

Towards computer-assisted maintenance of medical knowledge bases

Dario A. GIUSE

The Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA 15213, USA

Nunzia Bettinsoli GIUSE and Randolph A. MILLER

Section of Medical Informatics, Department of Medicine, University of Pittsburgh, B50 Lothrop Hall, 190 Lothrop Street, Pittsburgh, PA 15261, USA

Received November 1989

Abstract. The problem of knowledge acquisition is especially severe for medical knowledge bases. At least two strategies can help alleviate the problem: creating computer-based tools which assist experts in entering knowledge into the knowledge base, and promoting situations where several organizations contribute simultaneously and independently to a single knowledge base. This article describes a set of tools currently under development at the Section of Medical Informatics of the University of Pittsburgh which incorporate both strategies. The tools, which operate within the framework of the Quick Medical Reference (QMR) knowledge base, perform tasks ranging from checking syntax and semantics to suggesting initial skeletons for new knowledge base entries. Extensive use of interactive direct-manipulation techniques and built-in knowledge of many details of the knowledge base makes these tools very effective in supporting the creation and maintenance of new knowledge base entries.

Keywords: Knowledge acquisition, Knowledge engineering, Medical knowledge bases, Medical informatics, QMR.

1. Introduction

A major obstacle to the widespread use of computer-based medical expert systems is the difficulty in constructing high-quality medical knowledge bases. This problem is known as the *knowledge acquisition bottleneck* [1]. Although widely described in the literature [2, 12, 13], the problem is still far from being solved.

Various approaches to knowledge acquisition have been attempted over the years; Marcus [12] presented an in-depth review of several approaches. Creation and maintenance of early knowledge-based systems typically required a *knowledge engineer*, that is, a computer expert who was familiar with the problems and techniques of knowledge acquisition, and who could translate domain knowledge (provided by a domain expert) into a form usable by a computer. This approach presents two difficulties. First, experienced knowledge engineers are not always available. Furthermore, unless they have sufficient knowledge of the domain they may not be able to successfully communicate with the domain expert. For this reason, many of the more recent knowledge-based systems take steps which reduce, and if possible eliminate, the need for a knowledge engineer during most of the knowledge acquisition process.

The previous discussion applies equally well to all fields in which knowledge-based systems are used. In this paper, however, we concentrate exclusively on medical knowledge-based systems. The term *domain expert*, therefore, will indicate an expert clinician involved with the creation of a medical knowledge base. The term *knowledge engineer* will be used in the sense explained above.

Several strategies which facilitate knowledge acquisition have been employed in the construction of knowledge-based medical systems. An important strategy, embodied in some of the earliest systems such as MYCIN [24], is the separation of knowledge from inference engine. This separation allows modification of knowledge independently from the software which implements the inferencing aspects, and is adopted almost universally in modern systems. A second strategy is that of using knowledge about the inference engine itself to help a knowledge engineer or a domain expert refine and maintain the knowledge base; this principle was explored in TEIRESIAS [5].

A third strategy is that of using computer assistance to reduce the amount of human effort required for creation and maintenance of the knowledge base. This principle was emphasized by several knowledge-based systems, such as TEIRESIAS, ROGET [3], and SALT [12] (which is not specifically aimed at medical knowledge bases). Different possibilities have been explored. Some systems attempt to insure the consistency and completeness of knowledge as it is entered in the knowledge base. Other systems support the automatic generation of knowledge base entries from augmented transition networks [16]. Some systems support knowledge acquisition via generalization of sample entries provided by the expert; KNACK [11], for example, creates report-generator tools from sample textual reports provided by an expert. This is achieved by relying on a fixed model of the domain to restrict the possible generalizations.

Another principle is that of using the knowledge base itself to assist in refining the knowledge base. Some systems can suggest possible modifications to the knowledge base based on the system's performance on a set of problems. SEEK2 [22, 26] is a system which embodies this principle. Its knowledge is rule-based, and the system uses a library of cases to monitor its own performance and suggest changes that would improve performance. SEEK2 employs principles such as generalization and specialization to suggest plausible changes. The domain expert is responsible for reviewing the proposed changes and deciding when to incorporate them in the knowledge base.

The MUM system [8] incorporates a principle known as *design for acquisition*, which tries to facilitate knowledge acquisition by organizing the system in such a way as to reduce the representational gap between the system and the domain expert. One of key ideas in design for acquisition is to provide *task-level* primitives, i.e. to let the domain expert work with concepts that are already part of the domain rather than with combinations of system-level primitives. Yet another principle, exemplified by PROTEGE [19, 20] and KNACK [11], is that of using a single tool to generate families of related knowledge-based systems. Besides reducing the amount of effort required to create each system, this principle has the additional advantage that the tools thus created are automatically uniform and thus more amenable to standardized knowledge-engineering methods.

Many features of PROTEGE are relevant to the issues raised later in this article. One such feature is that the tools generated by PROTEGE are intended for the end user, i.e. a physician who is designing oncology protocols, rather than for the domain expert. Since the system is really a tool to build knowledge acquisition tools, the domain expert is chiefly involved at the intermediate level, i.e. using PROTEGE. Once a tool has been generated, the domain expert assumes a secondary role. Another feature that is worth mentioning is that both the system and the tools it generates

are graphical in nature. The interface is almost entirely menu- and form-based and is meant to match closely the 'natural' way of expressing knowledge in the domain, which is the design of oncology protocols. Note, however, that the tools generated by PROTEGE do not provide much support for long-term knowledge engineering. The knowledge itself is encoded into ONCOCIN [25] rules, and thus one would have to employ a different tool for long-term maintenance of the knowledge base.

2. Maintaining the QMR knowledge base

We will now concentrate on a specific knowledge base and present some results that apply to that system. The remainder of this paper describes our current efforts in developing knowledge acquisition tools for the Quick Medical Reference (QMR) system.

QMR [15, 17] is an expert system for the diagnosis of complex cases in the domain of internal medicine. Its knowledge base, based on the one developed for INTERNIST-I [18], is among the most comprehensive currently in use for general internal medicine [10]. Besides being a significant medical knowledge base in its own right, we feel that the QMR knowledge base exemplifies many of the challenges found in most knowledge-based systems.

The creation and maintenance of the QMR knowledge base is a major undertaking which has required about 25 person-years of effort so far [14]. While the number of diseases currently described in the knowledge base is about 600, we estimate that approximately 200 more will be needed to ensure coverage of the entire field of internal medicine. The task of maintaining the knowledge base as new information becomes available through the medical literature is even more problematic. It has become increasingly difficult for a single group to maintain a knowledge base this size manually. After considering different alternatives, we have identified two particular strategies [7] that can aid in expanding and maintaining the system. These two strategies are *computer assistance* and *multi-center creation and maintenance*.

Computer assistance reduces the amount of human effort required for creation and maintenance of a knowledge base by relying on computer tools to perform the more clerical aspects of the work, and by providing more efficient means to manipulate knowledge. This strategy was described in the previous section. The second strategy, i.e. multi-center knowledge base construction [7], has not been explored as widely. The most attractive feature of this strategy, which relies on geographically separate centers that independently contribute to a single knowledge base, is that it widens the base of available clinical expertise while it increases the available manpower. Although this strategy does not reduce the total amount of effort, it makes it possible to distribute it over different organizations.

It is fortunate that the two strategies can be adopted simultaneously, and in fact can be mutually reinforcing. Multi-center creation of new knowledge base units can greatly benefit from the availability of computer tools, which help clinicians who were not involved with the initial design of a knowledge base. Collaboration between multiple centers, at the same time, not only provides an incentive for the development of computer tools, but also releases a portion of the parent organization's resources which can be devoted to improving the computer tools.

It should be emphasized that not all knowledge bases present the same problems with acquisition and maintenance. Compare, for example, the problems found in the QMR knowledge base with those encountered in many rule-based expert systems. The

unit of the QMR knowledge base is the complete description of a disease and represents a sizable amount of knowledge. The unit of knowledge in rule-based expert systems, on the other hand, is the production rule, a much smaller entity. Strategies for semi-automatic improvement of a knowledge base which work well for rule-based systems may not work as well for a system such as QMR. Different systems, moreover, show different degrees of interdependency among different components of the knowledge base, and thus strategies that work for one system may not apply to others.

2.1. Disease profile: the unit of the QMR knowledge base

Before starting the description of the tools, a brief introduction to the QMR *disease profile* is needed to clarify some of the terminology. The disease profile is the basic unit of the QMR medical knowledge base [7]. A disease profile is a comprehensive description of a disease and is composed chiefly of *findings* (also called manifestations) and *links*. Findings correspond to items of history, clinical observations, or results of laboratory tests. Each disease profile contains the list of all the findings that are known to occur in that particular disease. Links express relations between the disease and other diseases in the knowledge base. Links can express different types of relation, such as a causal relation or a co-occurrence relation.

Two numbers, the *evoking strength* and the *frequency* [17], are associated with all findings and links in a disease profile. The evoking strength of finding *F* in disease *D* is analogous to the positive predictive value (the incidence of the disease is assumed to be that of the incidence in the population of patients presenting to a general internist in non-tropical climates). The evoking strength answers the question "When seeing a patient with finding *F*, how strongly should one consider diagnosis *D*?". The frequency of finding *F* in disease *D* is analogous to the sensitivity of *F* for *D*. It answers the question "Of all patients with disease *D*, how many will show finding *F*?".

A *profile submission* is the first stage in the creation of a disease profile. A profile submission is normally created by one individual after an extensive survey of the medical literature. Profile submissions are written in a format which is easy for humans to read, but is precise enough to allow a program to extract all relevant information.

Figure 1 shows excerpts from the profile submission for the disease *Cardiac Sarcoma Primary*. The line

— 0 2 AGE 16 TO 25

indicates the evoking strength (0) and frequency (2) for the QMR finding *Age 16 to 25*. The line

* 2 2 HEART ANGIOCARDIOGRAPHY VENTRICULAR FILLING DEFECT (S)

describes a newly created finding, i.e. a finding which was proposed for the disease *Cardiac Sarcoma Primary*. As the figure shows, each finding or link is supported by literature reference lines. These are short excerpts from the medical publications that were consulted during the creation of the profile submission. Each literature reference line begins with a reference number (in square brackets), followed by an indication of how relevant each reference is for a particular finding or link, and finally by the actual literature excerpt.

Disease: CARDIAC SARCOMA PRIMARY
 Number: 6.1.3.2
 Author: Coleen Andruss
 Institution: University of Pittsburgh
 Reviewer: Jack D. Myers, M.D.
 Completed: October, 1989

; Findings:

- 0 2 AGE 16 TO 25
- [5] 3 case report 23 year old
- [61] 3 case report 17 year old
 (13/59) 22% of author's reviewed case reports
- 1 2 HEART MURMUR SYSTOLIC APICAL
- [29] 3 case report systolic regurg. murmur in the tricuspid area, along left sternal border, and apex - RMS LA appendage protruding into pericardium, LA and mitral orifice
- [39] 3 case report - moderate systolic murmur at apex radiating to axilla - RMS LA and MV
- [50] 3 case report apical systolic murmur - RMS LA and MV
 (5/39) 13% author's cases reviewed
- * 2 2 HEART ANGIOCARDIOGRAPHY VENTRICULAR FILLING DEFECT <S>
- Group: Angiocardiography
- Import: 4
- Type: lab2
- [14] 3 case report angio. diagnoses RV tumor - RMS RV thru TV to RA
- [19] 3 case report Rt. heart cath - globular defects in RA, RV and PA - RMS RV, TV, PA
- [41] 1 case report angio large persistent filling defect RV outflow tract extending into Pulmonary valve orifice - RMS RV outflow tract
 (4/14) 29% author's cases reviewed

LINKS:

- ^ 1 2 CARDIAC SARCOMA PRIMARY (caus) CARDIAC TAMPONADE
- [2] 3 1/3 of 77 RMS cases reviewed showed effusion or tamponade
- [22] 3 (8/30) RMS reported hemorrhagic pericardial effusion of rapid onset - 27%
 (21/100) malignant sarcomas with hemopericardium
- [47] 3 blood tinged pericardial fluid; shock, abrupt cyanosis.
- [49] 1 case RMS - pulsus paradoxus, decreased breath sounds, dullness to percussion, rales, increased venous pressure.
- [54] 3 (3/8) primary sarcomas of the heart with tamponade - all were Angiosarcoma
 (7/49) 14% of author's reviewed case reports with evidence for tamponade

Fig. 1. Excerpts from the profile submission for the disease Cardiac Sarcoma Primary.

3. Knowledge acquisition tools for QMR

Construction of the Internist/QMR knowledge base has continued over the past fifteen years, following the guidelines originally laid out by Jack D. Myers. The QMR knowledge base is built through a careful process [6] which includes a comprehensive review of the medical literature, consultation with specialists, and an in-depth group review [15, 17]. The creation of each unit in the knowledge base requires an average of two to three weeks of full-time effort [4]. In addition to the creation of new knowledge base units, a second ongoing process is aimed at the continuous maintenance and improvement of existing units.

Over the past few years, we have developed computer-based tools to help in the acquisition and maintenance of the QMR knowledge base. All the tools run on inexpensive, widely available machines such as IBM and IBM-compatible personal com-

puters. The first phase in the development of these tools was completed in 1988, and the second phase is now underway. The tools developed during the first phase were described recently [7], and thus are only briefly presented here. The second phase of tool development will be described in more detail.

Both sets of tools are intended for use by the domain expert; the knowledge engineer never enters the process. This is largely possible because the QMR knowledge base is by now quite mature, and therefore most crucial knowledge engineering problems have already been solved. New knowledge base entries follow the general structure of the knowledge base, which is fairly stable. This enables physicians and medical students who were not involved with the design of the system to perform the major part of the knowledge base creation task.

3.1. Knowledge acquisition tools already completed

The tools developed during the first phase are specifically geared towards the initial creation of QMR disease profiles. The first tool, named CHECKLIST, is used to create an *initial checklist* for a disease profile. This is simply an initial suggestion for the structure of the profile, and is generated by consulting the existing knowledge base for the diseases that are most similar to the new one. This tool, which is used at the very beginning of the creation of a new profile, relies on the regularity of the underlying domain to yield a reasonable first skeleton for a new profile.

The second tool, named QMR-PARSE, is used to perform syntactic and semantic checks on a *profile submission*. A profile submission is the result of the knowledge acquisition process (normally performed by a single physician), just before it undergoes the group review which determines its final shape. QMR-PARSE acts in the manner of a compiler, checking each unit in a profile submission and generating appropriate error messages. In addition to a thorough checking of the syntax, this tool also enforces a few important semantic constraints which are built into the QMR knowledge base. QMR-PARSE may be used at any stage of the creation of a disease profile.

The development of these tools was guided by two principles. The first principle is that of supporting *multi-center* creation of a medical knowledge base, enabling separate organizations to contribute to knowledge base construction in a relatively independent fashion. This requires endowing the tools with many details of the QMR knowledge base, and with some knowledge about the knowledge acquisition process itself. This facilitates the unsupervised construction of knowledge base entries by physicians not associated with the parent organization.

The second principle is to rely upon the existing knowledge base in order to develop new disease profiles. Note that this is quite different from the strategy used in SEEK2 [21], for example. SEEK2, a rule-based system, uses a collection of cases to modify empirically its existing rules with the long-term goal of improving its performance. This is achieved through a heuristic strategy which analyzes existing rules and attempts to modify them to improve performance.

In the case of the QMR knowledge base, on the other hand, our goal is to extend the knowledge base significantly. This is basically a manual task which involves extensive consultation of the medical literature and expert clinicians, but the objective is to supplement human endeavors with computer assistance as far as possible. This is achieved by relying on the regularity of the knowledge base, and by using existing knowledge base units to guide the creation of new knowledge base units.

All profiles created in the past two years have benefited from the regular use of the tools. At the conclusion of the first phase, we conducted a study of the types of

errors made by different physicians preparing disease profiles, with special attention to the differences between novice and experienced subjects. The results of the study [6] suggest that degree of experience in creating disease profiles is correlated with a marked reduction in the error rate. Less experienced individuals, in particular, show many more 'mechanical' errors, i.e. errors that could be prevented by appropriate computer-based tools. These results confirmed that computer assistance can be greatly beneficial in the creation of QMR knowledge base units, and encouraged research on the second phase of tool development.

3.2. New tools: QMR-KAT

The design principles described for the original tools are being extended and incorporated into a new tool which is currently under development. This new tool, the QMR Knowledge Acquisition Tool (QMR-KAT), provides two major extensions:

- it replaces all batch-mode operations with direct, interactive operations; and
- it extends computer support beyond the profile submission stage.

The user interface style of the new tool is markedly different from the previous ones. Tools from the first phase were essentially batch-oriented. The expert would use a text editor to edit a disease profile in its various stages, and then run the profile through the tools for error checking. This had obvious disadvantages, especially for physicians who were not very familiar with computers, since the tools had to perform extensive error checking and could not provide immediate feedback as a profile was being modified. QMR-KAT eliminates this problem by adopting a direct-manipulation style of interaction [23], which allows it to edit the profile directly. Its user interface is based on windows, menus, and dialog boxes, and relies extensively on the use of a pointing device. Pull-down menus allow the user to manipulate all aspects of the profile being created, as well as to browse through the existing QMR knowledge base and the accompanying literature references.




Fig. 2. QMR-KAT, the QMR Knowledge Acquisition Tool.

Figure 2 shows QMR-KAT being used to create the profile submission for *Q FEVER*. Of the three windows shown in the figure, the Disease Profile window contains the

profile being created; compare this with the profile submission in Figure 1. This window contains the profile which is eventually submitted for the expert review. Entries that begin with a dash are findings already in the QMR knowledge base; entries that begin with an asterisk are newly created findings. Each entry is optionally followed by excerpts from the literature.

In order to add a finding which is already in the knowledge base, the user first obtains a list of findings and then copies the desired one into the profile window. This is shown in Figure 2, where the Work Area window contains a list of findings which was obtained by typing the string *occup* and asking the system for all possible matches. Any item in the work area can be copied to the profile simply by clicking over it with the mouse. A consequence of the direct-manipulation style of interaction is that syntax and semantics checking are no longer necessary, since only valid information may be entered at any stage. Only valid manifestation names, for example, may be retrieved from the knowledge base and added to the current profile. All numerical information entered by the user, in addition, conforms to the conventions already embedded in the knowledge base. This is enforced automatically by only offering the user valid choices.




Fig. 3. Enforcing knowledge base consistency rules:
Non-specific findings must be given evoking strength 0.

Figure 3 illustrates this feature of the program. Right before the state shown in the figure, the user asked to edit the entry for the finding *Exposure To Cat <s> Intimate HX*, which describes a history of cohabitation with cats. Since this finding is marked as non-specific in the QMR knowledge base, QMR-KAT only presents the value 0 (i.e. non-specific) as a possible choice for the evoking strength. The result is that a new profile being constructed is automatically in a consistent state at every point in time.

QMR-KAT has a wider scope than the previous tools. Rather than being limited to the profile submission stage, it extends well into the activity of long-term knowledge base maintenance. While the original tools did not provide support for bibliographical operations, for example, the new one does. An example of this capability is shown in Figure 4. As literature reference lines are added to the profile submission in the profile window, the user may request to see the bibliography reference for any of the relevant publications. This is done simply by clicking with the mouse over a reference number in the profile window. Figure 4 shows the result after the user clicked over the two




Fig. 4. Viewing bibliographical references in the References Area.

literature reference lines; the bibliography references for entries 1 and 7 are shown in the Reference Area window. These entries were obtained through a MEDLINE search [9]. Additional information, such as abstracts, for each publication may be obtained by selecting an item from a menu.

Manipulating bibliographical references is essential for the long-term maintenance of the knowledge base, which is based on analysis of the evolving medical literature. Support for bibliographical operations includes both an interface to existing systems for bibliographic retrieval such as Grateful Med access to MedLine MEDLINE, and the ability for the user to define and maintain local lists of bibliographical entries.

Long-term maintenance of a knowledge base requires modifications to existing units when new information becomes available in the medical literature. Moreover, additions to the knowledge base itself may require revision of some existing units. Support for long-term knowledge base maintenance in QMR-KAT takes the following forms:

- Ability to examine all details of existing disease profiles while modifying or creating new ones. This ability includes retrieving any portion of an existing profile into the Work Area and using it in the creation of a new disease profile submission.
- Ability to examine full profile submissions for disease profiles already in the knowledge base. In the past, only a portion of the data in any profile submission was actually incorporated into the final QMR knowledge base. All quotations from the medical literature, for example, were eliminated from the final profile. QMR-KAT allows all the original information to be retained and subsequently examined.
- Ability to re-utilize literature search strategies that were used during the creation of a disease profile. These are the keywords and MeSH logic that were employed during the computer-based literature review phase. Such search strategies are permanently saved as a profile submission is built, and are available for inspection afterwards. This makes it possible, for example, to re-run a computerized search during knowledge base maintenance. The results of the search show what relevant new medical information was published since the original search.

QMR-KAT, in conclusion, extends the original tools in several directions. It is an integrated tool for knowledge acquisition which supports both the initial creation of

a disease profile and long-term maintenance of the QMR knowledge base. It allows the user who is creating a profile submission to inspect the existing knowledge base and any portion of older profile submissions. Its direct-manipulation style of interface makes it easier to use for people without specific computer training. Support for bibliographical references greatly facilitates the task of the person who is creating a new profile submission.

While we have not yet conducted a formal evaluation of how QMR-KAT affects the creation of new profile submissions, the tool has already been used in the creation of a few disease profiles. These profiles were all generated by medical students who had never participated in the creation of the QMR knowledge base. Based on our previous experience, it appears that QMR-KAT makes it significantly easier to create new profile submissions. It appears, in particular, that the tool may reduce the time to create a profile submission and also reduce the number of errors in the final submission. We plan to conduct a formal evaluation which will compare the results of manual creation with those obtained by using the tools we developed during the first phase and with those obtained by using QMR-KAT.

4. Future work

Tools from the first phase automate the first level of knowledge acquisition, enforcing certain simple syntactic and semantic conventions which are used in the knowledge base. QMR-KAT will go one step further and automate the second level. We are currently in the process of evaluating which semantic aspects of profile creation are best suited for partial automation, based on our previous study [6] of actual profile submissions.

One capability we are developing, for example, is the ability to suggest appropriate frequency numbers for a finding. The information contained in literature reference lines from a profile submission is detailed enough to allow a program to compute the most likely frequency number. We intend to utilize this information and have QMR-KAT suggest those numerical quantities to the person who is creating a new disease profile. We expect that heuristics of this type may significantly improve the correctness and consistency of new profile submissions.

Much work remains to be done before the QMR knowledge acquisition tools are complete. We have already identified several directions along which QMR-KAT can be expanded. The following is a list of the major tasks, most of which we plan to address in the near future:

- The ability to send the output from QMR-KAT directly to the QMR knowledge base, once a profile submission has been approved through the expert review process. This task is scheduled for implementation in the immediate future.
- The direct integration of QMR-KAT with computerized literature search systems such as Grateful Med/MedLine. This will enable the results of computerized searches to be browsed and retrieved directly from within QMR-KAT. This task is being implemented at the time of this writing.
- The ability to manage and share personalized bibliographies along with those provided through literature search systems. Personalized bibliographies are useful both for material that appears in computer-based bibliography services and for unpublished material, such as consultation with specialists, as well as for the numerous

references used in knowledge base creation that were published prior to 1965, when MedLine began.

- The ability to constantly monitor new information being published in the medical literature. This capability, which is farther in the future, is required if the knowledge base is to keep abreast of developments in medical knowledge. The strategy we have designed is based on periodic scans of new literature as it becomes available through computerized MedLine services. Each new scan will utilize the pre-stored search strategies for a disease and will flag all new publications that match those search strategies. Such publications would then be presented to the domain expert for review and possible inclusion in the knowledge base.

Other strategies that were not described in this article should also help reduce the burden of long-term knowledge base maintenance. One such strategy would use a large number of pre-stored clinical cases to constantly monitor the performance of the knowledge base as knowledge units are added and modified. Whenever a modification alters the performance of the system on a pre-stored case, human intervention is required to determine how to handle the situation. Note that this strategy is only superficially similar to the one employed by SEEK2, because of the much larger granularity of knowledge base units in a system like QMR. This strategy, moreover, would be used not to suggest actual modifications to the knowledge base, but simply to alert the domain expert whenever a conflict arises. Dr. Jack D. Myers has already manually employed such a strategy for over a decade in maintaining the INTERNIST-I/QMR knowledge bases.

5. Conclusions

The scarcity of comprehensive, high-quality medical knowledge is a severe obstacle to the more widespread diffusion of medical expert systems. Creating and maintaining such knowledge bases is a time-consuming task which requires large investments of funds and personnel. Computer-based tools should be used to make the problem more tractable and improve the quality of the resulting knowledge bases.

We have presented some strategies for knowledge base creation and long-term maintenance. In spite of differences among different knowledge bases, many of these strategies apply to a wide range of situations and can be advantageous for different systems. Some of the best known strategies include direct-manipulation knowledge acquisition and maintenance tools, design for acquisition, and case-based automatic suggestion of incremental knowledge base modifications.

To support the continuing growth of the QMR knowledge base, we have developed a set of tools which embody several strategies for reducing the burden of knowledge base creation and maintenance. One such strategy is multi-center creation of knowledge base units, which alleviates the burden for individual institutions by spreading the work over several centers. A second strategy is that of direct manipulation of knowledge while constructing new units. A third strategy uses existing portions of a knowledge base to validate and correct new knowledge as it is entered. A final strategy uses heuristics which make quantitative suggestions based on new medical information presented to the system in a stylized form.

We have described two generations of tools developed at the Section of Medical Informatics of the University of Pittsburgh. These tools assist in various phases of the creation and maintenance of the QMR knowledge base, ranging from the initial

skeleton of new profile submissions to the long-term maintenance of the knowledge base.

While this paper concentrated mostly on a specific knowledge base and only discussed a few of the heuristics we are developing, the main strategies and problems we presented are applicable to a wide range of systems and medical knowledge bases. The next few years will see a great development in computer-assisted tools for the creation and long-term maintenance of medical knowledge bases, and future tools are likely to incorporate more than one strategy in order to reduce the knowledge acquisition bottleneck. Pending radical improvements in natural language processing by computers, which are probably still several years away, computer-assisted knowledge base creation and maintenance is the most promising answer to the problem of generating high-quality medical expert systems. Research on tools which can assist the domain expert in this activity is thus a vital part of AIM, and one which will assume even greater importance in the coming years.

Acknowledgements

The research described in this document was supported in part by the National Library of Medicine, Research Grant No. R01 LM04622. Dr. Miller's work is funded in part by Research Career Development Award K04-LM-0004 from the National Library of Medicine.

The authors gratefully acknowledge Richard A. Bankowitz, M.D., and Dr. Diana Forsythe for their careful review and helpful comments on a preliminary draft of this paper.

References

- [1] B.G. Buchanan, D. Barstow, R. Bechtal, J. Bennett, W. Clancey, C. Kulikowski, T. Mitchell and D.A. Waterman, Constructing an expert system, in: F. Hayes-Roth, D.A. Waterman and D.B. Lenat (eds.), *Building Expert Systems*, ch. 5, pp. 127-167 (Addison-Wesley Pub. Co., Reading, MA, 1983).
- [2] Association of American Medical Colleges, *Medical Education in the Information Age*, (Association of American Medical Colleges, Washington, DC, 1986).
- [3] J.S. Bennett, ROGET: A knowledge-based consultant for acquiring the conceptual structure of a diagnostic system, *Automat Reason* 1 (1985) 49-74.
- [4] N.B. Giuse, D.A. Giuse and R.A. Miller, Medical knowledge base construction as a means of introducing medical students to medical informatics, in: *Int. Symp. on Medical Informatics and Education*, Victoria, B.C., Canada (May 1989) 228-232.
- [5] R. Davis and D.B. Lenat, *Knowledge-based Systems in Artificial Intelligence* (McGraw-Hill, New York, 1982).
- [6] N.B. Giuse, R.A. Bankowitz, D.A. Giuse, R.C. Parker and R.A. Miller, Medical knowledge base acquisition: the role of the expert review process in disease profile construction, in: *Proc. 13th Annual Symp. on Computer Applications in Medical Care* (IEEE Computer Society Press, Washington, DC, 1989).
- [7] N.B. Giuse, D.A. Giuse and R.A. Miller, Computer-assisted multi-center creation of medical knowledge bases, in: *Proc. 12th Annual Symp. on Computer Applications in Medical Care* (IEEE Computer Society Press, Washington, DC, 1988).
- [8] T.R. Gruber and P.R. Cohen, Design for acquisition: principles of knowledge-system design to facilitate knowledge acquisition, *Man-Machine Studies* 26 (1987) 143-159.
- [9] P. Hewitt and T.C. Chalmers, Using MEDLINE to peruse the literature, *Controlled Clinical Trials* 6 (1985) 75-85.

- [10] E.T. Keravnou and J. Washbrook, Deep and shallow models in medical expert systems, *Artificial Intelligence in Med.* 1 (1989) 11–28.
- [11] G. Klinker, KNACK: Sample-driven knowledge acquisition, Technical Report CMU-CS-88-158, Computer Science Department, Carnegie-Mellon University (June 1988).
- [12] S. Marcus (ed.), Automating knowledge acquisition for expert systems, in: *Knowledge Representation, Learning and Expert Systems* (Kluwer, Norwell, MA, 1988).
- [13] N.J.I. Mars and P.L. Miller, Tools for knowledge acquisition and verification in medicine, in: *Proc. 10th Annual Symp. on Computer Applications in Medical Care* (IEEE Computer Society Press, Washington, DC, 1986) 36–42.
- [14] R.A. Miller and F.E. Masarie, Use of the Quick Medical Reference (QMR) program as a tool for medical education, in: *Int. Symp. on Medical Informatics and Education*, Victoria, B.C., Canada (May 1989) 247–252.
- [15] R.A. Miller, M.A. McNeil, S.M. Challinor, F.E. Masarie and J.D. Myers, The INTERNIST-1/quick medical reference project - Status report, *Western J. Med.* 145 (December 1986) 816–822.
- [16] P.L. Miller, S.R. Blumenfrucht, J.J. Rose et al., Expert system knowledge acquisition for domains of medical workup: an augmented transition network model, in: *Proc. 10th Annual Symp. on Computer Applications in Medical Care* (IEEE Computer Society Press, Washington, DC, 1986) 30–35.
- [17] R.A. Miller, F.E. Masarie and J.D. Myers, Quick Medical Reference (QMR) for diagnostic assistance, *M.D. Computing* 3 (5) (May 1986) 34–48.
- [18] R.A. Miller, H.E. Pople and J.D. Myers, INTERNIST-I, an experimental computer-based diagnostic consultant for general internal medicine, *New England J. Med.* 307 (August 1982) 468–476.
- [19] M.A. Musen, Generation of model-based knowledge acquisition tools for clinical-trial advice systems, Ph.D. Thesis, Stanford University, Dept. of Computer Science (1988).
- [20] M.A. Musen, Automated generation of model-based knowledge acquisition tools, in: *Research Notes in Artificial Intelligence* (Pitman, London, 1989).
- [21] P.G. Politakis, Using empirical analysis to refine expert system knowledge bases, Ph.D. Thesis, Rutgers University, Laboratory for Computer Science Research (1982).
- [22] P.G. Politakis, Empirical analysis for expert systems, in: *Research Notes in Artificial Intelligence* (Pitman, Boston, 1985).
- [23] B. Schneiderman, *Designing the User Interface* (Addison-Wesley, Reading, MA, 1987).
- [24] E.H. Shortliffe, MYCIN: A rule-based computer program for advising physicians regarding antimicrobial therapy selection, Computer Science Dept., Stanford University, Stanford, CA (1974).
- [25] E.H. Shortliffe, A.C. Scott, M.B. Bischoff, W. van Melle and C.D. Jacobs, ONCOCIN: An expert system for oncology protocol management, in: *Proc. 7th IJCAI*, Vancouver, B.C. (1981) 876–881.
- [26] S.M. Weiss, P.G. Politakis and A. Ginsberg, Empirical analysis and refinement of expert system knowledge bases, in: *Proc. 10th Annual Symp. on Computer Applications in Medical Care* (IEEE Computer Society Press, Washington, DC, 1986) 53–60.