B. de Faria Leão

Health Informatics Center,
Federal University of São Paulo
(UNIFESP),
Escola Paulista de Medicina,
São Paulo, Brazil

Synopsis

Knowledge Processing

AI in Medicine - Dream versus Reality

Much has been promised about AI applications in Medicine. Twenty years ago, when the first generation expert systems began to appear, expectations were very high. The belief was that by the 1990s knowledge-based systems (KBSs) would be largely in use. With very few exceptions, KBSs are still an academic issue, only in use where they had been created. Many factors contribute to this, such as: lack of integration, absence of international standards for knowledge representation, and difficulties in learning from experience on conventional KBSs. These issues have already been addressed by several authors [1,2,3].

Information Explosion

As the amount of data stored in electronic media increases, so does the necessity to organize and to retrieve data in a multidimensional approach. The Internet is a good example of that information explosion. Today, there are almost 100 million computers hooked to the net, offering different types of services. Anyone who has ever used Internet knows how exciting this can be and also how difficult it is to find information that is already there. Conventional database application developers are also discussing this issue. Multidimensional queries, business rules, OLAP, data mining and data warehousing applica-

tions are creating a new paradigm on how to interpret, retrieve and analyze data stored in databases. The first commercial applications are already available, offering users sophisticated tools to support better decisions. These are knowledge-dependent applications. In the relational data-base world they represent a natural evolution of the FAQs (frequent asked queries) concept. While conventional queries could only show data in two dimensions, these applications are able to gather and show data in multidimensional views, offering the user different scenarios to better illustrate existing data relationships.

Integration with already existing applications is, therefore, crucial for the flourishing of decision-support tools. This is, of course, easily done in domains with less complexity than medicine, where knowledge is based on few and well determined concepts.

The Medical Workstation Dream

Ideally, medical decision-support tools should be embedded in the medical workstation environment. They should be transparent, supporting the daily medical activities in patient care, research and education. The reminders and alerts should be there at the right moment with all the necessary information and links to other information sources in the network. There is still some place for the consulting tools, but this should also be integrated with the electronic patient

record. Furthermore, these decision-support tools should be able to learn from experience, adapt to the environment and discover new data associations, presenting them in a meaningful context.

To achieve this dream, many AI methodologies must be considered, such as: natural-language processing, neural networks, case-based reasoning and hybrid systems. These methodologies should, however, be considered as software components to be integrated in the already existing in-house systems. The five papers selected for knowledge processing chapter in this Yearbook, address different scenarios for these methodologies.

Natural-Language Processing

Considering that a great amount of data from the patient record is still free text, tools to interpret and formalize these narrative reports are essential. The paper by Hripsack et al. [4] describes an interesting approach where a general purpose natural-language processor is used to encode data from chest radiographs reports. These coded data are then fed automatically to a decision-support tool that aims at evaluating the efficiency and accuracy of the coded data generated by the natural-language processor. Besides the natural-language processor, the decision-support tool was also evaluated, revealing a performance similar

to that of physicians. The paper by Smart and Roux [5] also describes a knowledge-based tool for automatically extracting knowledge from medical reports written in natural-language. Here, instead of using a general-purpose language processor, a specific tool was created that extends the classical concepts of conceptual graphs operations, allowing for graph manipulation and unification. The results are also promising.

Case-based Reasoning

Case-based reasoning (CBR) is a new branch of symbolic AI, based on the theoretical model of the Dynamic Memory [6]. The main assumptions behind CBR are: it uses stored experience from solved cases to solve new ones; it is able to interpret and understand new cases, based on previous ones; through adaptation, CBR systems are able to compensate for the differences between the new and the old situations. This is an extremely interesting methodology for medicine. In fact, CBR is well suited to dealing with fuzzy and complex domains. There are already some isolated medical applications using CBR. The most difficult part in building CBR systems is the definition of the dimensions that describe the case. Based on these dimensions, complex index structures are created, allowing for case retrieval and adaptation. The paper by Evans [7] describes a CBR application to diagnose dysmorphic syndromes.

Knowledge Acquisition

Knowledge acquisition is identified as the bottleneck phase in the development of conventional KBS. There is no cook-book recipe on how to acquire knowledge from experts and translate it into rigid computational formalisms. The paper from Lanzola, Quaglini and Stefanelli [8] addresses this issue, based on an epistemological model of medical reasoning and a

computational model which represents the implementation details to perform a task described by the expert. The main challenge was to create an infrastructure that allowed for multiple intelligent agents to interact when solving problems.

The Connectionist Approach

Since knowledge acquisition is such a complicated and time-consuming task, automatic methods of knowledge discovery and learning are mandatory for KBS development. Some of these methods use symbolic algorithms of machine learning. However, the connectionist approach is being largely used with success in many domains, either isolated or integrated with conventional systems. The paper from Dorffner and Porenta [9] presents an extensive comparison between several feedforward neural network types in the context of a clinical diagnostic task. The main challenge when using a neural-network approach is to choose the neural network architecture and determine which evidence should be given to the model.

One recent and promising approach is to use hybrid systems. "Hybrid Connectionist Expert Systems" have been proposed to solve well-known problems in conventional expert systems. Classical expert systems provide an adequate methodology to represent knowledge and to explain the conclusions reached; however, learning and knowledge acquisition still remain a problem. On the other side, connectionist systems attract increasing interest for their inherent learning and pattern recognition capabilities. Hybrid connectionist expert systems were, therefore, proposed, integrating symbolic and connectionist paradigms, both supporting each other. The symbolic paradigm increases the semantics of the stored knowledge, simplifying its manipulation and comprehension. The connectionist paradigm, on

the other hand, offers the necessary resources for knowledge acquisition and refinement [10]. Hybrid systems offer several advantages and show impressive performance in several domains, including Medicine [11].

Distributing and Integrating Knowledge

All five papers present [4,5,7-9] interesting solutions to build decision-support tools for different medical contexts. All of them reflect a major research effort with sophisticated manpower and computational infrastructure behind them. All of them have embedded medical knowledge, either automatically extracted from free text, encoded in weight connections in neural networks, described in cases, or extracted directly from the expert. Which are the available methodologies and standards that could allow for the sharing of these KBS or part of them? The idea of having a software warehouse of components and services that could be shared in a distributed environment is fascinating. The CORBA standard from the Object Management Group (OMG) [12] offers the necessary infrastructure to achieve such interoperability. This paradigm is based on a distributed object environment with services and brokers to handle integration with different and heterogeneous applications. If this architecture is used to develop decision-support tools for medicine, integration with the EPR could be achieved. In fact, this has already been considered by the new Domain Task Force of the OMG, called CORBA-med, with the goal to achieve interoperability in health care.

The Arden syntax is a known available standard to write medical rules that allow knowledge sharing with different applications and institutions. MLM modules could easily receive an IDL layer (Interface definition language) and enter the CORBA world.

Also, knowledge that is now encoded in languages such as LISP and Prolog could have an IDL description, allowing for its integration with other applications.

Apart from the complexity of representing medical knowledge and adapting it to each local context, one must not forget that this effort has to go in parallel with interoperability features. The CORBA standard is an up-to-date methodology that allows this integration. If developers of medical decision-support tools also consider this issue while designing their systems, it will not take long to have a library of medical objects that could be shared worldwide. Each object would be an intelligent agent with a proper interface to communicate with others, encapsulating the details of its implementation. Software development is becoming an exercise of higher abstraction that demands knowledge and, above all, has to incorporate adequate interfaces to communicate with others. A distributed scenario with intelligent agents seems to be the obvious direction to go. In this scenario, the challenge is to build proper interfaces

to communicate with other agents. Standards are, again, the hot issue, not only to represent medical knowledge, but to communicate what is being done with this knowledge. For the medical informatics community there is a long way to go. The future of useful decision-support tools depends on our ability to develop applications that may participate in this new scenario of distribution and interoperability.

References

1. Shortliffe EH. The adolescence of AI in medicine: Will the field come of age in the '90s? *Artif Intell Med* 1993;5:93-106.
2. Uckun S. Artificial intelligence in medicine: State of the art and future prospects. *Artif Intell Med* 1993;5:89-91.
3. Safran C, et al. Effects of a knowledge-based patient record on adherence to practice guidelines. *MD Comput* 1996;13:55-63.
4. Hripcsak G, Friedman C, Alderson PO et al. Unlocking clinical data from narrative reports: A study of natural language processing. *Ann Intern Med* 1995;122:681-8.
5. Smart JF, Roux M. A model for medical knowledge representation application to the analysis of descriptive pathology reports. *Meth Inform Med* 1995;34:352-60.
6. Kolodner JL. Maintaining organization in

a dynamic long-term memory. *Cogn Sci* 1983;7:243-80.

7. Evans CD. A case-based assistant for diagnosis and analysis of dysmorphic syndromes. *Med Inform* 1995;20:121-31.
8. Lanzola G, Quaglini S, Stefanelli M. Knowledge-aquisition tools for medical knowledge-based systems. *Meth Inform Med* 1995;34:25-39.
9. Dorffner G, Porenta G. On using feedforward neural networks for clinical diagnostic tasks. *Artif Intell Med* 1994;6:417-35.
10. Kandel A, Langholz G, eds. *Hybrid Architectures for Intelligent Systems*. Boca Raton: CRC Press, 1992.
11. Leão Bdf, Reategui E, Guazzelli, A, Mendonca EA. HYCONES: A hybrid approach to designing decision support systems. *MD Comput* 1996;13:160-4.
12. Object Management Group. *The Common Object Request Broker: Architecture and Specification*. New York, John Wiley & Sons, 1992. (OMG Document 91.12.1)

Address of the author:

Beatriz de Faria Leão, MD, PhD,
Health Informatics Center,
Federal University of São Paulo
(UNIFESP),
Escola Paulista de Medicina,
São Paulo, Brazil.
E-mail: bleao.cis@epm.br