*Expert Systems and Information Retrieval

Gerard Salton *

(Reprinted from SIGIR Forum, Volume 21, Numbers 3-4, 1987)

Abstract:

The existing bibliographic retrieval systems are too complex to permit direct on-line access by untrained end users. Expert system approaches have been introduced in the hope of simplifying the document indexing, search and retrieval operations and rendering these operations accessible to end users. The expert system approach is examined briefly in this note and the conclusion is reached that expert systems are unlikely to provide much relief in ordinary retrieval environments. Simpler and more effective retrieval systems than those currently in use can be implemented by falling back on methodologies proposed and evaluated over twenty years ago that operate without expert system intervention.

1. The Need for the Expert System Approach

In the early years of computing, the user was effectively removed from the computational process and most automatic processes were carried out in a batch processing mode. The computer environment was revolutionized in the early 1970's by the introduction of hands-on, on-line computing where users could control the operations by using interactive computer terminals. It was expected that the on-line computer revolution would also affect the end-user in automatic search and retrieval situations by allowing the user population directly to initiate and control the retrieval process. Unfortunately, this has not proven to be the case for two main reasons:

a) many competing information retrieval systems have been designed with different access methodologies and different command languages; in these circumstances, the users are effectively required to learn a variety of different access languages and methodologies;

b) the standard Boolean search methods that operate with controlled indexing and search languages are difficult to master and remain inaccessible to normal end-users.

The current situation relating to the use of retrieval systems has been described by Cleverdon in the following terms: [1]

> "So there is now the situation where the desire of database producers and hosts to see a major increase in the use of on-line (searching) is doomed to comparative failure because of the user-hostile systems which they operate, and they are forced to court the favors not of the end-users but of the professional intermediaries who are, in general, the only people capable of using present-day on-line bibliographic databases in science and technology."

The solution currently envisaged in library and information retrieval consists in introducing the expert system approach: [2]

> "I believe that expert systems research is the new frontier and the next area of development in library information science. Expert systems will enable users to make more effective use of the automated systems and on-line databases that were designed and implemented during the past decade, and they will help the libraries be more productive and efficient in carrying out the many tasks involved in managing an information service center."

---

* Department of Computer Science, Cornell University, Ithaca, NY 14853.

This study was supported in part by the National Science Foundation under grants IST 83-16166 and IST 85-44189.

The expert system approach consists in asking human experts--for example, human search intermediaries--to supply the facts and the rules used by these experts in solving certain specific information retrieval problems. Improved systems could then be designed based on the knowledge derived from the expert human informers.

The applicability of the expert system approach to problems in information system design is briefly examined in the next section.

2. The Use of Expert Systems in Information Retrieval

The expert system approach may be based on two main components: the use of a stored *knowledge base* which captures in some organized way the facts of interest in a given subject area, and a *set of rules* used to derive problem solutions by judicious applications of the available knowledge. To derive problem solutions, interactive methods are normally used to traverse the knowledge base in accordance with individual user requirements by letting the user make appropriate choices of operations at particular times.

The introduction of expert system techniques in information retrieval is designed to produce flexible information handling operations that would effectively remove some of the limitations of the existing information handling systems. Among the contemplated uses of expert systems in information handling, the following applications may be of principle importance:

a) user studies designed to elicit the information needs and characteristics of each user with the aim of tailoring the search effort to individual user requirements;

b) studies of knowledge representation systems and of methods designed to extract the required subject content from the text of available document collections;

c) generation of refined systems for cataloguing, indexing , or classifying information items based on the organization and structure of the available knowledge base;

d) construction of automatic search intermediary systems designed to transform the user's statement of information need into a viable search formulation, and the subsequent implementation of interactive database searches;

e) generation of additional text transformation systems such as for example, automatic text summarization and text abstracting systems.

The implementation of some of these expert-type systems would be simplified if accurate information were available about the subject areas under consideration and about the manipulations needed in each case. The following sources of information might be used in particular: knowledge of the particular subject area and of the common sense world knowledge applicable to all subject areas, knowledge of standard library and referral sources, knowledge of user characteristics, and of course knowledge of the problem areas to be covered, such as, for example, the storage and retrieval techniques that might be used to carry out automatic searches.

In order to determine whether expert system techniques are applicable to the information retrieval environment, it may be useful to outline briefly the structure of existing expert system designs. The best known of the operational expert systems, and the only one in widespread use that appears to be commercially successful is the R1, or XCON, system designed to generate computer configurations obeying certain problem specifications by assembling individual components with stated capabilities. In the R1 system, the basic computer components and their properties represent the elements of the stored knowledge base, and rules used to assemble components into complete configurations represent the system expertise. [3] Because several hundred components exist and many dozens of rules of assembly, the human system

designer finds it hard to envisage all possible designs that satisfy a given problem statement, and for this reason the automated system substantially simplifies the design process.

In evaluating the operational expert systems such as R1, it must be noted that the definition and properties of each individual computer component is unequivocal, and that the rule set specifies precisely the necessary and sufficient conditions by which complete configurations are constructed from individual components. The computer configuration application thus represents a structured problem area from which major ambiguities and uncertainties are largely absent. The question arises whether this type of environment is reasonably descriptive of the information search and retrieval situation. The following analogies have been drawn in this respect: the choice of particular search words is assumed to be similar to the choice of a computer component; a complete search formulation is then similar to a complete assembly of components, and the constraints that operate on the assembly of components into complete configurations are identifiable with the syntactic and semantic rules that control the construction of well-formed query statements: [4,p. 120]

> "the search formulation to be processed on a document retrieval system is analogous to the set of diagrams produced by R1 used by the technician in assembling the system components. The constraints regarding the syntax and semantics of the query formulation are analogous to the constraints regarding the configuration of components."

A search system using the expert-system approach could then take the form of a structured search space that describes the subject area under consideration, together with rules to traverse the search space based on information provided at each selection point by particular system users. For example, to search a database in the area of cancer research, one might start with the initial notion of "cancer therapy", from where one might proceed to various therapy types such as "chemotherapy" or "radiotherapy" or "surgery", which in turn would furnish access to types of drugs used in chemotherapy, or various types of surgical procedures. Rule-based programming techniques would be used to control the user interface and to make choices of appropriate search terms for incorporation into the users' search formulations. [4]

Unfortunately, as Pollitt remarks himself, the analogy between the assembly of computer components into complete configurations, and the assembly of individual search terms into finished search statements, is not felicitous:

a) unlike the individual computer components, the search terms are often ill-defined and ambiguous, and their coverage is uncertain;

b) the rules used to construct complete search statements are not transparent and cannot be easily constructed; in particular, the normal "if A then B" type of rule ("if term A is used then term B should be deleted") is certainly not adequate in expressing the multiplicity of constraints that governs the search formulation process;

c) the traversal of a hierarchically organized subject schedule may have very little to do with a search formulation process actually used by experienced searchers, and the query statements produced by a rule-base process may not represent formulations that actually reflects the user's information needs.

In other words, in information retrieval the relationships between concepts may be elusive, and the connections between topics sufficiently idiosynchratic to preclude the construction of simple rules of assembly as required in rule-based processing. Furthermore, the hierarchical representation of knowledge in a particular topic area, and the need to make choices among sets of well-defined search paths may be confining, and may prevent the generation of useful query formulations. This is confirmed by a study of existing classification schedules and cataloguing rules, where it was found that many terms used in cataloguing are ill-defined, many defined concepts overlap, many distinctions and exceptions to the normal cataloguing rules are introduced on an ad-hoc basis, and in general the rules are not properly formulated and thought out. [5]

There is much evidence that in areas where document content and user information needs play a major role, the structured systems approach which underlies the implementation of useful expert systems is not applicable. Instead we are required to deal with ambiguous concepts, and rules that may apply to some degree and in particular circumstances, and whose consequents may not follow unequivocally from the antecedents. Of course, expert system designs have been proposed that can deal to some extent with reasoning under uncertainty. But no one has actually attempted to implement such systems in complicated situations involving natural language understanding, and there is no reason to believe that rules can be elicited for processes that remain elusive even for the expert. [6]

Instead of designing expert-based document processing systems that may be arbitrarily cobbled together, it may be much preferably to return to basics. Over twenty five years ago, fully automatic information retrieval systems were introduced that were based on automatic indexing products derived from natural language user statements. Both documents and user queries were represented by sets of properly weighted index terms. Global matching operations were implemented to compare terms attached to documents and queries, producing ranked document output in decreasing order of the computed query-document similarities. This ranked output could then be evaluated and used as input in a relevance feedback procedure to generate more refined user queries capable of producing improved search output. [7-11] These procedures are the main ingredients of the experimental Smart retrieval system. They are far simpler, and much more effective then any of the proposed expert-system approaches, and they may be expected to remain so for the foreseeable future.

References

[1] C.W. Cleverdon, Historical Note, Journal of the American Society for Information Science, 38:3, 1987, 152-155.

[2] H. Borko, Getting Started in Library Expert System Research, Information Processing and Management, 23:2, 1987, 81-88.

[3] J. McDermott, R1: A Rule-Based Configurer of Computer Systems, Artificial Intelligence, 19, 1982, 39-88.

[4] S. Pollitt, CANSEARCH: An Expert System Approach to Document Retrieval, Information Processing and Management, 23:2, 1987, 119-138.

[5] R. Hjerppe and B. Olander quoted in R. Davies, Outlines of the Emerging Paradigm in Cataloguing, Information Processing and Management, 23:2, 1987, 89-98.

[6] H.P. Dreyfus and S.E. Dreyfus, Mind over Machines, The Free Press, NY, 1986.

[7] G. Salton and M.J. McGill, Introduction to Modern Information Retrieval, McGraw Hill Book Company, NY, 1983.

[8] G. Salton, A Blueprint for Automatic Indexing, ACM SIGIR Forum, 16:2, Fall 1981, 22-38.

[9] G. Salton, Another Look at Automatic Text Retrieval Systems, Communications of the ACM, 29:7, July 1986, 648-656.

[10] G. Salton, E.A. Fox and E. Voorhees, Advanced Feedback Methods in Information Retrieval Journal of the ASIS, 36:3, May-June 1985, 200-210.

[11] G. Salton, C. Buckley, and E.A. Fox, Automatic Query Formulations in Information Retrieval, Journal of the ASIS, 34:4, July 1983, 262-280.