Symposium ■

Aspects of IAIMS Implementation That Require Further Research

Perry L. Miller, MD, PhD

A b s t r a c t   The Integrated Advanced Information Management System (IAIMS) program promotes an integrated approach to information management within a medical center. Since the IAIMS program was conceived, many of the initial IAIMS technologic needs have been quite widely achieved or are planned for implementation in many medical centers. At the same time, the IAIMS frontier is being steadily pushed to new issues that need to be addressed to achieve the full power of the IAIMS vision. The paper discusses 1) levels of integration where IAIMS has been successfully pursued to date and 2) challenging areas in which research is required to approach the full potential of the IAIMS in the future.

■ **J Am Med Inform Assoc.** 1997;4:March–April Supplement:S52–S61.

The Integrated Advanced Information Management System (IAIMS) program of the National Library of Medicine, designed to encourage an integrated institutional approach to information management within a medical center, was inaugurated in 1983.[1-4] The IAIMS vision was broadly defined to include the integration of information in all phases of a medical center's mission, including patient care, education, research, and administration. In addition, it was recognized that the goal of the IAIMS was not to confront these issues solely on a technological level, but rather to fulfill a need for a broadly based initiative involving a wide spectrum of people (clinicians, researchers, students, and staff) and institutional organizations, extending to the highest levels of the medical center's leadership.

Since the IAIMS program began, there have been dramatic changes in the power and cost of computing and communications, in the national network infrastructure, and in the infrastructures of most medical centers. As a result of these advances, many of the

IAIMS technologic needs envisioned in the early 1980s have been quite widely achieved, at least in part, or are being actively planned for implementation at many medical centers. For example, most major medical centers now have quite extensive computer networks. These networks may allow access from the desktop and from the clinical workplace to many information sources, including those on the Internet. In addition, many medical centers are planning or implementing clinical data repositories that will contain a wide range of clinical and administrative data that can be analyzed in an integrated fashion for diverse purposes.

As these major phases of integration proceed, the IAIMS frontier is being steadily pushed to new areas and new issues that must be addressed to achieve the full power of the IAIMS vision. For example, this vision is now leading researchers beyond issues involving hardware and software to address the problems of integration at the level of *language* and at the level of the underlying clinical *concepts* involved.

This paper first briefly reviews certain levels of integration where the IAIMS program has been successfully pursued to date. The paper then discusses challenging areas in which research is required to allow us to achieve the full potential of IAIMSs in the future. The goal of the paper is not to provide a comprehensive review of these areas. Indeed, if one defines these issues broadly, almost all research in the field of medical informatics can be seen as contributing to the ultimate fulfillment of IAIMSs. Instead, the goal of the paper is to discuss representative areas in which re-

search will make significant contributions to the overall IAIMS vision, focusing particularly on the clinical arena.

Areas of Past and Current IAIMS Focus

Establishing Underlying Network Connectivity

In the early 1980s, when the IAIMS concept was introduced, many important computer systems within a medical center typically ran on stand-alone computers accessed by phone modem or by dedicated lines. As a result, an early IAIMS focus was on building an institutional network that linked the major computer facilities and made them more readily accessible from many locations within the medical center and beyond.

"Front-end" Integration of Access from the Desktop

As network connectivity was established, an important area for integration involved facilitating access from the desktop machine, for example, by developing a unifying front end allowing *connection* to multiple information sources. Examples of this approach at different IAIMS institutions include Bio-SYNTHESIS,[5] Willow,[6] the Knowledge Workstation,[7] and NetMenu.[8] A more difficult task is to provide automated integrated *querying* of multiple information sources. The Intelligent Query Workstation[9] explored how to provide this capability, accommodating the existing interfaces of different sources. In contrast, the Z39.50 protocol[10] provides a common interface to multiple information servers, so that a client can query different sources in a uniform fashion.

"Back-end" Integration of Data for Use by Multiple Applications

A more powerful level of integration can be achieved by the "back-end" integration of the multiple computer systems and applications.[11] One way in which this form of integration can be achieved is by extracting data from a medical center's various transaction systems (e.g., from its clinical information systems and its administrative and billing systems) and storing copies of those data in a data repository. The transaction systems are optimized to allow the rapid response required in the busy clinical environment. In contrast, data repositories are typically network-based relational databases designed to support applications that access the data for such diverse purposes as financial and administrative analyses, clinical outcomes and quality-of-care studies, analysis of practice patterns, epidemiologic studies, and clinical research.

In addition to data repositories that allow many types of data to be analyzed in an integrated fashion, another form of "back-end" integration involves linking the various applications to one another in an increasingly robust fashion. One example involves allowing online patient data to be passed to a clinical decision-support system so that the system can be invoked automatically. Another example involves allowing automatic linkage from a patient record to those parts of an online textbook that are relevant to the patient's medical management.

Areas of Research to Achieve the Full IAIMS Vision for the Future

The type of "back-end" integration described above can be pursued at various degrees of sophistication. Indeed, the most sophisticated linkages are performed by humans using human intelligence. As a result, to achieve the most sophisticated type of linkage, a tremendously intelligent machine would be required, far beyond current capabilities. The practical challenge is to define an appropriate level of sophistication at which one can realistically expect to produce clinically useful integration. The remainder of this paper discusses a number of the current areas where researchers are working to develop such advanced capabilities. This research will help build advanced forms of integration that will incrementally approach the full potential of the IAIMS.

Research Involving Standardized Clinical Vocabularies

The need to develop standardized clinical vocabularies is well recognized. Such vocabularies are required to allow the different clinical information systems within a single medical center to be integrated, and also to allow standardized analyses to be made among multiple medical centers. Standardized clinical vocabularies will also allow clinical decision-support logic developed at one institution to be applied to the clinical data elements of other institutions.

A variety of controlled vocabularies have been developed, each typically focused on particular needs. The International Classification of Diseases, Ninth Edition, with Clinical Manifestations (ICD-9 CM),[12] Current Procedural Terminology (CPT),[13] and Current Medical Information and Technology (CMIT)[14] were developed for administrative analyses and billing. The Systematized Nomenclature of Medicine (SNOMED),[15] the READ Clinical Classification System,[16] and the Gabrieli Nomenclature[17] focus on the need to capture data from the clinical encounter. Medical Subject
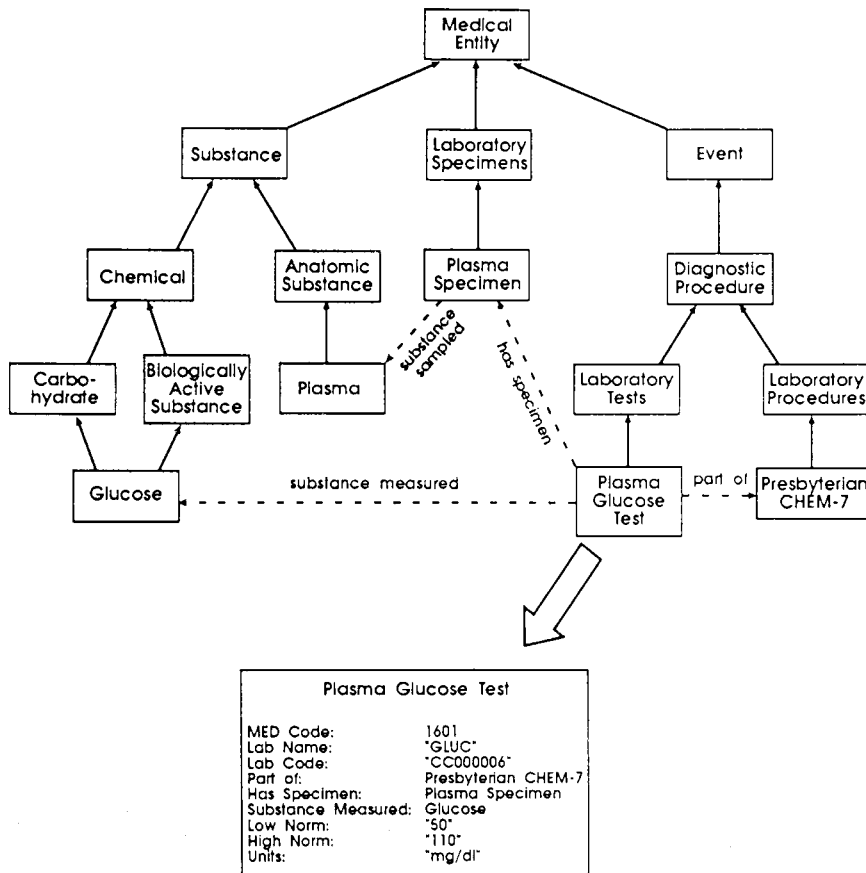
Figure 1 An example of higher-level structure in a controlled structured vocabulary. Here Columbia's Medical Entities Dictionary (MED) contains both hierarchical and non-hierarchical relationships that provide a context for each of its vocabulary entities, in this case "plasma glucose test." (Reproduced with permission from Cimino et al.[20])

Headings (MeSH)[18] were developed for bibliographic indexing and retrieval. The Unified Medical Language Project Metathesaurus[19] was designed to help link terms from various medical vocabularies.

A wide range of research issues arise in the creation of a standardized clinical vocabulary. A number of the challenges stem from the scope of the material being described. A key challenge will be to define a vocabulary to cover a patient's history and physical examination. Another challenge will be to cover all the clinically relevant findings possible in each of the highly specialized subfields of a discipline such as diagnostic imaging.

These issues are compounded by the need to develop a *structured* standardized vocabulary.[20] While it is important to develop 1) a standardized approach to naming the individual vocabulary elements (the terms, concepts, or entities), it is equally important to develop 2) a standardized higher-level structure (the hierarchical and non-hierarchical semantic relationships that place the terms into their broader context; see Figure 1), and also 3) the lower-level modifiers that provide the underlying definition of each term.

Hierarchical relationships between terms play an important role in allowing clinical decision logic to be defined in terms of groups of terms (e.g., cephalosporin antibiotics) so that the logic can concisely reference all such terms. Non-hierarchical semantic relationships are useful in formulating decision logic, and are also potentially useful in information retrieval.[21] Lower-level modifiers (e.g., defining the normal range of a laboratory test for a specific age range and over a specific period of time) are important for the decision logic to operate correctly and also to allow comparisons across different medical centers and different time periods.[22]

Determining how best to select a set of vocabulary terms, and how best to define the higher-level and lower-level structures in which those terms fit, is a major ongoing research area. The optimal solution depends on the many potential uses to which the clinical vocabulary will be put. Many of these uses are only beginning to be explored at the present time.

Identifying Relevant Network-based Information Sources

Another challenge is posed by the rapid proliferation of network-based information resources relevant to bi-

Figure 2 Yale's pilot implementation of the UMLS information sources map. In the top figure (a), the user has entered the term "cancer." The computer then searches its database of roughly 140 information sources and returns a set of potentially relevant sources (middle and bottom figures, b, c) ranked by "priority." The user may then select a source and request either further information or automatic connection. The user may also request that the set of sources retrieved be regrouped according to several other criteria in addition to the "priority" ranking shown here. (Reproduced with permission from Miller et al.[27])
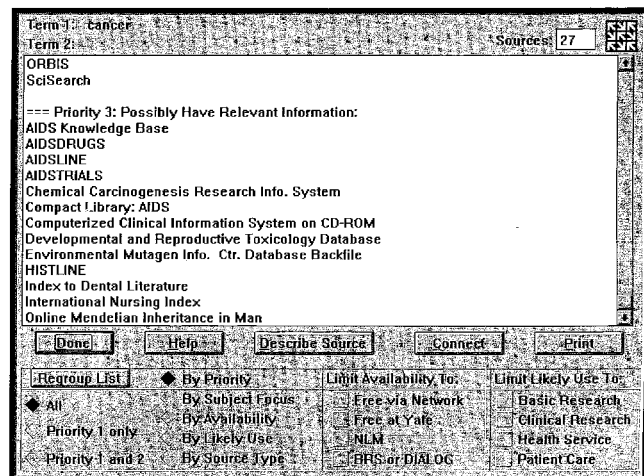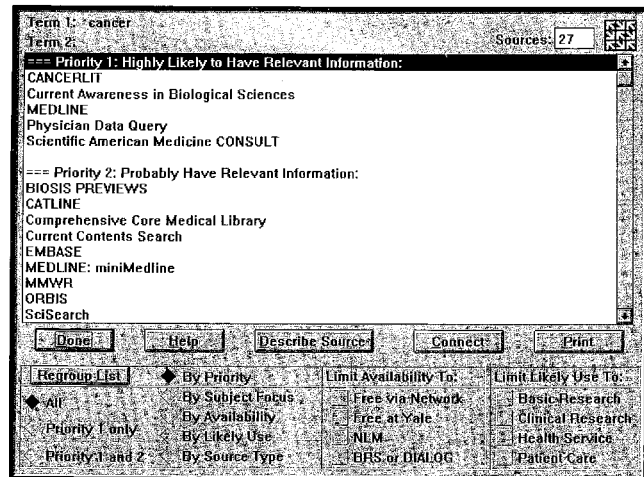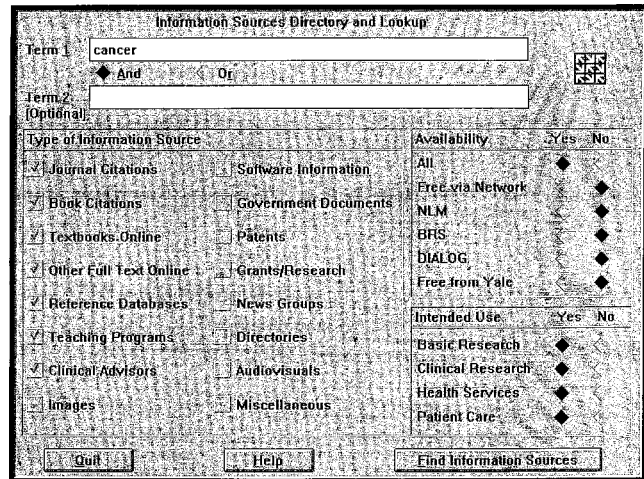

omedicine. These information sources vary widely in content, scope, quality, and accuracy. Many different types of information are available, including bibliographic citations, full texts of journals and textbooks, scientific facts, clinical and research images, computer-based expert clinical advice, and computer-assisted teaching. As the use of the World Wide Web grows, many more, increasingly diverse, types of information are being made available.

An important research area involves helping the user find information sources relevant to a particular question. There are a number of national networking and library activities, not focusing specifically on the health sciences, that are dealing with Internet directories and resource discovery.[23] Approaches to network navigation and resource discovery include Gopher, the World Wide Web (WWW), and Wide Area Information Servers.[24,25] Initial efforts to create directories of Internet-based resources included the Internet Resources Meta-Index maintained by the National Center for Supercomputing Applications (NCSA) on WWW, the InterNIC Directory maintained by AT&T, and the OCLC Catalog of Internet Services.[26] WWW browsers such as Netscape now provide a host of different search engines that organize and categorize WWW-based resources.

Since there is a rapidly growing number of sources of greatly varying quality, it will be important to be able to focus selectively on promising sources for particular purposes. One research project that has confronted these issues in biomedicine is the Unified Medical Language System's (UMLS's) information sources map (ISM), which has been implemented on a pilot basis (see Figure 2) and is currently in routine use at Yale.[27] The Sourcerer project[28] is exploring how best to implement this type of capability on the World Wide Web.

Accommodating Automated Transfer of Online Clinical Data to Clinical Decision-support Applications

For the past two decades or more, there has been extensive work on the development of expert consulta-

tion systems in medicine, and more recently on the closely related task of placing clinical guidelines into an interactive, patient-specific, computer-based form. Such clinical decision-support systems (CDSSs) will clearly be most useful if the patient-specific data that they require as input can be passed automatically from an online patient database.

For this to happen in a sophisticated way, a number of challenges must be faced in addition to the issues involving standardized vocabularies discussed above. One problem involves determining what data the CDSS needs. This is a potential problem since as a clinical field evolves, the specific set of data required by a CDSS may change over time. Ideally, one would like the CDSS, when invoked, to produce a list of the clinical data it requires as input, so that that data can be automatically extracted from the database.[29] Without this type of ''dynamic linking'' between a CDSS and a patient database, each time the CDSS is modified, its interface with each database where it is used will need to be modified as well.

A related question is how to handle the possibility that a given patient database may not contain all the clinical data items required by a CDSS. One solution is to design the CDSS to ask the clinician user for any data items not in the patient database. Another potential solution, however, arises from the fact that the granularity of the advice produced by a CDSS might be tailored dynamically to the clinical data items available.[30] Where this is possible, a CDSS might respond to the lack of certain data items by producing a somewhat larger quantity of textual output that covers the implications of the several different values that the missing data item(s) might have. The clinician would then decide which parts were relevant to a specific patient.

Structuring Textual Knowledge to Facilitate Patient-specific Retrieval

A closely related problem involves exploring how best to structure textual material to facilitate patient-specific retrieval. Historically, text has been organized linearly to facilitate reading material on the printed page. With the advent of computer-based text, more powerful interactive organization is possible. Hypertext represents one approach that allows the material to be browsed in a more flexible fashion.

When linking textual reference material to patient care, however, one would like to retrieve the material in a patient-specific fashion. If one takes a simpleminded approach to this problem and, for example, divides a conventional review paper into ''chunks'' and then retrieves only those chunks relevant to a specific patient, one risks having a presentation that is disorganized and lacks the overview and context required to appreciate the material fully.

One example of patient-specific retrieval is an expert system that takes as input a patient description and outputs a highly patient-specific set of recommendations. This approach typically requires a great deal of iterative refinement, validation with a large set of test cases, and periodic revalidation as the clinical field evolves. At the other extreme, a system might produce large sections of generally applicable text in a coarsely selective fashion, and let the clinician user decide what is relevant. There is also a spectrum between these two extremes. In many cases, a relatively coarsely granular approach may be quite satisfactory from a clinical standpoint. As a result, there are a number of interesting research issues exploring the tradeoffs between the granularity of the material presented, its clinical utility, its ease of verification and maintenance, and the ability to provide appropriate context for evaluating the material.

Indexing Images to Provide Integrated Access for Diverse Purposes

Clinical images are an important part of the patient record and provide a number of special requirements for computer-based systems. These include the need for adequate hardware and software to view the images, high communication bandwidth to transmit the images, and large-capacity memory to store a medical centers's clinical images.

An additional challenge involves exploring how best to index the images so that they can be retrieved and used robustly in a variety of applications. Such indexing would include keywords indicating the type of image and the imaging modality. The biggest challenge, however, concerns how best to index the image's content. Ideally, one would like to describe the clinically relevant features of the image in a coded fashion that could be used for a variety of purposes, such as input to a CDSS, to a quality of care evaluation system, to a teaching program, etc. One approach is to attempt to index an image's contents through natural-language processing of a radiologist's clinical report.[31] A complementary approach involves deriving mathematical features that can be computed from the images themselves, thereby facilitating the automatic indexing of a large number of images.[32,33] Another set of issues revolves around accessing the images. For example, using ''axes'' of similarity[34] (see Figure 3) allows retrieval of representative images similar to a given image for patient care or for educational purposes.

Figure 3 AXON, a research prototype built to assist in the interpretation of chest radiographs for patients with lymphoproliferative disease. Here the large image shows pulmonary infiltrates in a patient with leukemia, where the infiltrates are a manifesta4tion of leukemia itself. AXON explored how the retrieval of images could be performed along a variety of "axes of similarity," which in this case includes similar diseases (e.g., infiltrates caused by other lymphoproliferative diseases) and other etiologies (e.g., infiltrates caused by different bacterial and viral infections in patients with leukemia). (Reproduced with permission from Cohn et al.[34])
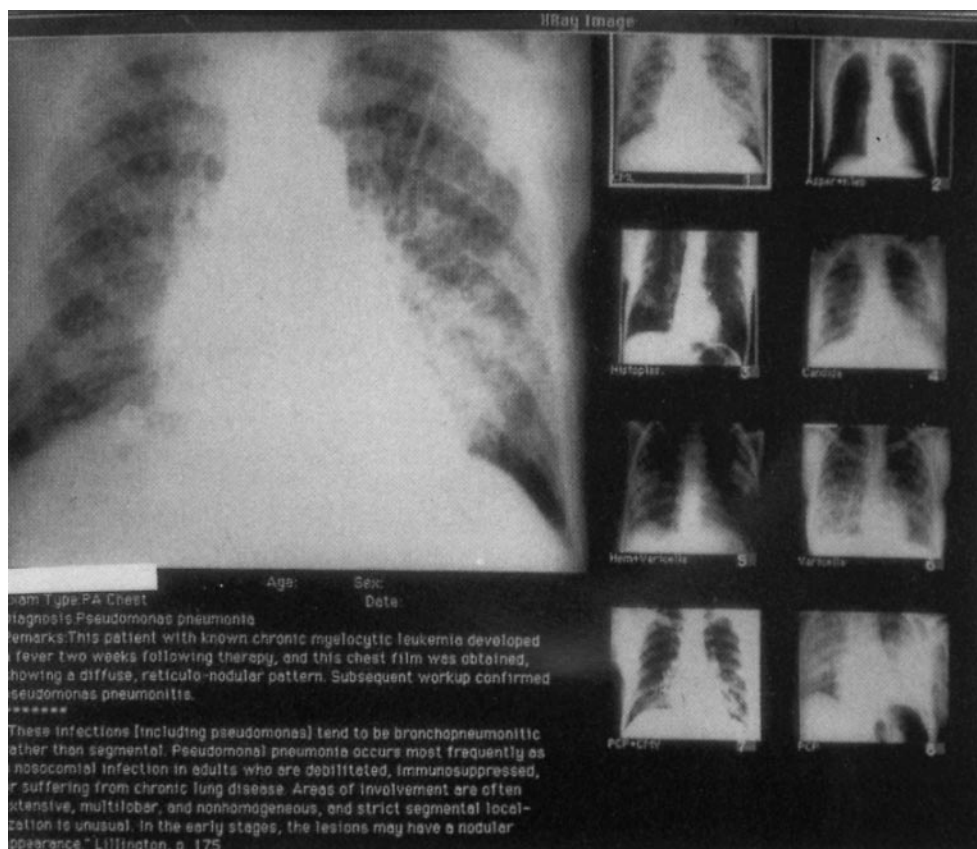
Allowing Conceptual Links from the Computer-based Patient Record to Online Reference Information

When a clinician reads a patient record of any size and complexity, a wide range of potentially relevant clinical topics typically come to mind. Many of these topics are explicitly referenced by the words used in the record itself, such as the patient's diseases and medications, and the procedures used. Many other topics, however, are implicitly inferred from the patient record based on the clinician's knowledge of risks, complications, special clinical considerations, etc.

When a patient record is online, one would ideally like the computer itself to be able to identify relevant clinical topics and to provide automated links to online information about each topic. Several systems have explored how to accomplish such automated conceptual linking between the computer-based patient record and online reference materials. HepaTopix[35] and PsychTopix[36] (see Figure 4) explored how such linkage could be achieved based on a conceptual schema of a clinical field (an outline of the key clinical topics), where each topic has associated 1) an "activation expression" of keywords used to indicate when that topic might be relevant to a patient record, and 2) "search expressions" to retrieve information about the topic from different information sources. An alternative approach was explored by CHARTLINE,[37] which used the UMLS Metathesaurus' Main Concept terms and their synonyms to retrieve bibliographic references relevant to a patient record, utilizing feedback from the user to help identify useful topics.

Graphic Visualization for Information Browsing and Linking

An interesting current area of research in information retrieval involves the development of graphic interfaces that allow the user to visualize the process of information retrieval. Particularly interesting is three-dimensional visualization.[38] This work involves the development of visual models, such as an information retrieval "cone tree," or an "information village" where buildings represent different information sources, to let the user better visualize the available information. Figure 5 shows an information retrieval cone used to help visualize retrieval from oncology information sources.[39] A potentially exciting research

Last-Name:R
First-Middle-Name:G
Sex:M
Date-Born:MM/YY/DD
Room:6-3XX
Consult-Number:YYMMDDN
Unit-Number:123-45-67
Admitted:MM/DD/YY

REASON FOR CONSULT: Has been acting "hypomanic"-- euphoric, up at night, and adjusted
his own cyclosporin dose.

CONSULT REPORT: Middle aged, white, married male in outpatient therapy on and off, who had
a heart transplant 2/5 for worsening cardiomyopathy. Did fairly well... episode of fever and rash
starting about 2/16 which responded to steroids... azotemia has improved.
   His current symptoms of euphoric mood, increased activity, overly familiar behavior with the
nursing staff, and insomnia started about 2/28. ... Notes of 3/3 indicate that he was adjusting his
own cyclosporin dose.
... He had been given a solumedrol "pulse" back on 2/20 and he was already on his current
prednisone dose (50mg/day).
   Patient has seen his outpatient therapist, Mr J S, MSW at least three times in the hospital....

   PAST PSYCH HISTORY: Went for "counseling" in college ... says he did well. ... Says
medications were never recommended. Denies any hospitalizations.

   FAMILY HISTORY: Says his sister was "schizophrenic" and hospitalized for a number of years
in her twenties. Denies other family history...

   SOCIAL HISTORY: Former salesman and corporate executive...

   MEDICATIONS: Cyclosporin Imuran Prednisone Benadryl Seconal Percocet Aspirin Persantine
Clortimazole troches Carafate Amphogel Mylanta Colace KCL FeSO4 Vitamin K Multi-Vitamins

   MENTAL STATUS EXAM: ... Mood upbeat-- "restraining myself a bit because of the nurses
comments." Was fully oriented. Speech-- bit grandiose and rambled. Talked of the book he's
going to write, mountain climbing, skiing, and how his course is going to be "like a space launch."
Cognition mildly impaired: couple of mistakes doing serial sevens and one mistake spelling
"world" backwards. Judgement mildly impaired.... No signs of sadness or lability...

   ASSESSMENT: Organic affective disorder, probably from steroids... Cognitive impairment
seems out of proportion to euphoria for manic depressive disease...

   PLAN: Haldol 0.5mg po qHS and 0.5mg po up to TID PRN to try to control insomnia,
behavior, and confusion. Discussed with attending and housestaff...

Topics organized by DSM-III-R [18]
      Organic Mental Disorders (DSM3R p97)
            Organic Mental Disorders Secondary to Physical Disorders (DSM3R p162)
                  Iatrogenic organic mental disorders
                        Steroid induced psychosis or mental changes
                  Organic mental disorder NOS (294.80)
      Mood Disorders (DSM3R p213)
            Bipolar Disorders
                  Bipolar disorder NOS (296.70)

Doherty M ; Garstin I ; McClelland RJ ; Rowlands BJ ; Collins BJ
A steroid stupor in a surgical ward.
Br J Psychiatry 1991 Jan;158:125-7

Kostic VS ; Levic Z
Psychiatric disorders associated with corticosteroid therapy.
Neurologija 1989;38(2):161-6

Viswanathan R ; Glickman L
Clonazepam in the treatment of steroid-induced mania in a patient
after renal transplantation [letter]
N Engl J Med 1989 Feb 2;320(5):319-20

Turner P
Steroid-induced psychosis [letter]
Lancet 1989 Oct 14;2(8668):923

d'Orban PT
Steroid-induced psychosis.
Lancet 1989 Sep 16;2(8664):694

Figure 4 This three-part figure
shows PsychTopix analyzing an on-
line psychiatric consultation report.
PsychTopix scans the textual report
looking for various combinations of
keywords.

In (a), several of the keywords found
have been underlined ("steroids,"
"euphoria,") "cognition ... im-
paired").

PsychTopix then selects, from a hi-
erarchical list of clinical topics, a spe-
cific set of topics (b) that it deter-
mines are potentially relevant to the
case.

Finally, the clinician clicks on one of
the topics and requests that a Med-
line search be performed, automati-
cally retrieving relevant citations (c).
(Reproduced with permission from
Powsner and Miller.[36])
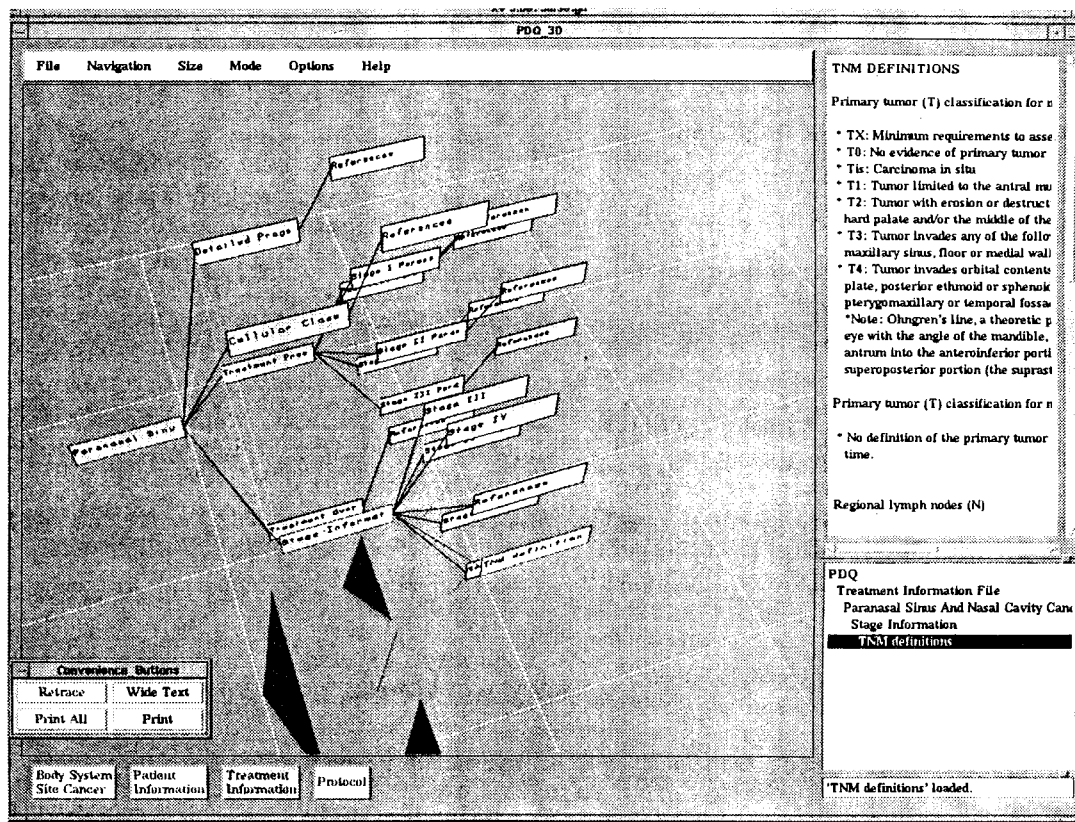
Figure 5 An information cone tree used to retrieve information about oncology. The graphic presentation is designed to help better visualize the process of information retrieval. (Reproduced with permission from Cole et al.[39])

area involves using this technology to let the user better understand and explore the wealth of biomedical information available via computer networks.

Tools for Language and Knowledge Maintenance

As IAIMS integration increasingly takes place at the level of language and clinical concepts, it will be important to develop software tools 1) to facilitate the acquisition and organization of all the domain knowledge involved and 2) to help with the maintenance of that knowledge as it evolves over time. The elements of a standardized clinical vocabulary, including the high-level and lower-level structures, will continually evolve, as will the clinical knowledge of medicine itself.

Difficult as it is to create clinical vocabularies and clinical knowledge bases, the major challenge will often be to maintain these resources over time. Computer-based tools will be essential for this process. A particularly interesting research area will involve developing tools that can be used on an inter-institutional basis, e.g., over the Internet, to facilitate such maintenance functions. The current Intermed Collaborative is an example of how such an inter-institutional ap-

proach might be implemented.[40] Researchers at several medical centers are using the Internet to share software and information resources and to build collaboratively new applications with shared components.

The Scope of the Present Paper

This paper describes a set of IAIMS-related research areas in the clinical arena. A list of research topics could similarly be outlined for other IAIMS areas such as teaching and bioscience research. There are also important issues that arise in applying new technologies in support of IAIMSs. Examples include 1) network authentication and access control using tools such as Kerberos,[41] and 2) the development of effective means of information commerce[42] in the rapidly emerging network-based marketplace. In addition, there are important research issues involving reengineering the clinical workplace, e.g., redesigning clinical workflow, to take maximum advantage of IAIMS concepts. The present paper has focused on representative research areas that enhance information integration in the context of patient care.

Conclusion

As IAIMS projects progress and as the available institutional infrastructure matures, the types of integration that are possible will become increasingly sophisticated. The vision of the IAIMS as articulated in the early 1980s has proven to be extremely prescient, and much of that vision is already becoming reality at many medical centers. This new infrastructure provides opportunities to move the IAIMS vision forward to meet further challenges. This paper has attempted to summarize some of the past and present foci of the IAIMS, and some of the challenging areas where research is needed to enable an increasingly robust implementation of IAIMSs in the future.

References ∎

1. Lindberg DAB. The IAIMS initiatives of NLM: institutional planning for advanced information services. J Am Soc Info Sci. 1988;39:105–6.
2. Lindberg DAB. The IAIMS opportunity: the NLM view. Bull Med Libr Assoc. 1988;76:224–5.
3. Matheson NW, Cooper JAD. Academic information in the health sciences center: roles for the library in information management. J Med Educ. 1982;57(part 2):1.
4. West RT. The National Library of Medicine's IAIMS grant program: experiences and futures. J Am Soc Info Sci. 1988;39:142–5.
5. Broering NC, Gault HR. BioSYNTHESIS: bridging the information gap. Bull Med Libr Assoc. 1989;77:19–25.
6. Ketchell DS, Fuller SS, Freedman M. Collaborative development of a uniform graphical interface. In: Frisse M (ed). Proc Sixteenth Annu Symp Computer Applications in Medical Care. Washington, DC: American Medical Informatics Association, 1992;251–5.
7. Lucier RE, Matheson NW, Butter KA, Reynolds RE. The knowledge workstation: an electronic environment for knowledge management. Bull Med Libr Assoc. 1988;78:248–55.
8. Shifman MA, Clyman JI, Paton JA, Powsner SM, Roderer NK, Miller PL. NetMenu: experience in the implementation of an institutional menu of information sources. In: Safran C (ed). Proc Seventeenth Annu Symp Computer Applications in Medical Care. New York: McGraw-Hill, 1993;554–8.
9. Cimino C, Barnett GO, Hassan L, Blewett DR, Piggins JL. Interactive query workstation: standardizing access to computer-based medical resources. Comput Meth Progr Biomed. 1991;35:293–9.
10. Lynch CA. The client–server model in information retrieval. In: Dillon M (ed). Interfaces for Information Retrieval and Online Systems: A State of the Art. Westport, CT: Greenwood Press, 1991;301–18.
11. Clayton PD, Sideli RV, Sengupta S. Open architecture and integrated information at Columbia–Presbyterian Medical Center. MD Comput. 1992;9:297–303.
12. United States National Center for Health Care Statistics (ed). International Classification of Diseases, Ninth Revision, with Clinical Manifestations. Washington, DC: National Center for Health Statistics, 1980.
13. American Medical Association. Physicians' Current Procedural Terminology. Chicago, IL: American Medical Association, 1995.
14. Gordon BL (ed). Current Medical Information and Terminology. Chicago, IL: American Medical Association, 1971.
15. Côté RA, Rothwell DJ, Palotay J, Beckett R, Brochu L (eds). The Systematized Nomenclature of Medicine. Northfield, IL: College of American Pathologists, 1993.
16. Saint Ives IF. The READ Clinical Classification System. Health Bull. 1992;50:422–7.
17. Gabrieli ER. A new electronic medical nomenclature. J Med Systems. 1989;13:355–73.
18. Medical Subject Headings. Bethesda, MD: National Library of Medicine, 1995.
19. Humphreys B, Lindberg D. Building the Unified Medial Language System. In: Kingsland L (ed). Proc Thirteenth Annu Symp Computer Applications in Medical Care. Washington, DC: IEEE Computer Society Press, 1989;475–80.
20. Cimino JJ, Clayton PD, Hripcsak G, Johnson SB. Knowledge-based approaches to the maintenance of a large controlled medical terminology. J Am Med Informat Assoc. 1994;1:35–50.
21. Miller PL, Barwick KW, Morrow JS, Powsner SM, Riely CA. Semantic relationships and medical bibliographic retrieval: a preliminary assessment. Comput Biomed Res. 1988;21:64–77.
22. Kannry JL, Wright L, Shifman M, Silverstein S, Miller PL. Portability issues for a structured clinical vocabulary: mapping from Yale to the Columbia Medical Entities Dictionary. J Am Med Informat Assoc. 1996;3:66–78.
23. Obraczka K, Danzig PB, Li SH. Internet resource discovery services. Computer. September 1993;8–22.
24. Krol E. The Whole Internet: User's Guide and Catalog. Sebastopol, CA: O'Reilly & Associates, 1992.
25. Powsner SM, Roderer NK. Navigating the Internet. Bull Med Libr Assoc. 1994;82:419–25.
26. Dillon M. Assessing information on the internet: toward providing library services for computer-mediated communication. Internet Res. 1993;3:54–69.
27. Miller PL, Frawley SJ, Wright L, Roderer NK, Powsner SM. Lessons learned from a pilot implementation of the UMLS information sources map. J Am Med Informat Assoc. 1995;2:102–15.
28. Rodgers RPC. Automated retrieval from multiple disparate information sources: the World Wide Web and the NLM's Sourcerer project. J Am Soc Infor Sci. 1995;46:755–64.
29. Miller PL, Paton JA, Clyman JI, Powsner SM. Prototyping an institutional IAIMS/UMLS information environment for an academic medical center. Bull Med Libr Assoc. 1992;80:281–7.
30. Miller PL, Frawley SJ. Tradeoffs in producing patient-specific recommendations from a computer-based clinical guideline: a case study. J Am Med Informat Assoc. 1995;2:238–42.
31. Friedman C, Alderson PO, Austin JH, Cimino JJ, Johnson SB. A general natural-language text processor for clinical radiology. J Am Med Informat Assoc. 1994;1:161–74.
32. Tagare HD, Vos FM, Jaffe C, Duncan JS. Arrangement: a spatial relation between parts for evaluating similarity of tomographic sections. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1995;17:880.
33. Zink S, Jaffe CC. Medical imaging data bases. Invest Radiol. 1993;28:366–72.
34. Cohn AI, Fisher PR, Mutalik P, Miller PL, Swett HA. AXON:

knowledge-based retrieval of radiographic studies using ''axes'' of clinical relevance. Comput Biomed Res. 1990;23: 199–221.

35. Powsner SM, Barwick KW, Morrow JS, Riely CA, Miller PL. Automated bibliographic retrieval based on current topics in hepatology: HEPATOPIX. Comput Biomed Res. 1989;22: 552–64.

36. Powsner SM, Miller PL. Automated online transition from the medical record to the psychiatric literature. Meth Info Med. 1992;31:169–74.

37. Gieszczykiewicz FM, Vries JK, Cooper GF. CHARTLINE: providing bibliographic references relevant to patient charts using the UMLS Metathesaurus knowledge sources. In: Frisse M (ed). Proc Sixteenth Annu Symp Computer Applications in Medical Care. Washington, DC: American Medical Informatics Association, 1992;86–90.

38. Robertson GG, Card SK, Mackinlay JD. Information visu-alization using 3D interactive animation. Communications of the ACM. 1993;36:57–71.

39. Cole WG, Sheretz DD, Tuttle MS, Hsu GH, Fagan LM, Carlson RW. Semantic visualization of oncology knowledge sources. In: Gardner RM (ed). Proc Nineteenth Annu Symp Computer Applications in Medical Care. JAMIA. 1995;2, September–October suppl:67–71.

40. Oliver DE, Barnett GO, Chueh HC, et al. InterMed: an Internet-based medical collaboratory. In Gardner RM (ed). Proc Nineteenth Annu Symp Computer Applications in Medical Care. JAMIA. 1995;2, September–October suppl, 1023.

41. Steiner J, Neuman BC, Schiller J. Kerberos: an authentication service for open network systems. In Proc Winter USENIX Conference, Dallas, TX, 1988; 191–202.

42. Lynch DC, Lundquist L, Digital Money: The New Era of Internet Commerce. New York: Wiley, 1996.