# Issues and Structures for Sharing Medical Knowledge among Decision-Making Systems: The 1989 Arden Homestead Retreat

Paul D. Clayton, T. Allan Pryor#, Ove B. Wigertz*, and George Hripcsak

Center for Medical Informatics, Columbia Presbyterian Medical Center, New York, NY 10032
#Department of Medical Informatics, LDS Hospital/University of Utah, Salt Lake City, UT 84143
*Department of Medical Informatics, Linkoping University, S581 83 Linkoping, Sweden

## Abstract

To address the issue of facilitating transfer and integration of the variety of computer-based programs which contain medical expertise, a retreat was held at Columbia University's Arden Homestead conference center June 16-18, 1989. The focus of this retreat was to explore ways in which the medical expertise contained in knowledge-based systems could be shared and expanded. During the three day meeting, the eighteen attendees from ten institutions discussed: (a) the need for better ways of mapping terminology used in one setting or program to terms with similar meaning that have been used in other programs, (b) the need for catalogues which list the variety of programs which are available, (c) a representational syntax and format for sharing modular medical knowledge, (d) the possibility of developing standards for interfacing program modules so that they could be "snapped" into place in a variety of systems, (e) methods for evaluating, validating and testing knowledge based systems, and (f) the legal and financial aspects of sharing systems which influence the care that is given to a patient.

We emerged from the retreat with a feeling that there was an enthusiastic but not unanimous consensus that sharing should occur in order to advance the field of medical information systems. We accepted an initial version of a working document for the representation of Medical Logic Modules (MLM's), appointed leaders for subcommittees to address the issues which had surfaced and settled upon an approach for dealing with the legal and financial aspects of the sharing process.

## Introduction

It is sometimes discouraging to see how much effort has been expended to build even a small part of the total variety of exciting and beneficial capabilities in medical information systems which have emerged during the 70's and 80's. We recognize that it is unlikely that any institution will ever internally and independently generate all of the capabilities which have been demonstrated or conceptualized.

The last fifteen years have seen the development and acceptance of knowledge based systems in medicine. The effort involved in authoring and testing the knowledge contained in the knowledge bases of various systems which have been developed is prodigious. The first logic for the HELP system was written in 1972 [1,2] and seventeen years later there exist about 5000 logic modules for that system. There are about 600 disease profiles which have been developed over approximately the same time span for the INTERNIST-1/QMR effort [3]. The DXplain project [4] has 2112 disease profiles which have been developed and tested over a five year period. The CARE system [5] at Indianapolis' Regenstrief Institute contains the logic for making approximately 1400 decisions and has been under development for approximately 16 years. There are many conceptual similarities among the knowledge bases in these systems, but the syntax for each knowledge base is different.

Our thesis is that the intellectual work of the authors of these knowledge bases need not be duplicated, and that to improve the breadth of coverage of the knowledge bases, more authors must be enabled to contribute to the effort. The Pittsburgh group has proposed a mechanism for expanding the QMR database using many contributors [6]. To facilitate the goal of sharing and transporting knowledge between institutions, we proposed [7] that a common structure be developed which would accommodate the logic of existing knowledge bases which are represented in a modular fashion.

Sharing does not necessarily mean that the knowledge which an institution or group has laboriously translated to a computer representation should be given away gratis, but does mean that the work of one group can be transferred to another environment.

Given these pressing needs and a structured agenda, a small, but representative group met at Arden Homestead, Harriman, New York. The retreat was sponsored by the CAMDAT Foundation, IBM, and the law firm of De Forest and Duer.

## Overview

The overview session of the retreat was lead by Ted Shortliffe and Jan van Bemmel. We discussed the benefits and obstacles to sharing and soon realized that we must talk about the representation of the knowledge that was to be shared. The first issue that arose was whether knowledge bases should be classified according to whether they were declarative or procedural in nature. It is possible for the expertise in a knowledge base to be used for multiple purposes other than the one for which it was first developed [8, 9, 10]. But most in attendance agreed that because the knowledge bases generally were



0195-4210/89/0000/0116$01.00 © 1989 SCAMC, Inc.

developed with a specific purpose in mind, the representation was strongly influenced by the application. This procedural influence made sharing across applications more difficult.

The CYC project [11] being lead by Douglas Lenat at MCC in Austin Texas was discussed. In that approach only declarative facts are taken from an encyclopedia and the inference engine is very general. We discussed the fact that there are no currently existing purely declarative databases in the medical area and because of completeness and other problems, it was not productive to classify systems according to whether they were procedural or declarative.

We settled on a classification scheme that was more agreeable to the group and which addressed the degree of modularity in the representation of the knowledge base. Here the term "modular" means that one may add the logic for making a single decision to an existing knowledge base in an incremental fashion. We broke medical decision-support systems into three broad classes.

The first class, called modular independent, is exemplified by the HELP or CARE systems, in which a single rule can be fired reminding someone to take a pap smear or that an aminoglycoside prescription is contraindicated because of a laboratory value which indicates that a patient has poor renal function. Even if no other rule existed in the knowledge base, these independent modules could be evoked and evaluated and their results displayed, appreciated, and acted upon.

The second class, called modular interdependent, includes DXplain, HELP, ILLIAD and QMR/Internist-1 programs in which the disease profiles are constructed independently but must be congruent. For example, if the ad hoc or pseudo-probabilistic score which is calculated to reflect the likelihood that a patient has a common cold is consistently be higher than the calculated scores for alternative diagnoses, then all patients will be classified as having a cold.

The last class is exemplified by some of the production systems and causal or semantic nets. Approaches such as HT-Attending [12], Pople's continued development of the INTERNIST-1 system [13], the CASNET/Glaucoma system for diagnosis in ophthalmology [14], and the pioneering work by Shortliffe in the MYCIN system [15], are based upon a set of interdependent rules or semantic links between entities or objects. In this approach, fragmentary knowledge severely restricts the performance of such systems, which for want of a better word, we will call interwoven. The rules in these systems may be modular but the need for completeness [16] sometimes means that the deletion of a single rule can ruin the performance or accuracy of an entire system. This final class also contains systems in which the knowledge is not modular but which still contain medical expertise. Examples of these systems are many of the computer-aided instruction programs.

We decided that it is definitely possible to have multiple individuals construct, evaluate, critique, and transport medical logic which is in the modular independent class. We also felt that because of completeness issues it would only be possible to transport systems in the interwoven class as intact systems. Modular interdependent knowledge, such as disease profiles for differential diagnoses, could be shared in a modular form but there was no great desire to promote a common standard for this type of system since there is currently a huge existing body of knowledge in a native representation.

Another consensus was that the full power of the programs which have been developed over the previous two decades would best be realized by interfacing them to existing clinical databases. In fact it was felt that the opportune time to move existing interdependent logic to a more transportable standard would occur when those systems were connected to clinical databases.

This consensus about connections to clinical databases led to the next obvious conclusion which had already been recognized and addressed in part by the National Library of Medicine - the need for a uniform medical language or at least the ability to map between terms in a clinical database or knowledge base and concepts in another system that were described by slightly different terms.

The final topic which we recognized as fundamental was the need to evaluate, validate and test knowledge bases. It is not desirable that everyone repeat the work of validating a knowledge base, but our discussion the following day showed that there was little consensus on how this validation should be performed.

We emerged with six definite areas in which we felt that we could make progress in sharing: (1) Increase the focus on ways of implementing the emerging Unified Medical Language System (UMLS), and create ways of mapping between the vocabularies of different systems. (2) Build a catalogue of existing systems so that interested parties could find out about existing efforts before re-inventing the wheel. (3) Pursue a standard syntax that would support creation and transport of modular independent medical expertise. (4) Set up a task force to study the ways of seamlessly sharing complete programs and systems. (5) Mount a focused effort to build upon ongoing efforts to test and validate medical knowledge. (6) Apply some common principles in dealing with legal and financial aspects of the sharing problem.

In the sections that follow, we shall report in detail on each of these issues and, when possible, the conclusions reached by the group. It should be recognized that over a three day period, the group covered a large amount of ground; it is impossible to give the intellectual credit to each individual during the ensuing paragraphs. In our summary, we occasionally recognize strong adherents and champions as some of the principles are discussed but recognize that there were many extremely valuable comments which we summarize without attribution. The authors of this paper wish to give credit for the content of the paper to the collective wisdom of the group.

## Vocabulary Definition

It was the consistent opinion of those who had been involved in sharing any degree of knowledge (Octo Barnett, Randy Miller, Ted Shortliffe, Homer Warner) that the understanding of the terms involved in describing medical expertise was the most basic concept which had to

be addressed before sharing could occur. It was noted that Jack Myers had to exert great effort to force the vocabulary in INTERNIST-1 to be consistent across diseases. The work of Roger Shank was also mentioned in which it was found that word for word translation did not work unless there was a conceptual framework present.

The problem of terminology presents itself when trying to define a disease. Because of the different temporal progression and symptom manifestation among patients, a single disease classification is not sufficient to describe their condition or determine their prognosis.

An example is the use of the phrase chest pain: does it mean crushing, substernal pain relieved by rest, i.e. anginal chest pain, or does it refer to sharp, pleuritic pain made worse by coughing? These two concepts have major implications for management and diagnostic logic and cannot be described by the same phrase. What would happen if the patient database simply stores a symptom under the heading "chest pain"?

Another example concerns negative manifestations? Can you infer from the fact that someone has stated "normal chest" that there is no tumor, or pulmonary hypertension, etc.? This leads into the whole area of default reasoning in which you assume some default value unless you are lead to believe otherwise.

The work of the National Library of Medicine has now produced a standard vocabulary which is to be demonstrated at the 13th SCAMC. All systems will not initially adhere to this controlled vocabulary but this standard will provide for a one to many mapping rather than having HELP map to QMR and DXplain, QMR map to DXplain and HELP, etc. Randy Miller was appointed to head a sub-task force to follow up on this discussion and develop suggested standards which were consistent with the needs which had been expressed and the UMLS work supported by the National Library of Medicine.

## A Catalogue of Existing Decision Support Programs

The first step toward sharing is knowing what else is available. Homer Warner had earlier [17] recognized that there are more decision support systems in existence than most people are aware of and started a catalogue of decision support technology. The group felt that this was a fundamental resource that should be expanded and maintained. Issues involved include whether to put something like this on-line so that people have electronic access to the catalogue. How would one support the creation and maintenance of such a resource? It was also questioned whether a professional organization such as the American Medical Association, the American College of Physicians, the American Medical Informatics Association, or the International Medical Informatics Association, or an educational institution, the National Library of Medicine, or a private foundation should best nurture the project. As Homer Warner had already embarked upon such a project, we felt that he should lead the task force which would further investigate and develop this capability.

## A Syntax for Sharing Modular Independent Logic

The need for a representation of medical knowledge that is not limited to one application or one institution was clearly recognized at the meeting. The use of object-oriented systems and expert system shells was discussed. In order to get started, however, the current standard focuses on independent modular knowledge. There has already been success in this area and there were several groups at the meeting who were interested in adding this capability to their clinical information system. The extension of the syntax to allow interdependent modular systems like QMR to access a patient database is currently under study. The extension of the syntax to allow the sharing of knowledge between interdependent modular systems (e.g., sharing disease profiles) is also under study, but it is recognized that this is problematic because of differences in inference engines and the use of constructs like properties.

A preliminary document describing the standard was distributed before the meeting, and a synopsis of the standard was presented at the meeting. The document describes the syntax and some of the semantics of units of medical knowledge called "medical logic modules" (MLM's). The term MLM was created because it is without prior connotation and expresses the fact that the field is medicine, that the MLM represents logic rather than vocabulary, and that the knowledge base is modular. Modules of medical knowledge have been referred to as "frames," "sectors," "profiles," "objects," etc, but these terms have been avoided for a number of reasons. For example, the word "frame" is concise, but it has created many misconceptions while communicating with people outside of the medical informatics field (e.g., vendors), since the word frame already has a long history in the AI world.

To be useful for sharing knowledge, the MLM standard must have several properties. (1) The MLM's should be able to be read and interpreted easily by medical experts with little computer training. (2) The MLM's should be unambiguous, so that the same module cannot be interpreted in two different ways. (3) The MLM's should be able to be written by medical experts with as little training as possible. Note that this is a different property than the readability. For example, it is easier to recognize a list of numbers like "1, 2, 3;" than it is to remember that the list is delimited by commas and terminated with a semicolon. (4) The MLM's should allow computer translation and maintenance. In fact, there is a trade-off between ease for the user and ease for the computer. While a standard that requires the user to write machine code will be a failure, so will a standard that allows the user to write natural language. The standard must be practical to both the user and the computer.

Broadly speaking, an MLM is broken up into discrete slots, each of which serves some particular purpose (Figure 1). A slot is composed of a slot name that serves to identify it and a slot body that contains information. The slots are grouped into three categories: maintenance, library, and knowledge.

The maintenance category contains those slots that specify maintenance information that is unrelated to the medical knowledge in the module. Its slots identify the name of the MLM, the author, the institution, the date it was created, etc.

The library category contains those slots pertinent to knowledge base maintenance that are related to the module's knowledge. Its slots identify the MLM's purpose, keywords, citations, and links to other knowledge sources like an electronic textbook.

The knowledge category contains the actual medical knowledge. The "data" slot defines medical terms used in the other slots. As the UMLS progresses, this slot should become simpler. The "evoke" slot can be seen as setting the context for a medical fact. The "logic" slot contains the fact itself, and the "action" slot contains the conclusion to be drawn if the fact pertains to the patient. The "logic" slot also allows more complicated calculations to be done.

The MLM document also describes the treatment of time and the data model. The syntax allows one to specify search intervals for database queries (e.g., "between 1 day before timeof(surgery) and 3 days after timeof(surgery)"). A relational data model is used, where rows represent instances in time (e.g., retrieving the last 3 platelet counts for a patient results in a column with 3 rows), and columns represent attributes of some entity.

At the present time the syntax is being implemented at LDS Hospital, Columbia-Presbyterian Medical Center, and Linkoping University with additional plans for testing the standard for completeness, ambiguity, and transferability at several other institutions. A current version of the document is available from George Hripcsak.

---

```
maintenance:
  title:       Agranulocytosis and
               Trimethoprim/Sulfamethoxazole;
  filename:    anctms;
  version:     1.0;
  institution: Zoo University;
  author:      Dr. Bonzo;
  specialist:  ;
  date:        7/20/1989, 8/7/1989;
  validation:  research;
  format:      USA;

library:
  purpose:     To display the Arden Homestead
               MLM Standard;
  keywords:    granulocytopenia; agranulocytosis;
               trimethoprim; sulfamethoxazole;
  citations:
     1. Anti-infective drug use in relation to the
     risk of agranulocytosis and aplastic anemia.
     A report from the International
     Agranulocytosis and Aplastic Anemia Study.
     Archives of Internal Medicine, May 1989,
     149(5):1036-40.
  links:   CTIM .34.56.78;
           MeSH agranulocytosis/ci and
             sulfamethoxazole/ae;

knowledge:
  type:        data-driven;
  data:
     anc := last 2 of(select
       absolute_neutrophile_count from pt_db)
       between now and 1 week before now;
     last_anc := last(anc);
     previous_anc := first(anc);
     tms := exist(
          select current_meds from pt_db where
          current_meds =
          trimethoprim_sulfamethoxazole );
  evoke:       on storage of anc;
  logic:
     IF tms and (last_anc < 1000) and
       (last_anc < previous_anc) THEN
     conclude(true);
  action:
     send "Caution:  patient's relative "||
       "granulocytopenia may be exacerbated"||
       "by trimethoprim/sulfamethoxazole."
     to user;
```

Figure 1. Sample Medical Logic Module

## Sharing Intact Systems

Randy Miller stated that modular declarative knowledge bases are passive objects; they require procedures to make something happen. The next question is whether we can mix and share passive modular knowledge base contents? The answer is probably yes since that is what we currently do with books. Can we mix procedures that manipulate the content of passive modules? The answer to this question was given to some extent by Paul Mongerson (Camdat Foundation) in comments before the meeting and by Bob Greenes when they strongly proposed that we need to share not just the medical logic modules, but also entire medical decision support systems. There are several options. (1) Use existing programs in their current form, letting users enter and leave them using conventional means like DOS or menu shells. (2) Use a shell program that provides a uniform user interface to existing programs. (3) Alter existing programs and write new ones so that they can be called as black boxes using a parameter list (which must include things like graphics and window displays) to transfer data. This would avoid having the user re-answer questions already asked by other programs.

The last option of having "black boxes" which conform to standards is the most desirable and Bob Greenes will head up a task force to keep momentum moving. Jan van Bemmel and Teun Timmers of Erasmus University are sponsoring a software engineering conference in Sept 1990 which will attempt to address some of these topics. There is some interest in having a similar meeting in the USA to address medical software.

## Validation of Knowledge Bases

In his opening comments Jan van Bemmel talked about the need to validate knowledge bases that might be shared. This is important if a hospital wants to use some MLM's or a complete

program such as DXplain, ILLIAD, or QMR. If a hospital imports a system, it must either evaluate the system itself or assess the amount of external evaluation the system has received. The original comments focused on the use of comprehensive databases against which the knowledge base could be tested. The analogy to ECG testing was discussed by several who had actually done the ECG comparisons. Given that one might have to have up to ten or twenty cases of every disease in order to test performance on non-classic patients, and that much extraneous information would have to be available which did not pertain to the disease which the patient had, it was projected that the test database might have to consist of up to 50,000 extraordinarily well documented patients in order to test the current state of diagnostic programs.

The problems of testing management MLM's were also discussed. If one uses an MLM for drug contraindications in a large population, one can go to the patient charts to see if the alerts generated by the MLM were true or false. The problem is finding those patients for whom the MLM inappropriately failed to generate an alert ("false negatives"). It is very difficult to find the false negatives because there is no way of knowing about them unless they are specifically recognized or constructed by some other method. There is still no guarantee that the MLM will not produce a false negative for other patients about which we are uninformed.

Others in the room felt that the best way to test a system was similar to what we do with people. Most well formulated educational examinations do not exhaustively ask a student about every detail she is expected to know, but rather start with simple cases and move progressively to more difficult ones knowing that somewhere there must be a boundary to one's knowledge. Since these boundaries are not uniform in every area, the exploration should occur over a range of subject matter. The benefits of a statistical sampling approach compared to the exhaustive testing of all possible combinations were discussed in a recent editorial [18]. Some espoused a somewhat less rigorous, but more pragmatic approach in which it was suggested that testing can occur simply by letting a variety and large number of knowledgeable individuals use the program with an awareness that it is imperfect and solicit their feedback. Once the recipients have a good feel for the capabilities of the system and have a strong awareness that the system is not expected to be perfect, it should be used. There is a lot of good that can come from imperfect systems (or people).

Larry Kingsland related that he had for some time endeavored to address this problem and that the National Library of Medicine had awarded a contract to consultants to address the issue. They were expected to deliver a completed report in the near future which might be the basis for testing diagnostic systems. We asked Dr. Kingsland to head the task force that would pursue this work further; but there was no consensus at the meeting that a standard for comparing, evaluating, or validating knowledge bases exists today. We hope that such standards can emerge and that developers can adhere to those standards and demonstrate that they have been prudent in their efforts to assess the quality of any knowledge bases which might be shared.

## Financial and Legal Aspects

The financial aspects of sharing were discussed. It was pointed out that many sites had spent years and thousands of dollars developing knowledge bases and that they needed income to perpetuate their work. The analogy to textbooks was drawn in which authors do receive royalties. The analogy breaks down with knowledge in an electronic format; its very strength - ease of access - contributes to the problems of controlling royalty usage.

The amount of money spent annually for malpractice insurance was mentioned and it was realized that hospitals would not be deterred from investing in medical decision logic if they could see that the number of claims could be decreased. Saving one or two claims a year in 100 major hospitals would provide a large amount of money to support the development of medical decision logic. Thus it was decided that sharing did not mean that expertise would be transferred without cost and that if the expertise was worth it, there would be a willingness to pay. Individuals and institutions would have to set pricing policies in an independent fashion.

The law firm of De Forest and Duer had researched the liability issues associated with sharing medical expertise that might be used as the basis for altering the treatment and hence, the outcome, of patients in institutions which employed medical decision making systems. Messrs. Tom Kusack and David Stein of this firm donated their time and efforts to prepare a presentation for this retreat. Our summary does not reflect legal counsel, but conveys the general inferences that the group made as a result of the lively discussion which ensued.

The bottom line to a recipient institution was presumed to be favorable. Based upon the requirement that a human must always accept ultimate responsibility for treatment received and the prospect of improved quality of care, it was felt that the decision making systems would avert more liability judgements than they would cause. The system might save 15-20 lawsuits for every one in which the computer expertise was judged to be incomplete or misleading. As such, the institutions could afford to take the risk of employing a system which is acknowledged to be incomplete and imperfect.

The problems get more complex when an allegedly harmed patient finds out that the imperfect computer suggestion that was followed by an imperfect human was written by Dr. Bonzo at Zoo University. In today's litigious society, it is natural to assume that Dr. Bonzo and Zoo University would possibly be named as defendants in any law suit. This is unfair to Dr. Bonzo because she has distributed copies of her logic to multiple institutions in the hopes of improving health care and at the same time has increased the breadth of her potential liability exposure.

Our approach going into the meeting was to get a standard agreement that one would be willing

to accept from another institution when exporting medical expertise and at the same time give to another institution when importing expertise. The main flavor of the agreement would require the recipient to indemnify the distributor for any liability (except in cases of fraud and maliciousness). In principal, the group accepted this approach but felt that in the near future every group would have to resolve these issues using specific agreements that would best be negotiated on a one on one basis. At present there is little indication of whether the author is protected or not. Until some court cases appear, it will be difficult to make blanket policies. For a more complete summary of the legal issues involved with using knowledge based systems, see Miller [19, 20] and Frank [21].

We concluded that the financial aspects are not currently an impediment to sharing and that the legal aspects are so formidable that they are best worked out on a one-on-one basis. We felt that the recipient or user of the system should generally (in the absence of any currently universally recognized standards for validating computer-based expertise) assume responsibility for the use of such systems and be willing to shield the originator of the logic modules or knowledge based programs. In the absence of this willingness to assume responsibility, there is a strong disincentive to share.

## Conclusions

We appointed sub-committees to expand the MLM representation standards, to explore the terminology problems in conjunction with the efforts by the National Library of Medicine to develop a Unified Medical Language System (UMLS), to generate lists of Decision Support Technologies, to begin to generate a standard for sharing entire programs in addition to MLM's, and to follow-up on and expand the current efforts to test and evaluate medical knowledge-based systems. These subcommittees will meet at least annually at SCAMC and may generate larger follow-up conferences to address standards for the particular subject matter. We have listed the sub-committee leaders with the intention that those interested in any one of the specific areas express their willingness to contribute to the leader so that their participation can be assured.

## Acknowledgements

The authors of this paper wish to give credit for the content of the paper to the participants of the Arden Homestead Retreat, and wish to thank the CAMDAT Foundation, IBM, and the law firm of De Forest and Duer for sponsoring the retreat.

## References

[1] Warner HR, Olmstead CM, and Rutherford BD. HELP - A program for medical decision-making. Comp Biomed Res 5:65-74 (1972).
[2] Pryor TA, Gardner RM, Clayton PD, Warner HR. The HELP System in Information Systems for Patient Care, Blum BI ed. Springer Verlag NY, pp. 109-128 (1984).
[3] Miller RA, Pople HE, and Myers JD. Internist-1, an experimental computer-based diagnostic consultant for internal medicine. N Engl J Med 307:468-476 (1982).
[4] Barnett GO, Cimino JJ, Hupp JA, Hoffer EP. DXplain: an evolving diagnostic decision-support system. J Am Med Assoc 258:67-74 (1987).
[5] McDonald CJ. Protocol-based reminders in the quality of care and the nonperfectability of man. N Engl J Med 295:1351-1356 (1976)
[6] Giuse NB, Giuse DA, Miller RA. Computer assisted multi-center creationof medical knowledge bases. Proc 12th SCAMC, Greenes RA ed., IEEE Comp Soc Press, pp. 583-509 (1988).
[7] Clayton PD, Pryor TA, Wigertz OB, Johnson SB and Hripcsak GM. Sharing medical knowledge for automated decision-making. Proc 12th SCAMC, Greenes RA ed., IEEE Comp Soc Press, pp. 591-594 (1988).
[8] Chin HL. Case-based tutoring from a medical knowledge base. Proc 12th SCAMC, Greenes RA ed., IEEE Comp Soc Press pp. 396-402 (1988).
[9] Clancey WJ. The use of Mycin's rules for tutoring. Int J Man Machine Stud 11:25-49 (1979).
[10] Clayton PD, Haug PJ, Pryor TA, Wigertz OB. Representing a medical knowledge base for multiple uses. Proc AAMSI Congress 1987 Levy AH and Williams BT eds., American Assoc for Med Sytems and Informatics Publishers (1987).
[11] Lenat D, Prakash M, Shepherd M. CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. AI Magazine 6(4):65-84 (1986).
[12] Miller PL. Medical plan-analysis by computer: critiquing the pharmacologic management of essential hypertension. Comp Biomed Res 17:25-40 (1984).
[13] Pople HE. Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnosis. Artificial Intelligence in Medicine, Szolovits P ed., Am Asso for the Advancement of Science (Westview) Washington DC, pp. 119-190 (1982).
[14] Weiss S, Kulikowski C, Safir A. Glaucoma consultation by computer. Comp Biol Med 8:25-40 (1978).
[15] Shortliffe EH, Computer Based Medical Consultations: MYCIN. American Elsevier New York (1976).
[16] Suwa M, Scott AC, Shortliffe EH. Completeness and consistency in a rule-based system. AI Mag 3(3):16-21 (1982).
[17] Anderson CL, Lundsgaarde H, Williamson J, Abaunza M, and Warner HR. A Decison support technology clearinghouse. Proc 12th SCAMC, Greenes RA ed., IEEE Comp Soc Press pp. 518-523 (1988).
[18] Penzias A. Teaching statistics to engineers. Science 244:1025 (1989).
[19] Miller RA, Schaffner KF, Meisel A. Ethical and Legal issues related to the use of computer programs in clinical medicine. Ann Int Med 102:529-536 (1985).
[20] Miller RA. Legal issues related to medical decsion support systems. Int J Clin Monitoring and Computing.
[21] Frank SJ. What AI Practitioners should know about the law. AI Mag 9(1):63-75 (1988).