EXPERT SYSTEMS IN LIBRARIES

edited by

Rao Aluri

Burr-Brown Corporation Library

and

Donald E. Riggs

Arizona State University

ABLEX PUBLISHING CORPORATION
NORWOOD, NEW JERSEY

Printed in the United States of America

Library of Congress Cataloging-in-Publication Data

Expert systems in libraries / edited by Rao Aluri and Donald E. Riggs.
p. cm.
Includes bibliographical references.
ISBN 0-89391-589-0
1. Expert systems (Computer science)—Library applications.
2. Libraries—Automation. I. Aluri, Rao. II. Riggs, Donald E.
Z678.93.E93E94 1990
025'.00285416—dc20 90-263
CIP

Ablex Publishing Corporation
355 Chestnut Street
Norwood, New Jersey 07648

CONTENTS

PREFACE

A new phenomenon known as "expert systems" is rapidly becoming part of the librarian's vocabulary. Expert systems offer a level of technology sophistication that has never been experienced in the library world. For the very first time, we now have machines that can "talk" intelligently with humans. Expert systems run on personal computers and are commonly carried on skeletal systems known as "shells." A shell is a software package designed for the creation of specific expert systems.

An expert system is a subset of artificial intelligence which is a subfield of computer science concerned with designing systems that perform human-like, intelligent functions. In order to meet the requirements of a genuine expert system, the intelligent capabilities of humans must be rivaled; in essence, the human's thinking/reasoning processes are cloned. For example, the expert system must: ask intelligent questions, solve problems, explain reasoning, and justify conclusions. An expert system is basically an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solutions.

Expert systems are being used in a variety of fields. More than $25 million per year is being spent by U.S. industries on expert systems research. The application of expert systems in libraries has only occurred within the past few years. One of the primary drawbacks to using expert systems in libraries is that of defining a workable domain. Using an expert system to teach and monitor the implementation of the second edition of the Anglo-American Cataloging Rules is doable. However, believing an expert system can perform all duties of a reference librarian is foolish. The interaction between the reference librarian and the user is a complex endeavor encompassing many unanticipated challenges and other human anomalies. Nevertheless, an expert system can improve the effectiveness of a reference librarian in many ways. Other library areas that can reap benefits from expert systems include: bibliographic instruction, indexing, and management. We will witness other uses of expert systems in libraries in the near future.

The Selection of Search Keys

Raya Fidel
Graduate School of Library and Information Science
University of Washington
Seattle, WA

INTRODUCTION

One of the tasks carried out by reference librarians when performing an online search of bibliographic databases is the selection of search keys. To understand the nature of this task it is best to examine its place in the process of online searching.

Online searchers answer requests from library patrons: The patron presents the searcher with a request that reflects the patron's information need. Once a searcher feels she understands the request well enough to answer it, she performs an online search in the relevant databases. The outcome of an online search is typically an answer set that includes bibliographic citations.

Patrons' requests have two major aspects. The first is the topic of a request which presents the subject matter of concern to the patron. For example, "the analysis of students' behavior during final examinations to determine the difficulty of the examination" is a topic of a hypothetical request.

The second aspect includes request characteristics that do not relate directly to the topic but rather to the purpose of the request, or to how the information is going to be used. For example, a patron may need a comprehensive search that retrieves all the relevant citations (high recall search), or she may be interested in just a few highly relevant citations (high precision search). In addition, at one time the patron may agree to consider articles about any kind of examinations—not only finals—and at another time she may be satisfied to receive citations to articles dealing with the analysis of students' behavior during final examinations, whether or not the analysis is used to determine the difficulty of the examination (low specificity search).

A classic online search includes the following procedure. The searcher interviews a patron to clarify the topic and the characteristics of a request. The searcher then develops a plan for searching the request online—a search strategy. The strategy specifies which databases will be searched and which terms (or search keys) will be used in each database. It can also include a more specific plan that determines the flow of the search: the terms that will be entered first, when to review some results, and what to do if the results are not satisfactory. The actual session at the terminal follows the search strategy but searchers may deviate from their original plan if it does not seem useful. Some requests may require a number of terminal sessions because searchers may decide to log off and reconsider their strategy, possibly with the help of the patron. At some point the searcher will decide to terminate the search and print the answer set that is given to the patron.

This chapter focuses on one segment of the search process: the selection of search keys.

To begin, what decisions are involved in the selection of search keys?

Searchers must first break down a request into its individual components or concepts. The request about students' behavior, for instance, includes four concepts: (l) analysis, (2) students' behavior, (3) final examinations, and (4) difficulty of the examinations. Each concept requires a set of search keys for its representation. Thus, searchers look for search keys that will best capture the literature on the topic of each individual concept or of the request as a whole, and, at the same time, will retrieve an answer set that satisfies other request characteristics such as recall, precision, or specificity.

What are the options in the selection of search keys?

There are two distinct types of search keys. Searchers may enter any desired term or phrase and retrieve citations that include the term or the phrase in their text (free-text searching). Or, searchers may decide to use search keys from a thesaurus—a list of descriptors, or subject headings that are used for indexing and retrieval (descriptor searching). For many databases both options are available, and searchers can also use them in combination: A concept may be searched using free-text terms as well as descriptors.

Thus, one of the decisions that searchers make during an online search is what type of search keys to use. They decide this when they plan their strategy and also during terminal sessions when they revise their strategy. The significance of this decision is clearly demonstrated in a recent review article (Svenonius, 1986) which analyzes the large body of literature on the subject.

Expert systems that perform like online searchers may be called intermediary expert systems (IES) because they mediate between the patron and the databases he searches. Ideally, patrons would be able to present their requests to IES which then would make expert decisions about the search process, and in particular about the selection of search keys. IES should interrogate patrons to elicit request characteristics, but the systems will use their expertise to make decisions about matters that are beyond the knowledge of patrons. For example, such an expert system should ask the patron whether high recall or high precision is required, and will use this information to decide whether to use free-text terms or descriptors.

Various intermediary systems are already available for public access, such as

CITE (Doszkocs, 1983); others, such as CONIT (Marcus, 1983) or CAN-SEARCH (Pollitt, 1984) are still being tested in experimental settings. Through such systems, patrons are freed from encounters with many peculiarities in databases and search systems—such as ORBIT, DIALOG, or BRS—and yet can benefit from a large range of capabilities. In particular, patrons can enter a request in a loosely structured format, preferably in natural language, using a sentence-like expression. An intermediary system processes the request terms, displays information to patrons, and asks for feedback. The information displayed may be in the form of a list of subject areas, databases, search keys, or actual citations from which patrons are asked to make a selection, possibly in ranked order. Interactions of this nature usually proceed until patrons terminate the session.

Some intermediary systems are actually helper systems: they provide menu-driven interaction and thus free patrons from learning the command language while still requiring them to make most of the decisions during a search process; or, they drastically simplify searching by reducing the number of options to a minimum. CITE, for example, leaves the selection of search keys to the patron: It displays a list of search keys that can be used for a request component—both descriptors and free-text terms—and asks the patron to select the terms. In contrast, CONIT simplifies the selection of search keys because it searches each search key as a free-text term, and under certain circumstances (that depend on the search system, not on the request) also as a descriptor.

Intermediary expert systems, on the other hand, attempt a more powerful form of patron assistance: They replicate the performance of an expert in online bibliographic retrieval by incorporating the knowledge of an expert with rules for making inferences on the basis of this knowledge.

Most current attempts to build IES for interface with bibliographic databases are based on text analysis. This means that decisions these systems make are based on analyzing the text of the citations that are stored in them. Such a system may have, for example, a rule that states: if the number of citations retrieved for one request component is less than 20, enter a descriptor that is most strongly associated with the original descriptor. Thus, suppose the descriptor *Students' Behavior* retrieves 18 citations. The system then examines all the descriptors that have been assigned to these citations and discovers that 11 citations include also the descriptor *Elementary Education*. Following its rule, it will next retrieve citations to all documents about elementary education, regardless of whether or not the patron is interested in this area, or whether she needs high precision or high recall.

The primary limitation of building a knowledge base on text analysis alone, then, is that it can deal only with the topic but not with the other characteristics of a request: Aspects such as recall, precision, or specificity requirements cannot be addressed solely by text analysis. Yet, such aspects are extremely important for successful interaction and are thus central to search processes when performed by online searchers. It is desirable, therefore, to develop a form of IES that would incorporate a simulation of the knowledge that online searchers gained through experience.

To build such systems, one must first collect and present in a formal manner the searching behavior of online searchers. This chapter is a report of a research project that began to investigate the rules that online searchers apply when they select search keys.

THE METHOD

To examine the rules for the selection of search keys about 50 online searchers were observed performing their regular, job-related searches. Searchers who have been searching for more than two years were recruited from among librarians in all types of library, searching a wide variety of subjects. They were studied one at a time, and were asked to verbalize their thought processes during their searching to the degree that speaking out loud would not interfere with their performance. These verbalizations, including the creation of search strategy, were recorded. At the end of the observation period (between 5 to 10 searches), each searcher was interviewed to reveal and clarify information not accessible to observation.

Data collected for analysis were about 300 printed search protocols with transcriptions and verbalized thought processes and additional explanation from the final interviews.

Each instance in which searchers had selected a search key was then identified and the reason for the specific choice was explicitly noted. Analysis of the first 10 search protocols generated a preliminary list of conditions under which a particular selection was made. For example, a condition for choosing a free-text term is: to enter straightforwardly a specific concept which might not be a trustworthy descriptor.

All the search protocols were then analyzed against this preliminary list of conditions. Each instance in which searchers had selected a search key was listed under the condition to which it applied. Instances whose condition could not be found suggested a new condition to be added to the list. This analysis revealed that some conditions were considered by most searchers, others occur quite rarely, and a few conditions reflect searchers' individual idiosyncrasies.

The list of conditions for the selection of search keys is presented in Figure 12.1 in a form of a decision tree, and is explained in the next section. This set of decision rules is called here "The Selection Routine."

The Selection Routine describes the most commonly selected path, but there might be complications. As such, it cannot always accurately predict the selection of search keys unless other factors and their impact are known. This Routine groups together similar conditions so it could be presented in a decision tree, but it is not meant to represent a necessary sequence in the selection process.

THE SELECTION ROUTINE

Two questions are central to the decision about how to represent a request term:

1. Is it a "good" term for free-text searching?
2. Can the term be mapped to a descriptor?

To map a term to a descriptor is to select a certain descriptor as the representative of the term, whether or not the descriptor matches exactly the term.

The first question is a little more complex and requires some explanation. Searchers consider a term to be a "good" term for free-text searching if it (a) usually occurs in a particular context, (b) is uniquely defined, and (c) is specific in the concept it represents. Such a term will be called here a "single-meaning" term. On the other hand, a term that occurs in *more* than one context will be called a "common" term. For example, in the request about the analysis of students' behavior during final examinations to determine the difficulty of the examination, terms such as *students* and *behavior* are single-meaning terms. The term *examination*, on the other hand, is a common term; it can occur in a subject-related context ("the best way to take student examination"), or in a descriptive capacity ("examination of students' responses"), or still further, it can be used very loosely to represent the concept of an inquiry of any kind. Similarly, the term *analysis* is a common term because it is likely to occur in many abstracts, whether or not the articles report about actual analyses.

When a term is a common term, searchers most often do not have much choice in the selection of search keys: If it can be mapped to a descriptor, searchers almost always enter the descriptor as a search key [A] (i.e., option [A] in Figure 12.1) because, by definition, it is not desirable to use a common term as a free-text search key.

Only rarely, and under special circumstances, would searchers elect to enter a common term as a free-text search key if it can be mapped to a descriptor [B]. In the 300 searches analyzed for this study, such a selection was made only twice. In both instances, searchers were using a term to limit the retrieval and assumed that to require that a descriptor is assigned would be too limiting. A searcher may, for example, combine all the concepts of the request about students' behavior but the "analysis" concept. Looking at the retrieval she concludes that precision is not high enough. ANDing the set with the descriptor *Analysis* is likely to eliminate citations that are possibly relevant. ANDing it with the free-text term *analysis*, on the other hand, is likely to eliminate a smaller number of citations because the term is common in abstracts. Thus, precision is improved and recall is not reduced drastically.

The common term that cannot be mapped to a descriptor almost always results in unsatisfactory retrieval. Searchers have no choice but to enter a free-text key [C]. Under such conditions, they often combine the term with other search keys that represent additional request components, to retrieve citations, select some

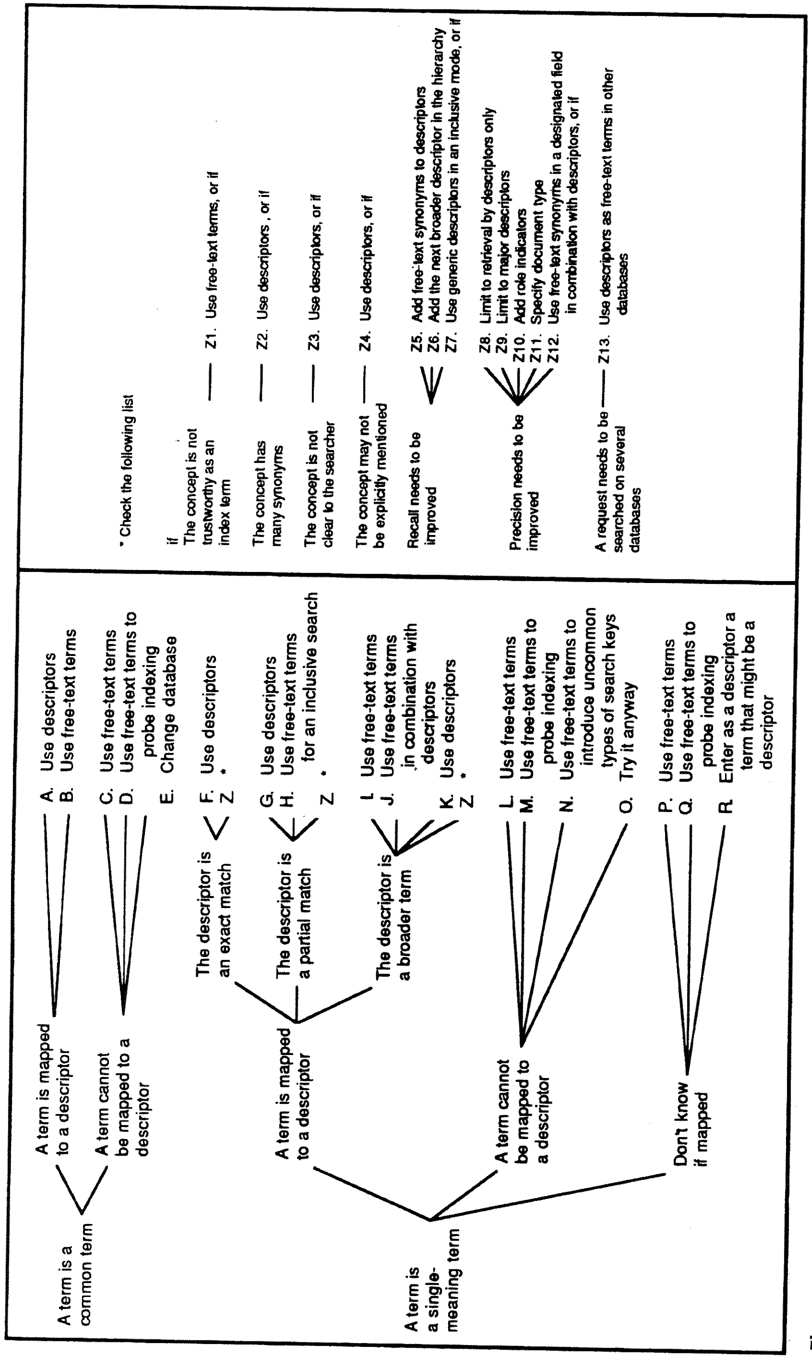

Figure 12.1. The Selection Routine.

relevant ones, and review their indexing in an attempt to find descriptors that might possibly be relevant [D]. For example, if the term *examination* cannot be mapped to a descriptor, one can AND the descriptors *Students' Behavior* and *Academic Difficulties* with the free-text term *examination*. Reviewing a sample of retrieved citations, one may find that all the relevant citations include the descriptor *Instructional Tests* in their indexing, thus suggesting that this descriptor is an appropriate choice for the representation of the concept *examination*.

Such probing does not always further the search and searchers may then decide to select a different database: one which does allow the common term to be mapped into a descriptor [E].

Single-meaning terms provide more options. If a single-meaning term cannot be mapped to a descriptor, searchers may enter a free-text term as the only search key [L], but they may also probe indexing of relevant citations to make sure that no adequate descriptor is overlooked [M].

In some cases, searchers may use a free-text key to search for a single-meaning term that cannot be mapped to a descriptor in a special way: They require that it occurs in a field other than the common ones, such as the journal title field [N]. Suppose a patron is interested primarily in the behavioral aspects of the request, and suppose the term *behavior*, or *students' behavior* cannot be mapped to a descriptor. Searchers may predict that searching for the occurrence of *behavior* in the text would retrieve a large number of irrelevant citations, and decide instead to retrieve citations to articles whose authors are affiliated with organizations which include the stem *behav* in their titles, or articles that were published in sources whose titles include this stem. If all else fails, searchers may enter a single-meaning term that is not mapped to a descriptor as a descriptor just to ascertain that their failure to map the term to a descriptor was not a mistake [O].

Another set of similar possibilities arises when searchers do not know whether a term is mapped to a descriptor. This may happen when the pertinent thesaurus is not available to searchers, or when they decide not to consult one. Here again, they can use free-text terms [P], they can use free-text terms to probe indexing [Q], or they can enter the term as a descriptor assuming that the term might be one [R].

The least problematic term is one that is single-meaning and also can be mapped to a descriptor. Such a term can be entered as a descriptor, as a free-text key, or as both. Here, searchers are free to deal with other request characteristics when selecting search keys. These characteristics are presented in the Z1-Z14 sequence. A single-meaning term that is mapped to a descriptor through any kind of match always includes a Z option. This option covers a variety of conditions, such as: The concept has many synonyms, recall needs to be improved, or a request needs to be searched on several databases. The Z options will not be discussed here; a somewhat dated but detailed explanation can be found elsewhere (Fidel, 1986).

A single-meaning term can be mapped to a descriptor through an exact match, partial match, or a term might be mapped to a broader descriptor. When a single-meaning term is mapped to a descriptor through an exact match, searchers may use the descriptor [F], or select to consider a variety of characteristics [Z].

Partial match usually implies mapping a request term to a narrower descriptor, or to a group of narrower descriptors. Searchers can always enter the descriptor, or the group of descriptors [G]. If suitable, however, they use a free-text key to inclusively search concepts that are not grouped together by the hierarchy of the controlled vocabulary [H]. If, for example, the request term *students* is mapped to descriptors such as *Foreign Students*, *College Students*, or *Undergraduates*, and a descriptor *Students* does not exist, the free-text key can be used to retrieve information about any type of student. It should be noted that in some search systems use of the free-text term *students* also would retrieve citations that are indexed with descriptors which include the term. This is a source for constant confusion for searchers who search a number of search systems regularly because the routine changes from one search system to another.

When a single-meaning request term is mapped to a broader descriptor, searchers may prefer to preserve the specificity of the request and use free-text search keys [I]. If they are concerned with the precision of the set to be retrieved, they enter free-text terms in combination (using the AND operator) with the broader descriptor to which it is mapped [J]. If, for example, the concept *final examination* is mapped to the descriptor *Instructional Tests*, searchers can AND this descriptor with the free-text term *final* and thus eliminate citations for articles about other types of examinations. Searchers, of course, may elect to enter just a descriptor even if it is broader than the original concept [K].

This concludes the description of the conditions and options presented in the Selection Routine. Our next question is how can this decision tree be used for the construction of IES?

APPLICATIONS TO EXPERT SYSTEMS

The Selection Routine presents the rules that searchers use when they decide which type of search keys to enter. The Routine by itself is not sufficient for an expert system to act like a searcher. IES needs more knowledge for that purpose. An expert system should be able to determine, for example, whether a term is a common term or a single-meaning one, if it is mapped to a descriptor and through what kind of a match, or when to ask patrons whether high precision or high recall is required.

This additional knowledge requires a variety of dictionaries, thesauri, and rules. A few examples of such knowledge and the manner in which it can be stored are discussed below.

Semantic Dictionaries

The most straightforward way for an expert system to check whether a given term is a common term or a single-meaning one is to consult a dictionary that would specify the status of each term. Unfortunately, such a dictionary does not exist and more research is needed before it could be constructed.

What are the issues that should be considered for the construction of such a dictionary?

From a linguistic point of view there are very few terms that are truly single-meaning. It is enough to browse through a regular dictionary to discover that most terms have more than one meaning. Thus the dichotomy common/single-meaning is not supported by linguistics, and most terms are neither the one nor the other.

On the other hand, searchers must make a decision whether or not a given term is "good" for free-text searching. For them the issue is a dichotomy, and they use their common sense and intuition to make the decision. Because the distinction between common and single-meaning terms is central to online searching, an expert system should also be able to make this distinction. Expert systems cannot use common sense or intuition, though, and one must devise rigorous routine for this decision.

Currently, we do not know what is the best way to determine for a given term whether it would be considered a common or a single-meaning term by online searchers. There are probably various methods to teach expert systems to make this distinction but the methods need to be tested and compared.

For example, one possible method to identify common terms for a database is to measure the frequency (the number of times) in which each term occurs in the database, and to designate terms that occur very frequently as common terms. Such a method is attractive because it is easy to implement and it does not require a full-scale dictionary. For each database an expert system needs to have a cutoff point beyond which a term would be considered a common term. A system may have a rule stating, for instance, that if the number of postings for a term is larger than 15,000 the term is a common term. While such a simple procedure is indeed attractive, we still need to test its validity. We need to examine whether common terms occur in high frequency and whether terms that occur in high frequency are always common terms.

Another method is to ask a panel of searchers-judges to consider each term in a database and to decide whether it is a common or a single-meaning one. While this method seems most reliable, it is too labor-intensive to be realistic.

As these examples illustrate, we are still in the dark: We have not found yet a method that is both feasible and reliable. Developing such a method requires extensive investigations. Since no such research is being carried out at present, either in information science or in linguistics, only the future can tell how soon will this issue be resolved.

Thesauri

Many databases have thesauri that list descriptors and lead-in terms: Terms that are not descriptors but are listed in the thesaurus and lead patrons—through "see" references—to descriptors. Searchers consult these thesauri when they are looking for descriptors. Expert systems can use the same thesauri to determine whether or not a term is mapped to a descriptor.

The thesauri to be used by IES, however, should be much more extensive than those used by human searchers, and may require special search algorithms. Ideally, patrons will present IES with requests expressed in natural language and in the form of a sentence. IES in turn should be able to identify the components of a request, and check the terms that represent each component in the thesauri of the relevant databases.

If an online searcher cannot find an entry to a given term, he can think about synonyms or other ways to represent a concept. To "think" about synonyms, expert systems require thesauri with extensive lead-in vocabulary that include any term that might be mentioned in a request.

In addition, the lead-in vocabulary would probably be more complex than that in a manual thesaurus. Some thesaurus look-ups are simple. Consider the concept *examinations* when it is not a descriptor. This term is an example of a simple case: a lead-in vocabulary is likely to have an entry for the term and guide a searcher to the right descriptor, say, *Instructional Tests*. Here, an online searcher and an expert system that can read the thesaurus would perform equally well. Suppose, however, that a request is about attitudes of students towards themselves. Here a searcher may pause to think about the best representation of the concept *attitudes towards themselves*. She may decide that another way to represent this concept is with the phrases *self image* or *self esteem*. In searching for descriptors, then, she would look under these phrases rather than under *attitudes*. An expert system should be able to make similar decisions.

For an expert system to consider different representations for a single concept, a sophisticated lead-in vocabulary with possibly the addition of some semantic rules are required. Like the distinction between common and single-meaning terms, thinking about various representations of a concept is guided by the searcher's linguistic abilities. A knowledge of the English language, and a special talent to manipulate the language, are considered by many an important requirement of online searchers. Expert systems can acquire this talent with the help of research in linguistics. Much research is being conducted in this area (natural language processing), but no conclusive prototype is available yet. It is possible, though, that such capabilities will be common to IES in the imminent future.

Options and Reasons

In addition to solving linguistic and semantic issues, searchers are looking for the search keys that would best satisfy a given request. An example can illustrate this function.

The Selection Routine does not provide a unique path for each set of conditions. For example, if a term is a single-meaning term and it is mapped to a broader descriptor, there are three options (in addition to the Z options): Searchers can use free-text terms [I], free-text terms in combination with descriptors [J], or they can use descriptors [K]. Suppose an expert system mapped a single-meaning term to a broader descriptor, how would it know which option to choose—which search keys to select?

To answer such questions, one must identify additional conditions that would determine under what circumstances should each option be selected. A substantial part of the study reported here was to uncover these additional conditions. It is beyond the scope of this chapter to provide a comprehensive description of these conditions, but one example can convey their nature.

The study found that searchers preferred to use free-text terms to represent a single-meaning term that is mapped to a broader descriptor under the following circumstances: (a) If a patron required retrieval with high precision, or (b) if they entered the broader descriptor and found that many of the citations retrieved were irrelevant. It was also found that the first condition occurred more frequently than the second one.

Alternately, searchers decided to enter a free-text term in combination with a descriptor (e.g., *final* AND *Instructional Tests*) only when they did not trust the indexing—thinking the descriptor was assigned loosely—and wanted to make sure that the specific term was present in the citations. This option, however, was rarely selected.

Searchers selected most often the third option: Once they had identified a descriptor for a single-meaning term, they entered it even if it was a broader descriptor. The most common reason for that selection was their knowledge that high recall was required. Another reason that was mentioned less frequently, but was still quite common, is the searcher's belief that it was better to search with descriptors. (It should be noted here that this argument was used consistently for particular databases.)

Three other reasons for selecting a broader descriptor that were mentioned by searchers relate to more specific situations. In a few instances searchers entered a broader descriptor after realizing that the free-text term had not provided enough retrieval. A few others used two broader descriptors that in combination represent the specific concept. For example, the descriptors *Instructional Tests* and *Graduation Requirements* are each broader than the term *final examinations*. Their combination with the AND operator, however, represents more closely the required concept. One searcher claimed that if the term is a single-word term, precision and recall would be improved if a broader descriptor was entered.

Once the reasons for selecting each option in the Selection Routine are laid out, one must ask: How can an expert system apply this knowledge when it selects search keys? IES will have a set of rules that will guide them in selecting options. To illustrate the nature of these rules, it is best to use an example.

An Example

Consider our sample request: "the analysis of students' behavior during a final examination to determine the difficulty of the examination." Suppose an intermediary expert system—named THE GOOD SEARCHER—has acquired the terminological and semantic talent needed for the analysis of requests. This system is able then to identify the components of the request, which are: (a) analysis, (b) students' behavior, (c) final examinations, and (d) examination's difficulty.

After looking in the thesaurus of an educational database the system finds out that:

1. The first concept, which is a common term, is mapped to the descriptor *Analysis*—an exact match.
2. The second concept is also mapped to a descriptor through an exact match.
3. The third concept is mapped to a broader descriptor—*Instructional Tests*.
4. The fourth concept cannot be mapped to a descriptor.

Following the Selection Routine, THE GOOD SEARCHER selects to represent the first two concepts with descriptors (options [A] and [F]). For the representation of the third concept, however, the system needs to make a selection among three options: [I], [J], and [K]. For that purpose it needs to know what are the precision and recall requirements, as described earlier. In response to the system's inquiry, the patron indicates that this is a high recall request. Following the reasons for the three options, THE GOOD SEARCHER decides to use the descriptor *Instructional Tests* (option [K]).

The last concept is a little more problematic because it is not mapped to a descriptor. While this condition raises four options ([L], [M], [N], and [O]), the system is programmed to always try option [M] first: use the free-text key to probe indexing.

The initial formulation of the request is then:

Analysis AND *Students' Behavior* AND *Instructional Tests* AND *difficulty*.

Because the first formulation was entered to probe indexing, THE GOOD SEARCHER asks the patron to provide some feedback. The system displays the 10 citations that have been retrieved and asks the patron to designate which citations are highly relevant. The patron marks six citations. The system's task now is to find descriptors that are assigned to these citations most frequently. It comes up with one descriptor: *Academic Standards*. The fourth concept of the request—examination's difficulty—can be now mapped to a descriptor.

The next decision the system is facing is whether to represent this concept only with a descriptor or whether to add the free-text term. Here, a general rule can be of help. It states: If more than 30% of the citations retrieved by the free-text key are relevant, use both descriptor and free-text; if less than 30% of the citations are relevant, use the descriptor only. Therefore, the system decides to represent the fourth concept with the combination: *difficulty* OR *Academic Standards*.

Using the modified formulation, THE GOOD SEARCHER finds 25 citations. Now it needs to find out whether the patron is satisfied with the retrieved set. After displaying the titles of the retrieved citations, it asks whether the retrieval seems relevant and whether there is enough. The patron is satisfied with the precision of the answer set but feels that some are missing: Higher recall is needed. Checking the options available to increase the recall of a retrieved set, the system discovers the rule: If a request has more than three components and higher recall is required, one request component should be eliminated. (This rule was not derived from the Selection Routine but from other stored knowledge about searching behavior.) The system cannot randomly select which concept to drop, however, and it asks the patron to rank the concepts in the order of their importance.

The patron response in decreasing order is: *students' behavior*, *final examinations*, *analysis*, and *examination difficulty*. Following its own rule, THE GOOD SEARCHER drops the fourth concept from the formulation and retrieves 53 citations.

To the system's request for evaluation of titles of retrieved citations, the patron asks that recall still be improved and expresses her satisfaction with precision. There are various options, and the system may select all or just some of them, depending on the specifics of the request. Under one situation it can add to the formulation the free-text term *final exam* [Z6], and under another it can retrieve all citations that have been indexed under a descriptor that includes the term *student*. (For simplicity's sake the specific situations are not explained here.) The system also considers option [B] for the concept *analysis* because it is a common term that is mapped to a descriptor, but it is not central to the request. In other words, *analysis* is a limiting factor—it qualifies the main subject of the request—and could be entered as a free-text term to increase recall.

After eliminating the other options to improve recall—through an interrogation of the patron—THE GOOD SEARCHER selects option [B] and retrieves 98 citations. At this point the patron is satisfied with both precision and recall, and the search is terminated.

CONCLUSIONS

What can we learn from the Selection Routine and its implementation for expert systems?

The first conclusion that stands out is that linguistic rules and semantic networks are essential to the operation of IES. No expert system could truly replicate the performance of human searchers before it can resolve semantic and linguistic issues.

Second, the Selection Routine demonstrates that once linguistic and semantic issues are resolved, IES can indeed be constructed to replicate the searching behavior of human searchers.

Third, such IES can identify which characteristics of a request are important to the success of a search and at which points in the search should the patron be asked about them. Most importantly, such IES can act on additional information they receive from patrons to improve search results.

Fourth, in some instances, such IES can *surpass* human searchers. IES have one substantive advantage over human searchers: They can process information much faster. Consider, for instance, the option [P]: to use free-text terms when it is *not known* whether a term is mapped to a descriptor. Obviously, it is better to know whether or not a term is mapped to a descriptor, but searchers decide not to check a term for descriptors because of various reasons. They may decide not to consult the thesaurus of a certain database because they think indexing—or the vocabulary itself—are not useful, because they are not available to them or because they have to search a request in a large number of databases.

While the first reason for avoiding a thesaurus look-up is justified under any circumstances, the rest are determined by cost-benefit considerations that are valid for human searchers but not for IES. An intermediary expert system should include all the thesauri of the databases it can search, and it should consult each of them for a request, regardless of the number of databases to be searched. While it probably should avoid consulting a thesaurus that is considered unreliable by most searchers, an expert system can perform better than a human searcher for a request that requires the use of a large number of thesauri.

In addition, IES are free from developing searching habits that might be counterproductive at times. A searcher may develop a habit of searching with free-text keys because most of the time this mode is required for the success of her searches. This searcher would probably elect to enter free-text keys, say, to search a database she rarely searches, even if a descriptor search is likely to produce better results. Because IES are based on the experience of a large number of searchers such individual idiosyncrasies are not likely to be incorporated into their knowledge. IES, then, can always select the path that seems most productive, free from individual habits.

Will IES replace human searchers?

Even though they can become very powerful, it is unlikely that the introduction of IES will eliminate the need for online searchers. It is obvious, however, that the role of online searchers will change and their expertise will be used differently. But this is another matter.

REFERENCES

Doszkocs, T. E. (1983). Automatic vocabulary mapping in online searching. *International classification, 10*, 78–83.

Fidel, R. (1986). Towards expert systems for the selection of search keys. *Journal of the American Society for Information Science, 37*, 37–44.

Marcus, R. S. (1983). An experimental comparison of the effectiveness of computers and humans as search intermediaries. *Journal of the American Society for Information Science, 34*, 381–404.

Pollitt, A. S. (1984). A "front-end" system: An expert system as an online search intermediary. *Aslib Proceedings, 36*, 229–234.

Svenonius, E. (1986). Unanswered questions in the design of controlled vocabularies. *Journal of the American Society for Information Science, 37*, 331–340.