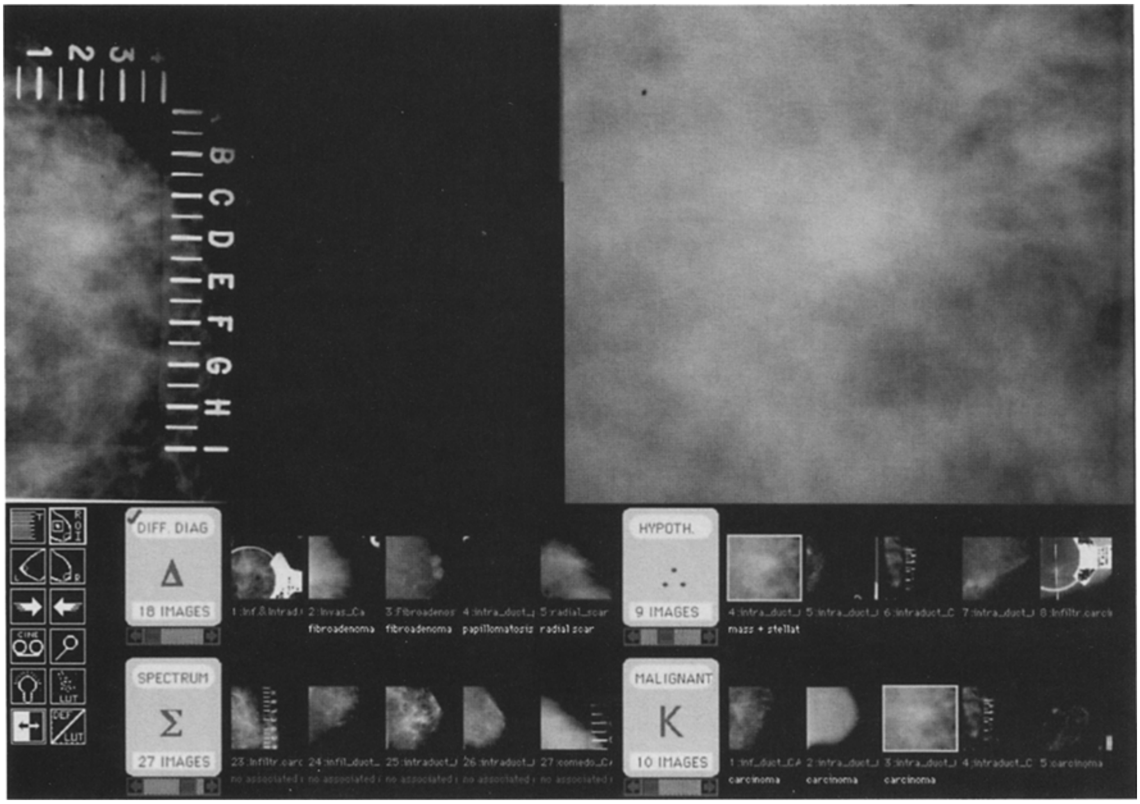

Fig 2. Axis display. The family of retrieved images shown in Fig 2 has been distributed among multiple similarity axes shown on the bottom half of the screen. Each of the four image sets shown contain multiple scrollable image tokens grouped and sequenced according to predefined, or user defined criterion. Here images are grouped by diagnosis (differential diagnosis axis), diagnostic hypothesis, disease spectrum, and malignant etiologies.

Hypothesis Axis

The number of images retrieved ranged from 0 to 11. In part, this reflected the rarity of some findings in the database. If the radiologist proposed a hypothesis that was unusual or was not well represented in the database, few or no retrievals occurred. This circumstance also occurred if the radiologist proposed an unlikely hypothesis (ie, a fat density lesion to be a carcinoma). This also reflected findings which were inadequately represented. Utility scores ranged from 5 to 7 with an average score of 5.6. There was one instance of 4. One incorrectly coded case was identified. Three cases were thought to be out of sequence.

Spectrum Axis

The number of images retrieved ranged from 2 to 27. The array was considered moderately helpful with scores ranging from 5 to 8 (average = 5.6).

Malignancy Axis

The number of images retrieved ranged from 0 to 11 cases. No images were retrieved for cases where findings are typically associated only with benign disease. In this case, the mammographers considered the axis to be of no clinical utility. There were no instances of inappropriately or incorrectly displayed images.

Fig 3. Hypothesis Axis. In this illustration, the major finding described for the index case are multiple smooth masses. If the radiologist proposes "cysts" as a provisional diagnosis, the Hypothesis Axis shows three cases matching the described findings (upper right image array). The Differential Diagnosis Axis (upper left image array) contains eleven examples of multiple smooth masses caused by a variety of etiologies. If the radiologist had proposed an unlikely diagnosis, the system would produce few if any retrievals.

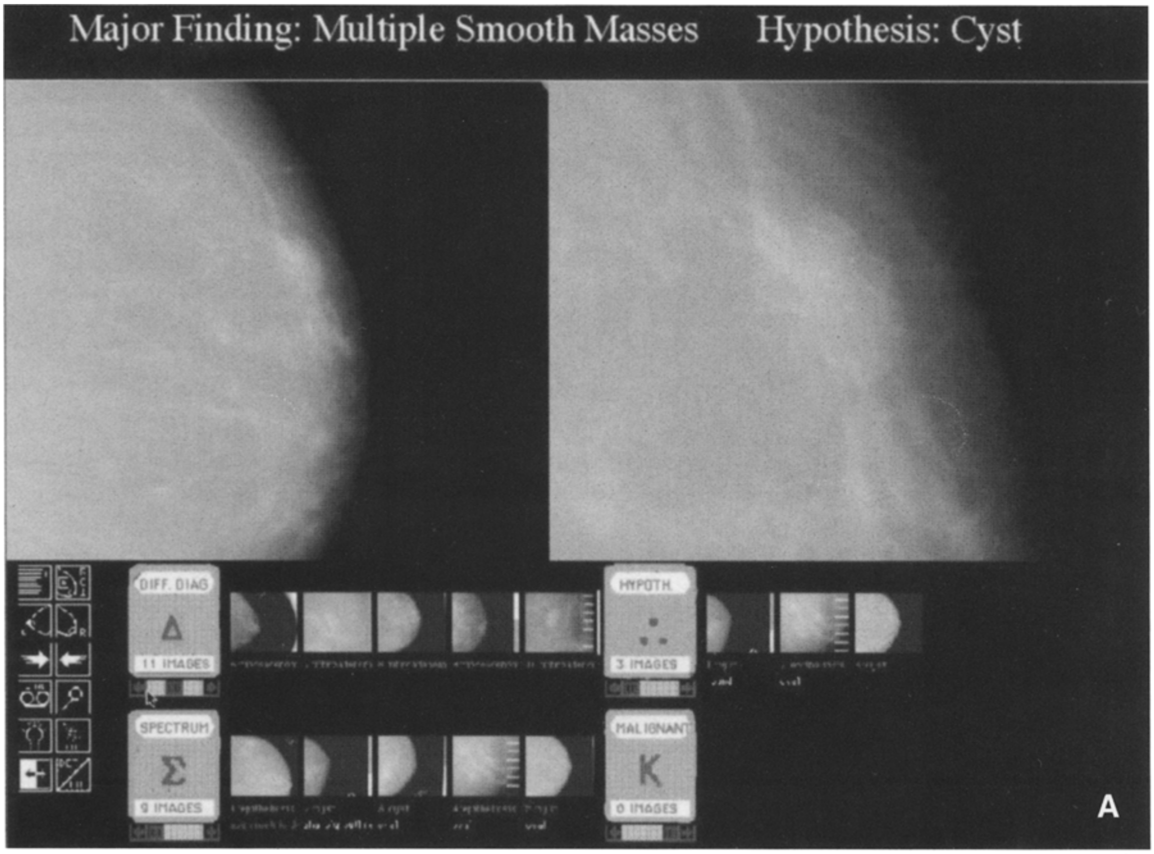
Major Finding: Multiple Smooth Masses
Hypothesis: Cyst
DIFF. DIAG
11 IMAGES
HYPOTH.
3 IMAGES
SPECTRUM
9 IMAGES
MALIGNANT
0 IMAGES
A

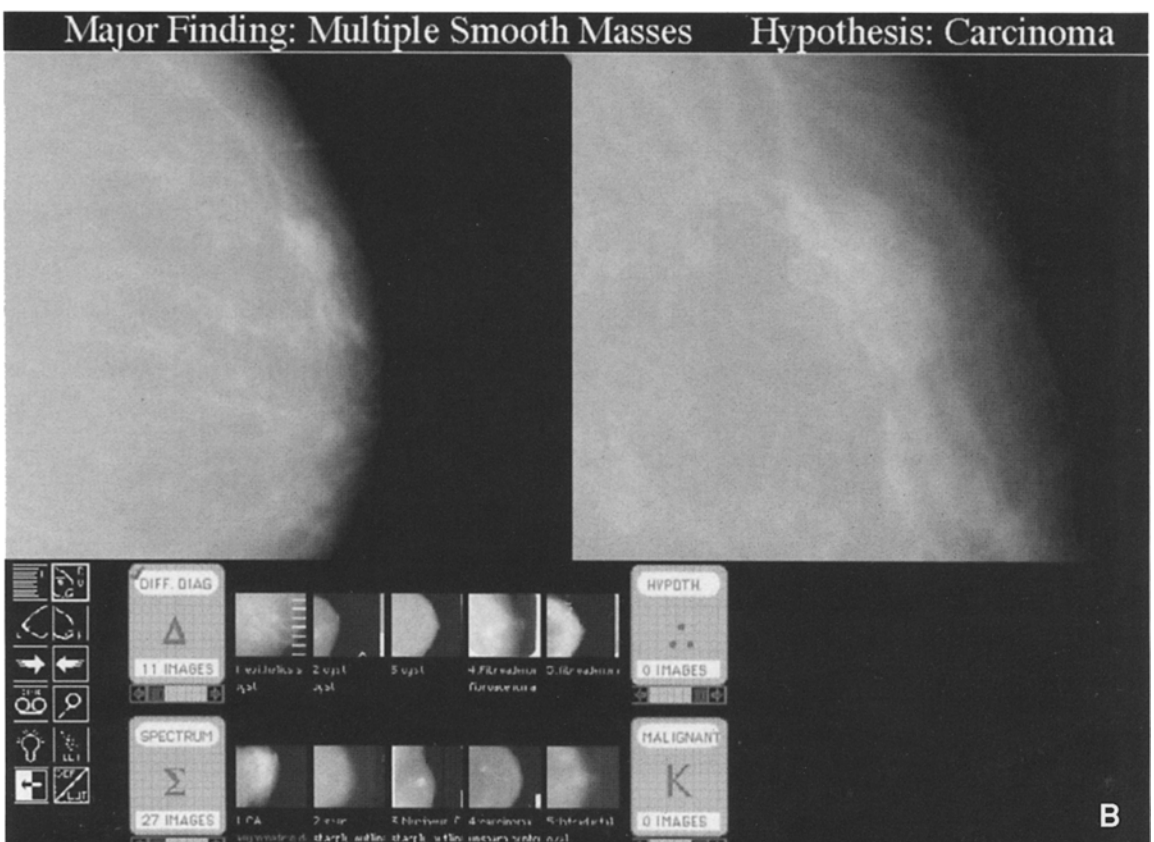
Major Finding: Multiple Smooth Masses
Hypothesis: Carcinoma
DIFF. DIAG
11 IMAGES
HYPOTH.
0 IMAGES
SPECTRUM
27 IMAGES
MALIGNANT
0 IMAGES
B

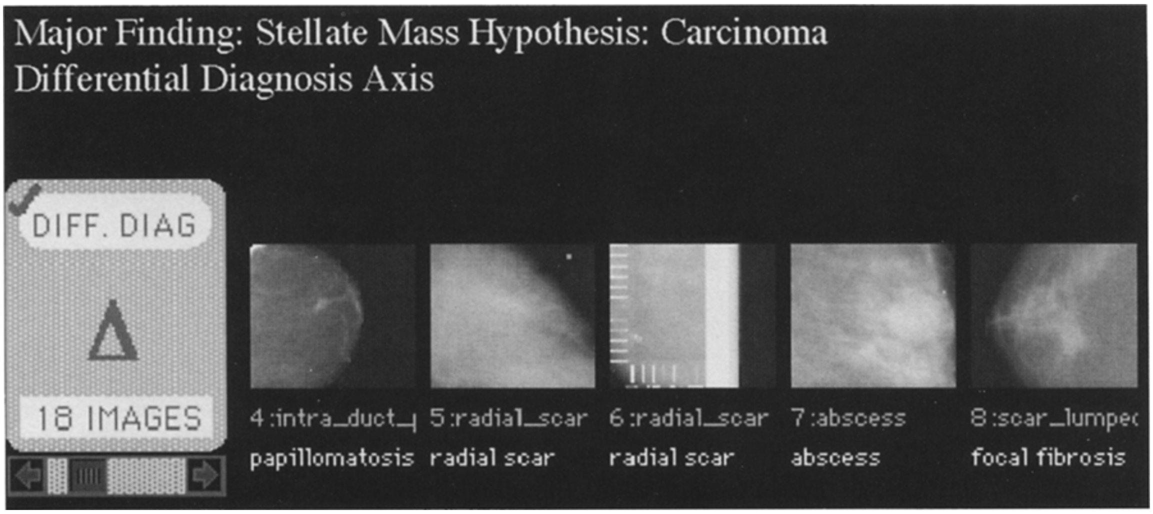

Fig 4. Differential diagnosis axis (close up). Images with similar features to the index case are displayed on this axis and are not constrained by diagnosis. This image array produces a visual differential diagnosis. In this case, the system searched for stellate masses with features similar to the index case. In addition to multiple carcinomas (not shown), other causes producing stellate masses are presented.

Best Match Axis

This axis always retrieved 141 images and the images were resequenced based on findings in the index case. The images were ordered in general with cases most like the index case presented first with increasingly dissimilar images presented with increasing distance down the array. The evaluators found this to be the second most helpful axis (clinical utility = 9, range = 4-9) but frequently disagreed with the ordering of the images and found that the large number of images presented detracted from the usefulness of the array.

User Defined Axis

From 1 to 34 images were presented on this axis depending on the specificity of the findings provided to the system. This was considered to be the most useful axis with an average score of 8.4 (range = 7-9). One image was considered to be inappropriately placed, and there were occasional incorrectly sequenced images.

DISCUSSION

Medical education often occurs at the wrong time. Education derived from general textbook and journal reading, or from attending continuing education courses, is the mainstay of continuing education for radiologists. Unfortunately, this knowledge may not be remembered readily when a specific clinical problem arises, particularly if the entity is rare or unusual. When time is not available to do case specific research on difficult cases, radiologists may resort to a variety of 'hedges' to deal with uncertainty. This may take the form of making

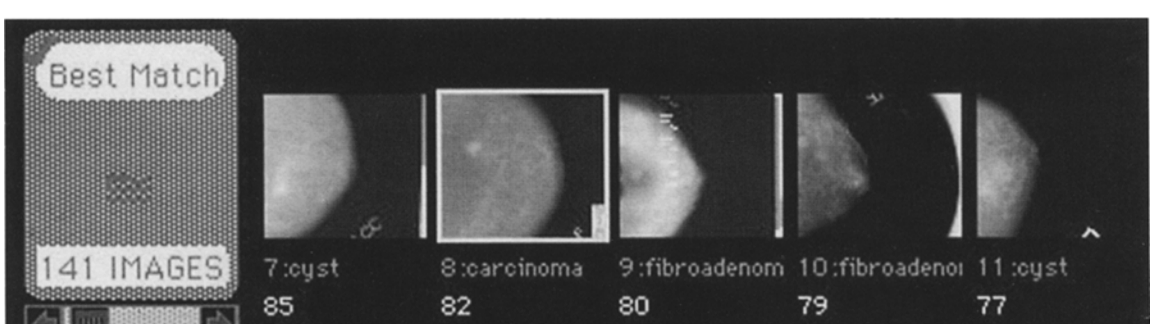

Fig 5. Best match axis.

Table 2. Subjective Estimate of Clinical Utility

	No. Cases	Differential Dx Axis	No. Cases	Hypothesis Axis	No. Cases	Spectrum Axis	No. Cases	Malignancy Axis	No. Cases	Best Match Axis
Stellate Mass										
Evaluator 1	18	8	9	7	27	5	10	6	141	9
Evaluator 2	18	9	9	6	27	5	10	6	141	8
High Density Mass										
Evaluator 1	5	8	2	7	27	7	2	8	141	8
Evaluator 2	5	8	2	5	27	7	2	9	141	9
Combined Density Mass										
Evaluator 1	3	7	1	6	2	5	0	6	141	8
Evaluator 2	3	7	1	6	2	5	0	6	141	8
Low Density Mass										
Evaluator 1	11	8	3	6	9	5	0	6	141	8
Evaluator 2	11	8	3	6	9	5	0	7	141	9
Radiolucent Mass										
Evaluator 1	1	4	2	5	2	6	0	1	141	8
Evaluator 2	2	8	2	5	2	5	0	3	141	9
Multiple Masses										
Evaluator 1	11	8	3	6	9	5	0	8	141	8
Evaluator 2	11	7	3	6	9	5	0	8	141	8
Clustered Calcifications										
Evaluator 1	28	7	11	5	27	6	11	8	141	8
Evaluator 2	28	10	11	4	27	8	11	6	141	4
Scattered Calcifications										
Evaluator 1	7	7	0	5	27	6	0	8	141	8
Evaluator 2	7	7	0	5	27	6	0	8	141	8

Note: Number is the number of images retrieved and placed in each array. The number under each axis name is the subjective clinical utility of the image set for the indicated prototypical case iterations. Key: 10 = Axis is very helpful; 5 = Axis is moderately helpful; 1 = Axis is not helpful.

general statements where specific diagnoses are possible and/or recommending additional examinations to establish a diagnosis with certainty. Our goal has been to develop computer systems that can maximize the diagnostic yield of specific examinations and to integrate these tools naturally into the image analysis and reporting process.

A variety of approaches to computer aided diagnosis (CAD) has been explored in radiology ranging from memory aids to expert advisory systems, from neural networks to CAD systems that process images directly and automatically identify nodules, masses, and calcifications.[3-7] A great deal of work has been done in mammography with the hope of developing tools that can point out potentially significant lesions, to characterize detected abnormalities, and to help radiologists confirm the significance of these findings.[8-15] We previously showed the ability to produce computer generated case discussions or "critiques" of radiological findings or differential diagnosis (ICON).[1] Decision making in radiology often depends more on visual pattern recognition than it does on cognitive decision making. For this reason we developed a system which could retrieve images from an image database in a context sensitive way (IMAGE/ICON).[2] We have now applied this approach to mammography (MAMMO/ICON) and automated image retrieval by providing a direct link to a speech recognition radiology reporting system. As a result, it is possible to dictate a written radiology report and simultaneously generate multiple searches of an image database. In effect, automatic context sensitive continuous education is available for every case when it is needed.

Retrieved images are organized into a series of image sets or axes that cut through a large number of potentially relevant images in a number of different ways. Each axis or image set is intended to anticipate pattern recognition issues that are likely to arise. For example, one commonly sees an abnormality and suspects that it may represent a specific diagnosis (a diagnostic hypothesis). MAMMO/ICON anticipates this clinical question by retrieving cases of the proposed diagnosis which have features similar to those seen in the index case and displays those images on a hypothesis axis. This allows the radiologist to be reassured that the suspected diagnosis may indeed cause the observed findings if similar cases are displayed. On the other hand, the radiologist might question the hypothesized diagnosis, if the system can find no examples of that diagnosis producing the observed findings. Another image axis retrieves all cases which share similar features with the index case regardless of diagnosis (differential diagnosis axis), a visual differential diagnosis or gamut.

The clinical evaluation that we performed was designed to determine whether MAMMO/ICON was functioning as designed. We learned that sometimes cases were not retrieved and placed appropriately in the expected image set because our coding lexicon was not robust enough to reflect the full breadth and subtlety of mammographic finding description. One of our central design philosophies was to sequence retrieved images in an intuitive and consistent manner. For example, a series of masses that were generally similar should be ordered in a smooth transition from sharp margin definition to fuzzy margin definition, from high mass density to low, and so on. When this failed to occur, it usually reflected deficiencies in our coding lexicon, or a lack of consistent coding standards by our coders. This last point highlights a key problem in any system that uses lexical data input to try and capture the complexity of visual observations. Word descriptions appear to be sufficient to produce broadly useful image retrieval, but we believe that image-based graphical data input will be a valuable way to refine image matching.

We did not try to evaluate the effect of MAMMO/ICON on a radiologist's accuracy. We did want to find out, however, if radiologists found this approach to be generally useful and which of the image axes were preferred. The mammographers who evaluated this system most often chose to explore the image database in an unstructured way by defining specific image features (the user defined axis). They also appreciated the best match axis, which preferentially retrieved images that closely matched the index case in all respects. Radiologists who are less experienced in mammography may find other image arrays to be of greater value.

The effectiveness of our approach depends on the composition of the database. Findings in the index case might be highly evocative of a specific diagnosis, but would not trigger a meaningful response if there were no cases of the evoked diagnosis in the database. Similarly, the number of examples retrieved may imply the strength that a particular diagnosis is suggested. That is, if the index case stimulates many retrievals of diagnosis X, but few of diagnosis Y, one might assume that X is a much more likely diagnosis. The current version of this system employs a relatively small database that highlights these problems. Large databases drawn from actual clinical experience are likely to have more meaningful implications, and may reflect local variations in disease incidence. Ultimately, conclusions based on actual practice and experience may lead to previously unrecognized associations between findings and disease entities.

The input to this system uses discrete speech recognition technology.[16] Predefined trigger phrases produce complete sentences or sentence fragments which are concatenated into a final report. These same trigger phrases are mapped to input feature descriptors used to search the database. Speech recognition technology is evolving rapidly[17] and natural free speech dictation will ultimately be feasible. The simple one-to-one mapping strategies we employ here may not be sufficient to provide reliable input values where context cannot be predicted. Nevertheless, the current system illustrates the power of direct linking of report dictation to automatic information retrieval in real time.

Automatic access to image libraries is even more attractive today as images are becoming available from many sources. Vast databases of digital radiological images are beginning to accumulate in picture archiving and communications systems (PACS), the Internet,[18] personal computer-based teaching files, and CD-ROM educational programs. In theory, the same approach could be used to search these other databases providing immediate access to a great deal of valuable reference material at the time of clinical decision making. For this to happen, new and more powerful image coding schemes will be required. Most radiologists use the American College of Radiology (ACR) code which is suitable for comparatively small image collections. The ACR code cannot represent detailed finding descriptions, finding quantification, and the relationships of findings and anatomic structures. It is unable to capture the kind of detailed finding descriptions, and relationships of findings that are required for complex image retrieval and sorting. We used a specially created image description scheme that is not generalizable. Other specially created lexicons of finding descriptions are useful in specific domains like mammography,[19-22] but can also not be generalized to other radiological disciplines. Systemized Nomenclature of Human and Veterinary Medicine (SNOMED), a coding method originally developed to code pathology findings, is now being extended to other domains in medicine.[23] SNOMED may be suitable for radiological images.[24,25] The Unified Medical Language System Project is a longer term effort to encode all anatomic, pathological, and clinical information in a consistent relational format that may eventually supersede other coding methodology.[26,27] Standardized structured reporting is likely to be incorporated into the Digital Imaging and Communications in Medicine (DICOM) standard. It will be critical that adequate coding schemes such as SNOMED be adopted soon so that the growing archive of radiological images will not be lost and the power of automatic image retrieval can be realized. As previously noted, accurate application of any of these coding schemes will be necessary to keep systems like ours from being misleading when applied to foreign databases.

REFERENCES

1. Swett HA, Miller PL: ICON: A computer-based approach to differential diagnosis in radiology. Radiology 163:555-558, 1987

2. Swett HA, Fisher PR, Miller PL, et al: Expert system controlled image display. Radiology 172:487-493, 1989

3. Kahn CE: Artificial intelligence in radiology: Decision support systems. Radiographics 14:849-861, 1994

4. Swett HA, Giger ML, Doi K: Computer Vision and Decision Support, in Perception of Visual Information, Hendee WR, and Well P (eds) Berlin, Springer-Verlag, Berlin, Germany, 1993, pp 272-315

5. Giger ML, Doi K, MacMahon H, et al: Computer-aided detection of pulmonary nodules in digital chest images. Radiographics 10:41-52, 1990

6. Yin F-F, Giger ML, Doi K, et al: Computerized detection of masses in digital mammograms: Analysis of bilateral-subtraction images. Med Phy 18:955-963, 1991

7. Chan HP, Doi K, Vyborny CJ, et al: Improvement in radiologists' detection of clustered microcalcifications on mammograms: The Potential of computer-aided diagnosis. Invest Radiol 25:1102-1110, 1990

8. Endo T, Kido C, Horita K, et al: Clinical evaluation of assistant diagnostic system for mammograms using the auto-analyzing method. Radiat Med 10:50-54, 1992

9. Huo Z, Giger ML, Vyborny CJ, et al: Analysis of spiculation in the computerized classification of mammographic masses. Med Phy 22:1569-1579, 1995

10. Wei D, Chan HP, Helvie MA, et al: Classification of mass and normal breast tissue on digital mammograms: Multiresolution texture analysis. Med Phy 22:1501-1513, 1995

11. Kegelmeyer WP, Pruneda JM, Bourland PD, et al: Computer-aided mammographic screening for spiculated lesions. Radiology 191:331-333, 1994

12. D'Orsi CJ, Getty DJ, Swets JA, et al: Reading and decision aids for improved accuracy and standardization of mammographic diagnosis. Radiology 184:619-622, 1992

13. Getty DJ, Pickett RM, D'Orsi CJ, et al: Enhanced interpretation of diagnostic images. Invest Radiol 23:240-252, 1988

14. Swets JA, Getty DJ, Pickett RM, et al: Enhancing and evaluating diagnostic accuracy. Med Decis Making 11:9-18, 1991

15. Vyborny CJ: Can computers help radiologists read mammograms? Radiology 191:315-317, 1994

16. Swett HA, Fisher PF, Mutalik PG, et al: Expert systems and clinical radiology: Integrating intelligent text and image display with voice activated reporting. Proceedings of the First International Conference on Image Management and Communication. Washington, DC, June, 1989 (suppl)

17. Herman SJ: Accuracy of a voice-to-text personal dictation system in the generation of radiology reports. AJR Am J Roentgenol 165:177-180, 1995

18. Mezrich RS, DeMarco JK, Negin S, et al: Radiology on the information superhighway. Radiology 195:73-81, 1995

19. Kopans DB: Standardized mammography reporting. Radiol Clin North Am 30(1):257-264, 1992

20. Heilbrunn KS: The American College of Radiology's mammography lexicon: Barking up the wrong tree? AJR Am J Roentgenol 162:593-594, 1994 (comment)

21. Baker JA, Kornguth PJ, Floyd CE: Breast imaging reporting and data system standardized mammography lexicon: Observer variability in lesion description. AJR Am J Roentgenol 166:773-778, 1996

22. D'Orsi CJ: The American College of Radiology mammography lexicon: An initial attempt to standardize terminology. AJR Am J Roentgenol 166:779-780, 1996

23. Rothwell DJ, Cote RA, Cordeau JP, et al: Developing a standard data structure for medical language—the SNOMED proposal. Proc Annu Symp Comput Appl Med Care 1993, pp 695-699

24. Friedman C, Alderson PO, Austin JH, et al: A general natural-language text processor for clinical radiology. JAMIA 1:161-174, 1994

25. Bell DS, Pattison-Gordon E, Greenes RA: Experiments in concept modeling for radiographic image reports. JAMIA 1:249-262, 1994

26. Lindberg DA, Humphreys BL, McCray AT: The Unified Medical Language System. Methods Inf Med 32:281-291, 1993

27. Tuttle MS, Nelson SJ: The role of the UMLS in 'storing' and 'sharing' across systems. Int J Biomed Comput 34:207-237, 1994