

# Assessing and Enhancing the Value of the UMLS™ Knowledge Sources

Betsy L. Humphreys    Donald A. B. Lindberg, M.D.  
William T. Hole, M.D.

National Library of Medicine  
Bethesda, MD

## Abstract

*The goal of the UMLS Project is to give practitioners and researchers easy access to machine-readable information from diverse sources. Assessment of the first experimental versions of the UMLS Knowledge Sources is essential to measuring progress toward that goal and to identifying needed enhancements. As of July 30, 1991, copies of the first edition of the UMLS Knowledge Sources had been distributed to 143 individuals and institutions; 66 had provided initial feedback information. The information received indicates that the UMLS Knowledge Sources will undergo broad testing in the patient care, medical education, library service, and product development environments. Preliminary data support the hypothesis that expanded coverage of routine clinical concepts is needed. Key enhancements planned for 1992 and beyond include expanded coverage of ICD-9-CM and CPT.*

## Introduction

The goal of the Unified Medical Language System® (UMLS) project is to make it easy for practitioners and researchers to locate, retrieve, and integrate relevant information from disparate machine-readable sources. [1] Although use of single databases such as MEDLINE can be of significant help in solving practice or research problems [2], some questions can only be answered by integration of information from a variety of machine-readable sources, including the scientific literature, patient records, factual databanks, and knowledge-based expert systems. Efficient access to multiple sources requires the ability to use information found in one source, e.g., a diagnosis in a machine-readable patient record, in an automated search of another source, e.g., MEDLINE. The barriers to such integrated access continue to be formidable; they include the variety of ways the same concepts are expressed in different machine-readable sources and by users themselves and the difficulty of identifying which of many existing databases have information relevant to particular questions.

The UMLS approach to overcoming these barriers is to develop generally applicable Knowledge Sources [3] that can be used by a wide variety of applications programs: to compensate for differences in the way concepts are expressed in different machine-readable sources, to identify the information sources most relevant to a user inquiry, and to negotiate the telecommunications and search procedures necessary to retrieve information from these information sources. While it will probably be necessary for individual system developers to add to the central UMLS Knowledge Sources to support particular applications, the existence of these Knowledge Sources should reduce the local effort required to build effective systems for retrieving and integrating biomedical information. Use of the UMLS Knowledge Sources may also enhance the value and impact of existing medical informatics applications by making them more visible and accessible.

The initial development of the UMLS Knowledge Sources has been a distributed national experiment involving multiple sites and people with a broad range of expertise. Adequate assessment of the Knowledge Sources will require a similarly distributed approach. To achieve their potential for increasing the productivity of medical informaticians and for improving access to information for practitioners and researchers, the UMLS Knowledge Sources must be applied and assessed in a variety of environments and then expanded and enhanced based on the feedback from these applications. We have therefore placed a heavy emphasis both on producing the first editions of the UMLS components as rapidly as possible and on developing a distribution mechanism that promotes extensive experimentation and feedback.

## Assessing Value

Because the UMLS Knowledge Sources will be used at many different sites and for many different purposes, an overall assessment of their value will be a complex undertaking involving the aggregation of information gathered about many specific applications. All the

standard methods of evaluation applicable to medical information systems [4] could eventually be applied to assessment of the value of the UMLS Knowledge Sources. These methods include: marketplace outcome; operations research, involving rigorous measurement of predetermined process and/or outcome variables; cost-effectiveness analysis, comparing the cost of alternatives for meeting the same objective; cost-benefit analysis, determining whether the benefits to be derived warrant the expenditure of the resources required to achieve them; technology assessment, determining broad societal effects and policy implications; and studies of scientific impact, discovering the level of effect on subsequent discovery and scientific understanding [4].

Only some of these methods are feasible or appropriate during the current phase of UMLS development, however. The UMLS Knowledge Sources are currently in the product development and testing stage and are likely to be modified and expanded significantly over the next several years. In the short term, use of operations research techniques is likely to be unsuccessful since both the Knowledge Sources and any systems making use of them will be undergoing continual change. It is also too early in the UMLS development cycle for a broad-based technology assessment or for any study of scientific impact.

At this stage assessments of value will necessarily focus on market outcome, feedback to verify whether the Knowledge Sources are suitable for their intended functions and to determine where content enhancement should be made, and a number of formal and informal cost-effectiveness and cost-benefit analyses. The basic market outcome measure is the number of individuals and institutions willing to obtain copies of the UMLS Knowledge sources and to experiment with them. This can be viewed as a measure of the potential value of the Knowledge Sources and of the likelihood that NLM will receive sufficient feedback to identify and correct any important shortcomings. Feedback from those applying the UMLS Knowledge Sources will help to answer key questions about their structure and content. Our questions include: can the Metathesaurus structure represent the useful and important relationships among its source vocabularies without loss of integrity of individual vocabularies? Can the UMLS Knowledge Sources actually be used to facilitate retrieval from disparate sources of data? What specific additions in content are needed to facilitate particular applications?

A reasonable hypothesis about the UMLS Metathesaurus is that its first experimental edition does not have sufficient coverage of the clinical vernacular for

signs, symptoms, laboratory tests, and procedures to support some of the clinical applications for which it is potentially useful. This is the likely outcome of the way the Meta-1 was constructed [5]. The core set of terms for Meta-1 includes all of NLM's Medical Subject Headings (MeSH) which reflects the terminology in the published literature and is weaker in specific findings and routine laboratory tests and procedures. The core set also contains all terms from DSM-III-R [6], a set of 776 heavily used terms from three COSTAR [7] ambulatory care sites, and a set of 50 common laboratory procedures. The DSM-III-R and COSTAR sets emphasize terms for diagnoses and clinical problems. All other terms in Meta-1 (e.g., those from SNOMED, [8] ICD-9-CM, [9] and CPT [10]) are there by virtue of lexical matching to the core set of terms. Those who apply Meta-1 to certain types of clinical applications are expected to identify concepts and terms that need to be added. Another hypothesis is that the UMLS Semantic Network is insufficiently specific in some clinical areas. Specific feedback on these issues will assist in setting priorities for the gradual expansion of UMLS content.

Cost-effectiveness analysis will be required to determine whether the UMLS Knowledge Sources can reduce the amount of local resources required to achieve a given level of performance in a particular application. An example of this effect has been reported by Hersh [11] who found that Meta-1 performed just as well as a special AIDS thesaurus he had constructed for the SAPHIRE automated indexing and retrieval system. Other analyses are needed to determine whether, and under what circumstances, use of the UMLS Knowledge Sources is more cost-effective than use of other sources, and methods.

Cost-benefit analysis is probably the appropriate mechanism for assessing the relative value of specific parts of the UMLS Knowledge Sources, for choosing among alternative additions to UMLS coverage, as well as for assessing the value of new or improved applications that are inspired or made possible by the availability of the UMLS Knowledge Sources. Examples of new applications already inspired by the UMLS Knowledge Sources include Coach™, an expert assistant program designed to be used with Grateful Med® [12], NLM's microcomputer software package for searching the MEDLARS databases, and the PsychTopix [13] system for retrieving citations and abstracts relevant to particular psychiatric consultations.

An important by-product of the initial feedback from UMLS experimenters will be the identification of UMLS

applications that warrant more structured and rigorous evaluation.

### Feedback Mechanisms

In October 1990, NLM issued the first experimental editions of two of three planned UMLS Knowledge Sources: Meta-1, the first version of a Metathesaurus™ of concepts and terms from seven different biomedical vocabularies and classifications [15], and the UMLS Semantic Network [16] which represents relationships among the semantic types or categories to which all terms in Meta-1 have been assigned. The first versions of the UMLS knowledge sources are available free of charge under the terms of an experimental agreement, in three different physical formats on a set of two CD-ROMS. The experimental agreement requires each recipient to keep NLM informed of how the UMLS Knowledge Sources are being used and to assess the initial versions in light of their own professional requirements. Experimenters are prohibited from redistributing the UMLS Knowledge Sources except as an integral part of a locally developed application and must demonstrate any application to NLM before distributing it. The agreement warns that future versions of the UMLS Knowledge may differ substantially in content and format from the first experimental editions. Those requesting the UMLS Knowledge Sources are first sent sample records with documentation and copies of the Experimental Agreement. Upon receipt of the signed Experimental Agreement, NLM distributes the complete Knowledge Sources, which are available on CD-ROMs only.

NLM has developed a general strategy for obtaining feedback from those volunteering to experiment with the UMLS Knowledge Sources which will be modified as experience dictates. The transmittal letter for the CD-ROMs includes the telephone number and E-mail address of an NLM staff person to whom all technical questions are to be directed. Recipients of the UMLS Knowledge Sources are sent an initial feedback form to complete some weeks after they receive the CD-ROMs. The form requests a narrative description of the projects in which the UMLS Knowledge Sources will be used and collects more structured information about:

- physical formats and specific sections of Knowledge Sources to be used,
- hardware and software to be used,
- data involved, e.g., clinical, bibliographic, full-text,
- processes or functions to be improved, e.g., query interpretation, indexing or coding data, and
- general application environments, e.g., patient care, medical education, library service.

The form also gives experimenters an opportunity to indicate whether NLM has permission to include their names on a public list of UMLS experimenters or to circulate information about their planned experimentation and asks whether they are interested in participating in meetings to share results and provide feedback to NLM.

NLM has established an internal database of information about recipients of the UMLS Knowledge Sources into which responses to the initial feedback questionnaires are entered. Analysis of these data will help NLM to plan feedback sessions that will bring people working in similar areas together to share results, to describe how well the UMLS Knowledge Sources currently meet their needs, and to identify the specific extensions in content or changes in structure that would improve the utility of the Knowledge Sources in their applications. NLM is exploring the possibility of holding such meetings in geographically dispersed locations, such as the Regional Medical Libraries in the National Network of Libraries of Medicine (NN/LM). Other feedback mechanisms include informal "user group" sessions held in conjunction with appropriate professional meetings and one or more electronic bulletin boards.

NLM has recently issued another round of UMLS contracts [16] which will also play a key role in assessing the utility of the UMLS Knowledge Sources and contributing to their further development and enhancement. The projects conducted under these contracts will include more structured evaluation components than can be required of those who are not receiving NLM funding for their experimentation.

### Early Feedback and Assessment Data

As of July 31, 1991, some 320 people had requested UMLS sample records; 143 requestors from 29 U.S. states and 19 other countries had signed the experimental agreement and received the full UMLS Sources. This is a very gratifying show of interest, especially since the initial UMLS Knowledge Sources are intended primarily for system developers and are not accompanied by application programs. Based on the level of demand for the first experimental edition, the potential value of the UMLS Knowledge Sources appears to be very high.

An analysis of the 66 responses to the request for initial feedback received as of July 30, 1991 indicates that most recipients of the UMLS Knowledge Sources are currently planning their experiments and have not yet

begun extensive use of the Knowledge Sources. This is not surprising, given the size and complexity of the UMLS components. The tables that follow summarize data from the 66 responses. Multiple answers were permitted to all questions so the percentages shown add total more than 100%.

Most experimenters expect to apply the UMLS Knowledge Sources in more than one environment:

- 61% Patient Care
- 48% Medical education
- 50% Library Service
- 26% Biomedical Research
- 45% New/improved products

Many experimenters also expect to work with multiple types of data:

- 52% Clinical data
- 55% Bibliographic data
- 33% Full-text
- 9% Research data
- 15% Curriculum data

and to apply the UMLS Knowledge Sources to the improvement of several processes or functions:

- 50% User query interpretation
- 55% Search Strategy formulation
- 33% Selection of information sources
- 63% Indexing or coding data
- 38% Natural language processing
- 24% Computer assisted instruction
- 74% Information retrieval
- 11% Other functions

The responses to the initial feedback questionnaires seem to indicate that many UMLS experimenters intend to apply the UMLS Knowledge Sources to the problem for which they were designed, i.e., linking and retrieving different types of data. The high percentage interested in indexing and coding data implies an understanding of the substantial benefits to be gained by some level of vocabulary control. When the respondents are segmented by combination of the environment and type of data selected, the 30 who selected the patient care environment and clinical data show a greater than average interest in indexing and coding data and computer-aided instruction. The 26 who selected both library service and bibliographic data show a greater than average interest in user query interpretation, search strategy formulation, selection of information sources, natural language processing, and information retrieval. The 8 who selected both medical education and curriculum data exceeded the average in selecting user query interpretation, indexing and coding, and information retrieval.

In many cases, experimenters expect to use multiple

hardware platforms:

- 47% PC-Compatible /DOS
- 23% SUN/UNIX
- 59% Apple Macintosh
- 21% Other

Since the browsers provided operate on Macintosh computers, many developers are using Macintosh hardware to become more familiar with the content of the UMLS Knowledge Sources and are also extracting data from one or both of the ASCII versions for use on other hardware platforms. For similar reasons, about one third of respondents indicate that they will use Hypercard, usually in addition to other software. A wide variety of software was listed including many locally developed programs. The overwhelming majority of respondents are willing to have their names and addresses included on a public list of UMLS recipients and are interested in participating in feedback sessions. About 67% are willing to have NLM distribute the information they provided on the feedback questionnaire.

Although most UMLS experimenters are still in the planning phase, some initial assessments of the structure and content of the first editions of the UMLS Knowledge Sources have been received. These comments support the hypotheses that additional clinical terminology is needed in the Metathesaurus and that it may be useful to extend the level of specificity in some sections of the UMLS Semantic Network.

#### Current Enhancement Plans

Given the lead time required to set up meaningful experiments using the UMLS Knowledge Sources, NLM is relying primarily on comments and data from NLM staff and from the groups performing UMLS research and development under contract to identify specific UMLS enhancements to be made in 1991 and early 1992. We expect to begin receiving more substantive feedback from many UMLS experimenters by the spring of 1992 and will be able to factor this input into changes made to the UMLS Knowledge Sources for 1993.

For 1991, a major UMLS development goal is to design and begin development of a robust operational system for ongoing maintenance and expansion of the Metathesaurus. To allow concerted effort toward this objective, the content changes reflected in Meta-1.1 have been limited to the correction of errors in content or format discovered by UMLS experimenters and the incorporation of the 1991 MeSH additions and changes. There are no changes to the Semantic Network for 1991. The major addition to the 1991 edition of the UMLS

Knowledge Sources is the inclusion of the first version of the Information Sources Map.

Current plans for Meta-1.2, to be released in 1992, call for additional high priority clinical terms from CPT and ICD-9-CM, expanded coverage of toxic chemicals that present health hazards when released into the environment, and addition of the existing French and/or German translations of MeSH developed by NLM's International MEDLARS partners. This last enhancement is a first step toward improving the international utility of the Metathesaurus. In 1989 federal legislation gave NLM expanded responsibility for improving information access in the field of health services research, with special emphasis on technology assessment. Accordingly, we expect to enhance the UMLS Metathesaurus by adding the Universal Medical Device Nomenclature that is maintained by ECRI, a non-profit organization that evaluates health devices, tracks adverse effects of such devices, and produces a variety of information products related to health devices and clinical guidelines and standards. This addition will occur in 1992 or 1993. Health services research interests will also affect the selection of the additional ICD-9-CM and CPT terms and codes added to the Metathesaurus in 1992 and beyond.

#### References and Footnotes

1. Humphreys BL, Lindberg DAB. Building the Unified Medical Language System. In: Kingsland LC III, ed. Proceedings of the 13th annual SCAMC Washington, D.C.:IEEE Computer Society Press, 1989;475-80.
2. Use of the Critical Incident Technique to Evaluate MEDLINE. Springfield, VA: National Technical Information Services, 1989. (PB90-142522/GBB).
3. Lindberg DAB, Humphreys BL. The UMLS Knowledge Sources: Tools for Building Better User Interfaces. In: Miller RA, ed. Proceedings of the 14th annual SCAMC, Washington, D.C.: IEEE Computer Society Press, 1990;121-125.
4. Lindberg DAB, The Growth of Medical Information Systems in the United States. Lexington, MA:Lexington Books, 1979; 69-78.
5. Sherertz DD, Olson NE, Tuttle MS, Erlbaum MS. Source Inversion and Matching in the UMLS Metathesaurus. In: Miller RA, ed. Proceedings of the 14th annual SCAMC. Washington, D.C.: IEEE Computer Society Press, 1990, 141-145.
6. Diagnostic and Statistical Manual of Mental Disorders. DSM-III. 3rd. ed. Washington, D.C.: American Psychiatric Association, 1980.
7. Barnett GO. The application of computer-based medical record systems in ambulatory practice. *N Engl J Med* 1984 June 21;310(25):1645-49.
8. Cote RA. ed. Systematized Nomenclature of Medicine. Skokie, IL: Collection of American Pathologists, 1979.
9. CPT 1989: physician's current procedural terminology, fourth edition. Chicago: American Medical Association, 1988.
10. The international classification of diseases: 9th revision, clinical modification: ICD-9-CM. 2nd ed. Washington, D.C.: U.S. Health Care Financing Administration, 1980. [For sale by Supt. Docs., U.S. Government Printing Office.]
11. Hersh W. Evaluation of Meta-1 for a Concept-Based approach to the Automated Indexing and Retrieval of Bibliographic and Full-Text Databases. Abstract. Second Annual Educational and Research Conference of AMIA.
12. Haynes RB, et al. Online access to MEDLINE in clinical settings. A study of use and usefulness. *Ann Intern Med* 1990 Jan 1;112(1):78-79. San Francisco, 1991.
13. Powser SM, Miller PL. Linking Bibliographic retrieval to Clinical Reports: PsychTopix. In: Kingsland LC III, ed. Proceedings of the 13th annual SCAMC. Washington, D.C., IEEE Computer Society Press, 1989; 431-435.
14. The general UMLS research and development contractors for July 1991 - July 1993 are: Brigham and Women's Hospital, Columbia University, Massachusetts General Hospital, University of Pittsburgh, with subcontractor, University of Utah, and Yale School of Medicine, Lexical Technology, Inc. is the principal contractor for the development of the Metathesaurus Maintenance and Enhancement System.
15. Tuttle MS, et al. Using Meta-1--the First Version of the UMLS Metathesaurus. In: Miller RA, ed. Proceedings of the 14th annual SCAMC. Washington, D.C.:IEEE Computer Society Press, 1990, 131-135.
16. McCray AT, Hole WT. The Scope and Structure of the First Version of the UMLS Semantic Network. In: Miller RA, ed. Proceedings of the 14th annual SCAMC. Washington, D.C.: IEEE Computer Society Press, 126-130.