

## Sharing Medical Knowledge for Automated Decision-making

Paul D. Clayton, T. Allan Pryor, Ove B. Wigertz, Stephen B. Johnson, George M. Hripcsak

Center for medical Informatics, Columbia Presbyterian Medical Center NY, NY;  
Dept. of Medical Informatics LDS Hospital/University of Utah, Salt Lake City,  
Utah; Dept. of Medical Informatics, Linchoping University, Linchoping Sweden

**Abstract:** Medical knowledge in a computer processable representation can be shared. The issues involved in distribution include: 1. the representational format (syntax) of the knowledge, 2. availability of a comprehensive editor, 3. development of convenient methods for translating the distributed knowledge base into the native language and database environments which vary from institution to institution, 4. quality control, 5. the need for users to review and assume responsibility for the content of any logic which they obtain, 5. willingness of authors to share knowledge, 6. willingness of institutions to act as clearing houses (libraries) for such collections of knowledge, 7. techniques for knowledge base management, and 8. liability. We describe some of the problems and potential solutions associated with each of these issues.

It has been argued that the primary source of the power of the new generation of medical decision-making systems is the knowledge contained in the knowledge base. A major problem is the fragmentary and incomplete nature of this knowledge base. Although different investigators and institutions have all contributed substantial and useful segments of the medical knowledge which is in a computer useable format, it is generally agreed that no single group will easily complete the entire spectrum of medical knowledge. It is also evident that no single system will emerge as a de facto standard for implementing automated medical decision-making. Medical logic for a wide variety of decisions has begun to emerge (diagnosis, contraindication alert, management suggestions, management critiques, data interpretation, monitoring alerts, test suggestions, formulation of questions to be asked etc.).

In light of these observations it appears that the best solution is to create mechanisms which will allow knowledge bases to be transported among institutions in a uniform representation (A uniform language for medical decision-making). Such mechanisms would enhance the breadth of applications by allowing a wider variety of authors to contribute medical logic and by enabling an institution or practice to gain access to this

store of medical expertise. This approach would also allow common protocols for testing, validating, and managing the resulting knowledge base. The knowledge would be distributed in modular, ASCII coded output which will be suitable both for direct scrutiny by medical personnel who are not required to have programming skills and for translation to one of the supported systems.

Medical knowledge in a machine processable representation should also emerge as a suitable format for library collections just as the knowledge is today stored in books and journal articles. Multiple authors could contribute to and criticize the contents of the medical knowledge base, libraries could better manage the knowledge, and efforts for validating and evaluating the quality of the knowledge base could be coordinated. Commercial vendors would be expected to support the standard and thereby broaden the base of users and contributors.

### Representational format.

One attractive representational format is modular, procedurally oriented frames with slots for defining variables used and the logical criteria for defining those variables. A crucial issue involves the question whether the modular format with combinations of procedural and declarative knowledge that this unified representation would imply, would be sufficiently flexible to accommodate enough of the knowledge from existing and developing medical decision-making systems. We recommend the frame structure because the majority of currently available logic models (e.g. DXplain, QMR, HELP, CARE,) can be incorporated within this structure and because frames for suggesting diagnoses, alerts, and suggestions for testing and therapy as well as critiques of therapies may be incorporated within a single representation and editor. This approach is probably insufficient to accommodate the current research work on causal networks, but the corpus of currently applicable knowledge and the fact that so much can be accomplished within this representation appear sufficient to justify the modular approach at the present time.

### The availability of editors for creating frames in a unified syntax.

As part of the representational format for frames, it is desirable to develop a uniform, high-level (ASCII) syntax for representing medical expertise used in decision-making that can be used by a wide variety of contributors. This syntax could be distributed via a widely available high level editor. This editor should run on a variety of personal computers. Such editors (stand alone PC versions) are currently being written for the HELP and QMR representations and at Columbia we have an operational prototype editor which is not model specific but which will accommodate multiple models.

### The development of translators.

A broadly available editing capability would allow many individual users to enter her/his expertise into the medical knowledge base in the ASCII level representation but in a format that would allow eventual maps/links to a particular system's native data dictionary e.g. the use of the word "cough" in a particular module of decision logic must be mapped into the way information about "cough" is stored in the particular system in which the decision logic will be evaluated. Evolving capabilities derived from the uniform medical language efforts would facilitate this interactive mapping process. For each of the variety of useable systems in which decision-making capability would be implemented, there would need to be a translator/compiler written to convert this high level representation into machine executable form. In the worst case, users could manually translate the knowledge using an application specific editor. Each user desiring to gain access to the library of medical knowledge would have to write only one translator which mapped to the high level language rather than one for each of the other systems which contained substantial amounts of knowledge. The transformed, mapped logic is then compiled into the object language for the native environment.

### Quality control.

The quality of the contributed knowledge base must be evaluated and validated. It will be necessary to assess the quality of decision frames and to compare differential diagnosis lists generated by different models. Each author/contributor must assume some responsibility for the quality of a contributed frame by providing appropriate scholarly references to justify the logical criteria. The author should also give some background about the previous use of the frame. Each contributor will specify whether a particular frame has been routinely used in

clinical practice or whether it has been newly developed. Institutions with on-line clinical databases could run the frames against retrospective databases and report their experience. Rather than saying "This frame is good... This frame is bad", it may be possible to say "In our patient population we found that 80% of the alerts generated were judged to be appropriate." The problems of false negatives remains an issue, but some databases may be developed in which gold standards are available.

Potentially, a count of the number of institutions or practices which have instituted the use of the frame could also be reported. Such a directory would have interesting ramifications for notification of a product defect; all of the users could be notified. However unless some highly automated knowledge base management techniques evolve, the reporting aspects of this scenario tend to make it unlikely in the near future.

For diagnosing diseases, a frame usually calculates some type of a score (Bayes, sum of weights, etc.) for each disease and then ranks each disease to form a list of differential diagnoses. For example, even similar models may use different scoring algorithms so that comparison between lists is not straightforward. We propose the following method for generating a list of differential diagnoses even though the scores for the diseases may be calculated via frames which use different scoring algorithms. Based upon a large clinical database of patients with and without a particular disease, one could generate a histogram for each frame which showed the frequency (vertical axis) with which each score (on a normalized scale of zero to one hundred) was obtained for populations of patients with and without the disease. In a particular patient, one could then use the frequent distribution of normalized scores or odds ratios (rather than the algorithmic score) to rank the differential diagnoses. Using these histograms, it would also be possible to construct ROC curves (Metz 78) to compare the quality of a particular diagnostic frame to other frames even though they did not use similar scoring algorithms.

### The need for end users to review the logic.

Each frame would be distributed to a requesting institution in a non-executable ASCII format. In order for the frame to be incorporated into a clinical information system with a data-driven evoking mechanism, there are certain user based decisions which must occur in addition to the database mapping function. If the logic of a frame suggests that a potassium supplement be given when the patient is receiving digitalis and also has a "low" serum potassium level, some qualified

institutional representative must decide what the threshold value for hypokalemia is in that particular practice or hospital.

This qualified individual or group of individuals must also decide which of all the potential frames should be included in the knowledge base for the institution. In our experience, pharmacists decided which pharmacy alerts should be generated, radiologists decided which criteria to use when suggesting that an intravenous pyelogram should not be performed because renal function is "poor", etc. This issue of responsibility and oversight becomes a bit more tricky when the logic gets into diagnoses which span several specialty areas of medicine.

#### Willingness of knowledge base authors to share.

It is impossible to speak for all authors in this matter. One can only suppose that there will be a variety of opinions. Some have declared openly that they are willing to share their knowledge base provided that the liability issues can be solved. Others may feel that a knowledge base is tantamount to software (Homer Warner coined the phrase "medical ware") and should be bought and sold. Authors do receive royalties for writing textbooks--one of the traditional ways of distributing medical expertise. Some authors may want royalties every time their frame is run. If the benefits of the program can be shown to outweigh the costs of the software, there is no reason that people would not pay for the knowledge base. We prefer distribution which includes only the cost of collecting and disseminating the media especially during the development phases of the industry.

#### Willingness of institutions to act as libraries or clearing houses.

It appears that there are several groups who are willing to try and distribute knowledge bases. The AMA supports access to DXplain over AMA net but does not distribute the software directly. This may be so that they can maintain the quality control without restricting access to the use of the system. The COMDAT Foundation is involved in placing versions of QMR in hospitals for evaluation purposes. Drs. Miller and Myers have also distributed copies of the INTERNIST1/QMR database to investigators who have collaborative research interests. One assumes that professional organizations such as the American College of Physicians and the American Medical Association are studying the potential roles that they might play in the eventual distribution and/or certification of the logic in these knowledge-based systems.

In the absence of agreement by the National Library or professional organizations to distribute the knowledge base, the other opportunities appear to be some confederation of Universities or private commercial interests. Columbia University, LDS Hospital/University of Utah, and Linchoping University in Sweden are trying to cooperatively establish a standard format and build a library of knowledge frames which will be available for distribution.

#### Knowledge base management.

The practice of knowledge base management is in its infancy although individuals have started to address the issues. How do you know that by putting one new frame in the knowledge base, you won't invalidate the results of other previously tested logic. This problem was especially difficult in the production rule oriented systems and required special programs to be written within the MYCIN system (Suwa, Scott and Shortliffe 1982). This problem is solved to some extent by the modular nature of the frames as opposed to the control issues involved in a production system.

Suppose an author of the frame for diagnosing hepatitis updates the weights used in the scoring algorithm or adds a new test. How do all the users of the old frame get notified that there is now a "new, improved" version of the knowledge base. This appears to be a sticky issue, but in essence is no different than the issue of new editions of textbooks. All previous subscribers would probably want to get updated on a periodic schedule than receiving change by change reminders. Others might respond that this periodic update approach defeats the whole purpose of having electronic communication rather than printed material. The issue will probably best be resolved by letting the recipient decide how rapidly to respond to the changes and whether there is truly on-line electronic communication or whether the knowledge gets sent around on tapes and floppy disks.

#### Liability issues.

Perhaps the largest current obstacle to the distribution of knowledge bases is the liability which may be associated with giving advice to potentially thousands of users. As far as the Food and Drug Administration is concerned, there appears to be a distinction between devices that offer advice to a competent human recipient and those which directly affect the patient. They have apparently decided to license only those devices which directly feed an action back to a patient without human intervention. What the FDA decides for regulation of medical devices will probably not apply to potential civil suits. It has been suggested that a

data-driven evoking model (e.g. HELP, CARE, etc) are different than passive stand alone systems which the user must actively seek for assistance. The former systems make active suggestions even when the user does not ask for assistance and therefore may be treated differently than the latter category which could justifiably be defended by saying that the expert system is the same as a textbook or a library and the reader must be responsible for sorting out the quality which differs among printed materials.

By reporting only experiential based results of individual decision frames instead of pronouncements, an institution can give others some feeling for the quality of a frame without accepting full responsibility for the content of the frames. Karen Harris Keeter (Keeter 88) has suggested that any commercially available knowledge translators/compiler "have a key which must be activated, before the ASCII level distributed text could be compiled into the native environment. This key would be a slot which gives the name of an institutional representative who speaks for a committee or individual who reviewed the logic in a frame and accepted it as appropriate for that institution and the date of this review." Such a key would insure that someone in each institution assumed institutional responsibility for the fact that the logic used meets current professional standards.

The authors and distributors of the knowledge base may be able to minimize their exposure by including some type of coverage under the current terms of their malpractice liability, writing logic which reflects the current standard of care, and requiring that recipients of the knowledge base obtain the logic in ASCII format which is tantamount to receiving a book or journal. The institution then becomes the entity which took the available published knowledge and converted it into an active data driven advice giver. The distributing organizations, by performing some level of quality control, can help the user institutions decide which frames to implement.

Using these defenses would help insulate the contributing authors and the distributors of knowledge bases, but would not absolutely preclude the possibility of a suit. In today's product liability atmosphere it is common that everyone connected to the product will be named as a defendant even though there was no identifiable negligence on the part of certain participants.

It is however interesting to note that at the same time potential liability problems are hindering the dispersion of medical knowledge in this format, a recent court ruling has found that a physician can be found negligent for not consulting available electronic aids. In this case the court said "... With the

demand of their profession, no one can expect doctors to have all material information stored in their minds... A literature search will put a physician on notice of these risks" (Pemberton 1986). Based upon the experience of the users of the HELP system in Salt Lake City, it was the general consensus of users that the reminders generated by the system reduced rather than increased the liability for the user. Hence from the user's and institution's point of view, it appears that one must balance the risk between using an imperfect source of assistance versus using nothing at all. We assume that many will choose the imperfect source of assistance and face their attendant risk in order to reduce the overall liability which an institution may face.

**Summary.** In spite of the potential obstacles which have been discussed, we think that the alternatives are even more distressing. In essence each laboratory would continue to independently develop its own knowledge base and the ability of non-development sites to acquire medical knowledge would be limited by the capabilities of the particular system which they may have purchased. Duplication and the inability to share would severely hamper the development of the field.

#### References

- K.H. Keeter, Atlanta, Georgia. personal communication Feb 1988.
- P.R. Manning, D.W. Petit, "The Past Present and Future of Continuing Medical Education," JAMA 258:3542-3546 1987.
- C.E. Metz, "Basic principles of ROC analysis" Seminars in Nuc. Med 8 (no 4):pp283-298. 1978.
- S. Pemberton, "Physician responsibilities in an information society," Health Care On-line 2(#2):1-3. Nov 1986.
- M. Suwa, A.C. Scott, E.H. Shortliffe, "Completeness and consistency in a rule based system," The AI Magazine 3:16-21 (autumn) 1982.