

Searchers' Selection of Search Keys: I. The Selection Routine

Raya Fidel

Graduate School of Library and Information Science, University of Washington, Seattle, WA 98195

The selection routine is a formal decision tree that represents the intuitive rules searchers use when they select search keys, textwords or descriptors. The case study method provided the data through: (1) observation of 47 professional online searchers performing their job-related searches; and (2) analysis of verbal and search protocols involved. Each option in the selection of search keys presents the use of a certain combination of textwords and descriptors which searchers choose because of request or database requirements, or because of their own beliefs. The routine delineates the terminological conditions which lead to the selection of each option. It is the first formal presentation of human knowledge that can be incorporated into the knowledge base of intermediary expert systems.

Introduction

Online searching behavior has attracted much attention among researchers because of the current discrepancy between the level of technological developments as compared to theoretical advancements. New and increasingly sophisticated technology is being developed and put to use at an ever-growing rate, but the scientific understanding of human-machine interaction and of the search process is in its infancy (Saracevic et al., 1988). On the positive side, the widespread use of technology provides a real-life laboratory for studying online searching behavior.

The research project reported here investigated online searching behavior manifested by *actual* searches of bibliographic databases, as performed by *professional* online searchers. The study explored the process of search-key selection, and attempted to represent this process in an empirically based model that is specified in formal terms. Such a model is valuable for basic

research, for the training of online searchers, and for the design and development of information retrieval systems that can be searched by end users.

Problem Definition

One of the tasks in an online search 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.

Reality of online searching is more complex than it may seem. To start, the information need of a user is often ill-defined and difficult to determine accurately. For the purpose of this research project, however, we assumed that users expressed to the intermediaries information needs that *are* defined.

Further, users' requests were regarded as having two major aspects: semantic and pragmatic. The semantics of a request is the *topic* of a request; it presents the subject matter that is of concern to the user. For example, "the analysis of students' behavior during a final examination to determine the difficulty of the examination" is the topic of a hypothetical request.

The pragmatic aspect of user requests concerns the *purpose* of a request, or the *use* to which the information will be put. For example, depending on the anticipated use, a user may need a comprehensive search that retrieves all the relevant citations, or she may be interested in just a few highly relevant citations, or only in recent citations. Usually, searchers interview users to clarify both the semantic and pragmatic aspects of each request.

The classic online search includes the following procedure. Once the searcher understands the request well enough to answer it, a plan is developed for the search—a search strategy. This strategy specifies which databases will be searched and which terms (or search keys) will be used to search each database. It can also include a more specific plan that determines the flow of the search: Which search keys to enter first, when to review some results, and what to do if the results are not satisfactory. Next, guided by the search strategy, the searcher actually performs the search and retrieves

This article is based upon work supported by the National Science Foundation under Grant Number IST 85-09719.

Received January 17, 1990; revised April 30, 1990, June 27, 1990; accepted September 7, 1990.

©1991 by John Wiley & Sons, Inc.

citations, but the searcher may deviate from the original plan if it does not seem useful. Some requests may require a number of terminal sessions. A searcher may logoff to reconsider the strategy, possibly with the help of the user. At some point, the searcher decides to terminate the search and to print the answer set that will be given to the user.

Thus, the intellectual components of a typical online search can be classified into three basic categories: (1) definition of query structure; (2) selection of search keys; and (3) feedback review. The second category, the selection of search keys, is the focus of this study.

To select search keys for a request, a searcher must first break down a request into its individual components, or concepts. Consider the request: "What analyses were used to determine if students' behavior during a final examination is related to the difficulty of the examination?" It includes four concepts: (1) analysis; (2) students' behavior; (3) final examinations; and (4) examination difficulty. Each concept requires a set of search keys for its representation. Thus, the searcher looks for search keys that will best capture the literature on the topic of each individual concept or of the concepts in combination, and at the same time retrieve an answer set that satisfies other request characteristics, such as recall, precision, or timeliness.

There are two distinct types of search keys: textwords which are used in free-text searching and descriptors from a controlled vocabulary. In many databases, both options are available. They can also be used in combination: textwords for some concepts and descriptors for others or both types of search keys for the same concept.

We studied the *decisions* that searchers made when they selected search keys, either when devising search strategy or during a terminal session. We aimed to uncover the reasons for the selection of each *type* of search key. The study resulted in a formal model describing these decisions. This model could be incorporated into the knowledge base of intermediary expert systems.

Intermediary Expert Systems

Intermediary expert systems attempt to provide a powerful form of user assistance: They seek to 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.

Studies examining users searching their own requests with no intermediary assistance show repeatedly that while users seem to master the command language with no difficulties, they lack the expertise needed for formulating search strategies (e.g., Sewell & Teitelbaum, 1986; Kirby & Miller, 1986). Therefore, every intermediary expert system that is being developed today must include a component that supports decisions about

search strategies and, in particular, about the selection of search keys.

In a database that offers both controlled vocabulary and free-text searching, such systems must examine each concept in the request topic and consider its representation as a descriptor, as a textword, or as both. An intermediary expert system should be able to take into account request (and user) characteristics that are beyond the topical description of the search but important for the selection of search keys. The purpose of this study is to provide a model that is pertinent to the implementation of this requirement.

The importance of human expertise to the design of intermediary expert systems is still controversial. It seems, however, that the notion that intermediary systems should be based on knowledge acquired from human experts is gaining increased recognition. Croft, for example, maintains that the formalization of the knowledge used by human intermediaries is one of the open problems of research in expert systems for information retrieval (Croft, 1987), and Daniels mentions it as the most promising method for the construction of user models (Daniels, 1986). In addition, a few prototypes, such as PLEXUS (Vickery et al., 1987), and EX-P (Smith et al. 1987), are already based on knowledge acquired from human experts.

The research reported here analyzed searching behavior of human intermediaries and then presented this behavior in a formal model. It thus represents the first step in incorporating experience gained by human intermediaries into knowledge bases of intermediary expert systems.

The Objectives of the Study

To begin a systematic investigation of searching behavior, I first completed a study of online searching behavior using the case study method (Fidel, 1984). I observed eight experienced human intermediaries doing their regular, job-related searches, and recorded their spoken thought processes. Analysis of data collected in this preliminary study uncovered the "selection routine," which is a presentation of rules for the selection of search keys in the form of a decision tree (Fidel, 1986).

While the selection routine clearly indicated that formal rules could be extracted from human experts, it was incomplete in the preliminary study. First, there were a number of conditions that led to more than one option. For example, if a term was a common term, that is, not appropriate for free-text searching, and it was not mapped to a descriptor, one was left to decide whether to use textwords to probe indexing, or whether to change database. Clearly, there might be additional conditions that would determine which of these options to select, but these conditions were not revealed by the preliminary study. Secondly, the eight searchers who were observed for that study were experts in the life

sciences literature. To build a selection routine that is applicable to bibliographic retrieval in *every* subject area, the searching behavior of human intermediaries in a *variety* of subject areas had to be investigated.

Thus, the objective of this study was to refine and validate the selection routine. A description of this routine is given following a discussion of the method.

The Method

The case study method with controlled comparison (Diesing, 1971) was used to investigate the selection of search keys. Briefly, in this method a case is analyzed to construct a model of the investigated phenomenon based on one case. An additional case, which is similar in a definite sense to the first case, is then analyzed and is fitted into the model created by the first one. Discrepancies are resolved either by increasing the level of generality in which the elements of the model are expressed, or by adding elements to the model. The modified model of the investigated phenomenon is now based on two cases. Additional cases are analyzed, one after the other and representing a gradual increasing diversity, to further refine the model and to expand its applicability. Models constructed by the case study method with controlled comparison are never complete in an absolute sense: The more cases are analyzed, the more general the model becomes. Such models are dynamic, however, in that they can be modified and expanded to fit new developments and discoveries in the investigated phenomenon. A detailed description of the use of this method in the investigation of online searching behavior is available elsewhere (Fidel, 1984).

The data for this project were collected through observation and interviews. Each member of the research team observed searchers when they were doing their regular, job-related searches and asked them to think out loud as they worked. The searchers' spoken words were recorded and transcribed, and together with certain written material, such as the search protocol and the request form, served as a basis for the analysis, which was primarily a protocol analysis.

Protocol analysis was used in the project to identify and analyze each instance in which a search key was selected. Once such an instance was established, verbalizations of thought processes, previous and preceding moves in the search, and recorded search strategy were used to explore the conditions that led to that particular selection.

Each member of the research team interviewed each searcher he or she observed immediately after the sequence of observations for that searcher had been completed. Before the interview, the whole team analyzed all the transcribed protocols of searches performed by a searcher to identify issues that were inaccessible to observation or those that needed clarification. In the interviews, searchers were asked to explain their reasons

for selecting individual search keys. The interviews were then transcribed. Answers of searchers were checked for validity by comparing them with other types of evidence. For example, if a searcher explained that a certain search key was selected to improve recall, the search protocol was examined to ascertain that the search key was indeed used for that purpose.

Search protocols were systematically analyzed, one after the other, to identify incidents where a search key was selected. Each such incident was then fitted into the decision tree, following the method of controlled comparison. This method facilitated ongoing modifications to the selection routine.

The study team selected for observation 39 experienced online searchers who had at least two years of searching experience, and ordinarily searched databases that provide both free-text and descriptor searching. Further, to improve the generality of the model, searchers were selected from a wide spectrum of subject specialties. The sample included 13 searchers from the humanities and social sciences, 21 from the science and technology area, two medical librarians, and three public librarians who could be called upon to search in any subject area.

Each searcher was asked to allow observation for five searches. While most searchers were actually observed for five searches, the number of searches per searcher, in a few cases, varied from four to eight. The total number of actual searches analyzed for this project is 201.

A detailed description of data collection and analysis can be found elsewhere (Fidel, 1988).

The Selection Routine

The number of modifications to the original selection routine was not large. Although the original routine has already been described in detail (Fidel, 1986), the refinements introduced in this study require a complete description of the revised version here. That is, the selection routine presented here is based on the analysis of searching behavior of 47 searchers: those who participated in the preliminary study and those who participated in this study.

The modified selection routine is presented in the form of a decision tree (Fig. 1). The refinements that were introduced, that is, the new options, are presented in the description of the selection routine which follows. In addition, Table 1 lists the options in the selection routine and the associated conditions.

The first criterion for decisions about the selection of search keys is whether a term is a common term or a single-meaning term. A single-meaning term is a term which is "good" for free-text searching. It usually occurs in a particular context, it is uniquely defined, and it is specific to the concept it represents. A common term, on the other hand, is a term that is not suitable for free-

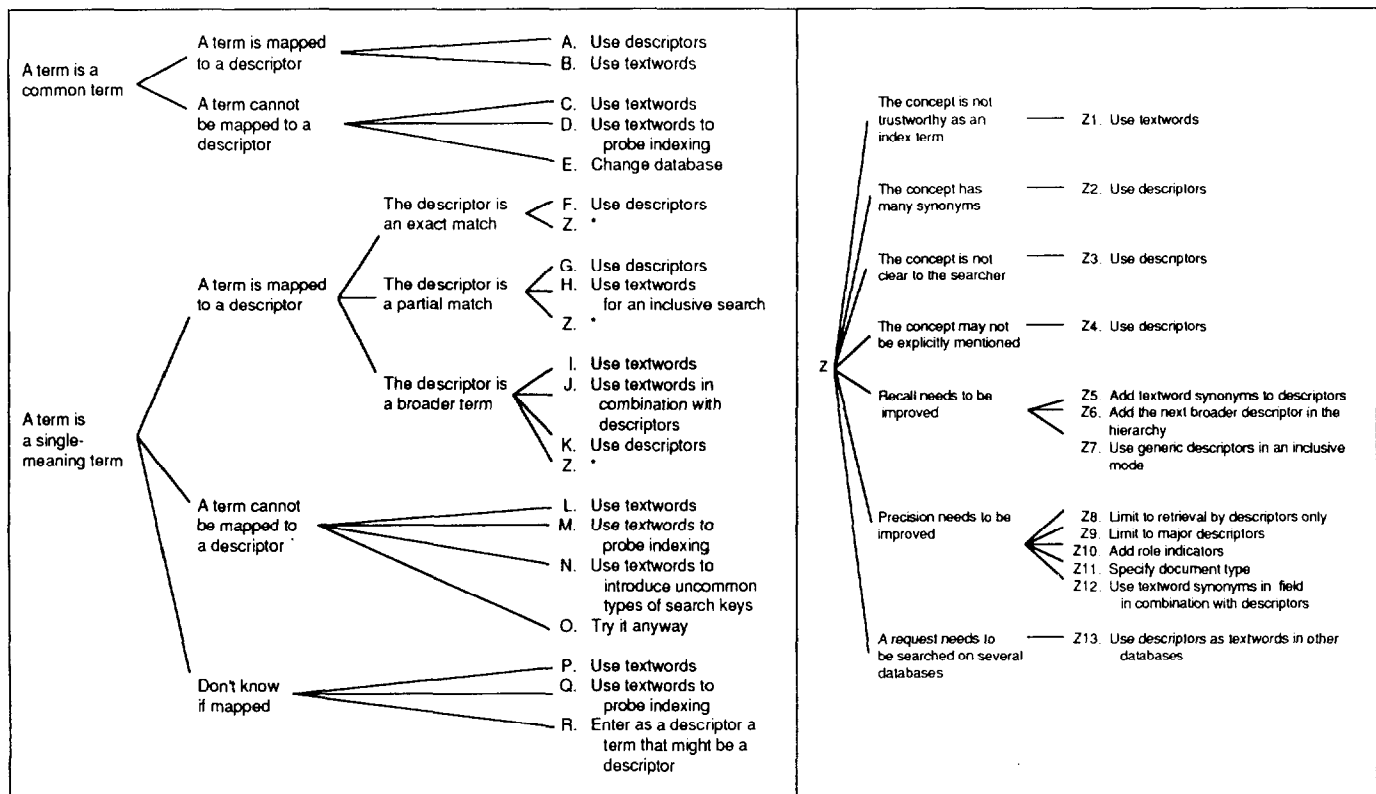


FIG. 1. The selection routine.

text searching. Such a term usually occurs in *more* than one context, or it has a broad and fuzzy meaning.

For example, in the request about the analysis of students' behavior during final examinations, the terms *students* and *behavior* are single-meaning terms. By contrast, *analysis* and *examination* are common terms because they may represent different concepts, depending on the context. To be more specific, the term *examination* can occur in a subject-related context ("the best way to take a student examination"), being synonymous with *test*. It can be used to represent the concepts of *perusal* or *study* ("examination of students' responses"), in which case the term *examination* could appear in titles and abstracts of articles that are about other subjects. Further, it can be used very loosely to represent the concept of an inquiry of any kind.

The definitions of "single-meaning terms" and "common terms" are pragmatic: they depend on the terminology used in the searched text, and they are not grounded in linguistic or philosophical theories. For example, a term might be a single-meaning term in one database but a common one in another. Though pragmatic, the nature of the distinction between single-meaning terms and common terms warrants further investigation.

The second criterion for the selection of search keys is whether or not a term that represents a request concept is mapped to a descriptor. A searcher maps a term to a descriptor when she has decided that this descriptor best represents the concept, whether or not there is

an exact match between the term representing the concept and the descriptor. This criterion generates three conditions: a term is mapped to a descriptor, a term cannot be mapped to a descriptor, and the searcher does not know if the term can be mapped.

These two criteria—whether a term is a single-meaning or a common term and whether or not it can be mapped to a descriptor—are central to the selection routine because they deal with the relationship between concepts and terms: the concepts that need to be represented and the terms that can express them. Since controlled vocabularies are designed to resolve problems in expressing concepts in query formulations, it is important to examine these relationships when analyzing the selection of search keys. That these two criteria are central does not imply, however, that they are always used by searchers first and before they examine other factors, such as the constraints of the request or of the database. The *priority* given to criteria used in the selection of search keys is situational: it may be determined by the nature of each request, or by the searcher's individual preferences. The question of priority was not examined in this study.

The selection routine, as presented in Figure 1, is formulated with respect to terminological considerations. Searchers who participated in the study, however, mentioned reasons, other than terminological, for the selection of search keys. These reasons fell into three categories: request-related, database-related, and

TABLE 1. Options in the selection routine.

Option	Conditions	Option	Conditions
Descriptor Searching		Free-Text Searching, cont.	
Use descriptors	<p>A term is a common term + it is mapped to a descriptor [A].</p> <p>A term is a single-meaning term + it is mapped to a descriptor + the descriptor is an exact match [F].</p> <p>the concept has many synonyms [Z2].</p> <p>the concept is not clear to the searcher [Z3].</p> <p>the concept may not be explicitly mentioned [Z4].</p> <p>the descriptor is a partial match [G].</p> <p>the descriptor is a broader term [K].</p> <p>it cannot be mapped to a descriptor [O].</p> <p>it is not known if mapped [R].</p>	Use textwords to probe indexing	<p>A term is a common term + it cannot be mapped to a descriptor [D].</p> <p>A term is a single-meaning term + it cannot be mapped to a descriptor [M].</p> <p>it is not known if mapped [Q].</p> <p>A term is a single-meaning term + it is mapped to a descriptor + a request needs to be searched on several databases [Z13].</p>
Add the next broader descriptor in the hierarchy	A term is a single-meaning term + it is mapped to a descriptor + recall needs to be improved [Z6].	Use descriptors as textwords in other databases	A term is a single-meaning term + it is mapped to a descriptor + the descriptor is a partial match [H].
Use generic descriptors in an inclusive mode	A term is a single-meaning term + it is mapped to a descriptor + recall needs to be improved [Z7].	Use textwords for an inclusive search	A term is a single-meaning term + it cannot be mapped to a descriptor [N].
Limit to retrieval by descriptors	A term is a single-meaning term + it is mapped to a descriptor + precision needs to be improved [Z8].	Use textwords to introduce uncommon types of search keys	Other Combinations
Limit to major descriptors	A term is a single-meaning term + it is mapped to a descriptor + precision needs to be improved [Z9].	Use textwords in combination with descriptors	A term is a single-meaning term + it is mapped to a descriptor + the descriptor is a broader descriptor [J].
Specify document type	A term is a single-meaning term + it is mapped to a descriptor + precision needs to be improved [Z11].	Add textword synonyms to descriptors	A term is a single-meaning term + it is mapped to a descriptor + recall needs to be improved [Z5].
Free-Text Searching		Add role indicators	A term is a single-meaning term + it is mapped to a descriptor + precision needs to be improved [Z10].
Use textwords	<p>A term is a common term + it is mapped to a descriptor [B].</p> <p>it is not mapped to a descriptor [C].</p> <p>A term is a single-meaning term + it is mapped to a descriptor + the concept is not "trustworthy" as an index term [Z1].</p> <p>the descriptor is a broader term [I].</p> <p>it cannot be mapped to a descriptor [L].</p> <p>it is not known if mapped [P].</p>	Change database	A term is a common term + it cannot be mapped to a descriptor [E].
		Use textword synonym in a designated field in combination with descriptors	A term is a single-meaning term + it is mapped to a descriptor + precision needs to be improved [Z12].

searcher-related. The last category includes general rules or assumptions that were habitually used by an individual searcher. We turn now to the description of the selection routine.

A Term is a Common Term

When the study searchers encountered a common term they almost always looked for a descriptor to represent it. A common term, however, may or may not be mapped to a descriptor.

A Common Term is Mapped to a Descriptor. When a common term is mapped to a descriptor, searchers almost always entered the descriptor as a search key [A] (i.e., option [A] in Figure 1) because, by definition, it is not desirable to enter a common term as a textword key.

There is one exception to this rule: Searchers may decide to enter the term as a textword [B] when the term was used as a limiting factor, and they perceived that a descriptor might be too restrictive. For example, in the request about the analysis of students' behavior during final examinations, a searcher combined the terms *students' behavior* with *final examinations*, using the AND operator. Adding the requirement that all citations be also indexed under the descriptor *analysis* might be too limiting, and the searcher decided to use *analysis* as a textword key, a somewhat less restrictive requirement.

A Common Term cannot be Mapped to a Descriptor. A common term that cannot be mapped to a descriptor almost always results in unsatisfactory retrieval, even when used in combination. Searchers, however, have almost no choice but to enter a textword [C]. Although searchers can enter such a textword just to check the indexing of relevant articles, two reasons were cited for a direct use of a common textword key. The first related to the request and the second related to the database searched. First, if a request includes a relatively large number of concepts—that is, the Boolean operator AND occurs more than two or three times in the query formulation—precision will not suffer if a common term is entered as a textword [CR1] (i.e., the first Request-related reason for option [C]). Second, if a request will be searched on a number of databases, it might be too costly to ascertain the relevant indexing in each database [CD1] (i.e., the first Database-related reason for option [C]).

If a request requires searching only one or two databases, however, searchers can enter the textword to probe indexing [D]. One method of probing the indexing is to enter the textword key in combination with other search keys, in order to retrieve citations, to select some relevant ones, and to review their indexing in an attempt to find descriptors that might possibly be relevant. For example, if the term *examination* cannot be mapped to a descriptor, one can devise a formulation (using the AND operator) that combines the descriptors *students*, *analysis*, and the textword *final* and

examination. Reviewing a sample of retrieved citations, one may find that all the relevant citations include the descriptor *instructional tests*, 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 in which the common term is mapped to a descriptor [E].

A Single-Meaning Term That is Mapped to a Descriptor

When a single-meaning term is mapped to a descriptor, it can be mapped through an exact match, through a partial match, or to a broader descriptor.

When the Descriptor is an Exact Match. The most direct use of a descriptor to represent a single-meaning term is when a term is exactly matched with a descriptor [F].

When the Descriptor is a Partial Match. A descriptor is a partial match when it includes the request term but other terms as well, in which case the descriptor is usually narrower in meaning. For example, the term *tests* is mapped to the descriptor *educational tests* through a partial match. Searchers may elect, however, to substitute the descriptor for a request term anyway [G]. They select this option because: (1) the term has been added by ORing it with other terms in one facet of the formulation in order to increase recall [GR1]; (2) the descriptor was spotted as an index term assigned to relevant articles [GD1]; or (3) the searcher prefers to use descriptors and the selected one is the best match [GS1] (i.e., the first Searcher-related reason for the option [G]).

When it seems appropriate, however, searchers use a textword key to inclusively search concepts that are not grouped together by the hierarchy of the controlled vocabulary [H]. This option is selected exclusively to improve recall [HR1]. If, for example, the request term *students* is mapped to descriptors such as *foreign students*, *college students*, or *undergraduates*, and the descriptor *students* does not exist, the textword *student* can be used to retrieve information about almost any type of student.

It should be noted that in many search systems, such as DIALOG or BRS, use of the textword *student* also would retrieve citations that are indexed with descriptors which include the term. In other systems, entering such a term retrieves only citations whose indexing includes this term. This is a source for constant confusion for searchers because routines change from one search system to another, and in one search system over a period of time.

The Descriptor is a Broader Term. When a single-meaning request term is mapped to a descriptor broader in meaning, searchers may prefer to use textwords because they correspond to the request component more accurately [I]. Some searchers do so to increase

precision [JR1]. In contrast, other searchers select this option because they believe that, generally, the use of textwords increases recall [IS1].

Further, a concern for precision may lead searchers to AND textwords with the broader descriptor to which the request term is mapped [J]. While precision is an important reason for the selection of this option [JR1], searchers may also use such a combination if they do not trust the indexing of the database [JD1].

Searchers, of course, may choose to enter the broader descriptor alone [K]. Most often they select this option to increase recall [KR1]. Entering a broader descriptor for one concept of a request is useful for recall enhancement in a variety of circumstances. Searchers may want to have an initial set that is broad because the request includes a relatively large number of concepts, or because the combination that is required by the request is especially limiting (if, say, the concepts are not likely to occur together). Another situation which calls for a broader descriptor is when an inclusive search, one that requires entering a descriptor as well as its narrower terms, is required to secure recall. For instance, if the request term *disabled students* cannot be mapped to an exact descriptor, a searcher may enter the broader descriptor *students* and add (using the OR operator) its narrower terms, descriptors such as *college students* or *undergraduates*, to perform an inclusive search. Inclusive searching is sometimes called "exploding" or "cascading."

Depending on the terms, searchers may enter a broader descriptor also as a precision strategy [KR2]. Suppose in a request about the income of college presidents the term *presidents* is mapped to the broader descriptor *administrators*. A searcher may prefer to search under the descriptor because the citations retrieved would be about college administrators and their income, hopefully some about presidents. The set retrieved with the exact textword, on the other hand, is likely to include many citations for items about the income of other college employees—an income that is determined by the president. Under such circumstances, the searchers perceive that, used as a textword, the particular term would generate a set with low precision (if, for example, it is a single-word term). In addition, searchers may enter a broader descriptor when it is used only as a factor that limits or qualifies the request [KR3].

The indexing in a particular database may also help searchers to select a broader descriptor. They may enter such a descriptor if it is found in the indexing of relevant citations [KD1], or because they generally prefer to use descriptors [KS1].

Additional Factors. A single-meaning term that is mapped to a descriptor, through any kind of match, provides searchers with more choices than those provided by terms not so mapped. If searchers think that a particular descriptor is assigned inconsistently by indexers, they may consider the use of the textword key to be more trustworthy [Z1]. Or, they may prefer to

enter a descriptor when: a term has many synonyms [Z2]; a concept and its use are not clear to the searcher [Z3]; or a concept is likely to be implied rather than explicitly mentioned in the searched text [Z4].

More options for meeting recall and precision requirements exist when there is a match between a term and one or more descriptors. Search keys can be used to increase recall in three ways: a searcher may add textword synonyms to descriptors [Z5]; add the next broader descriptor in the hierarchy [Z6]; or use generic descriptors in an inclusive mode [Z7].

Searchers elect to *increase recall* by adding textword synonyms to a descriptor when they see the need to complement indexing [Z5R1]: They want to include citations that mention the concept, in either titles or abstracts, even though the descriptor was not assigned to them. For some searchers this is the most straightforward approach to ensure recall: When a term is specific, they require that as a search key it would occur in the descriptor, title, and the abstract fields. Searchers in the study selected this option at times because the user, who was present at the terminal, specifically insisted on using textwords as well as descriptors [Z5R2].

Database-related considerations may also lead searchers to the selection of the option that combines a descriptor with textword synonyms. Searchers may decide to use textword synonyms because they plan to search a number of databases [Z5D1] and wish to use the same query formulation across databases. Or, they may add textword synonyms because they do not trust the indexing [Z5D2].

In contrast, adding the next broader descriptor in the hierarchy is selected as an option only when the searcher thinks that the user will be interested in material indexed by the broader descriptor as well [Z6R1].

The use of generic descriptors in an inclusive mode might be desirable for a number of reasons. When searchers create a set that they wish to combine with other sets, using the AND operator, in order to limit the scope of the retrieval, they may use a generic descriptor for that set so the limiting set is not too restrictive [Z7R1].

Databases and their thesauri also play an important role in the choice of this option. A searcher who is interested in material about undergraduate students, for example, may want to secure high recall and retrieve all citations which are indexed under any descriptor which includes the term *students*, whether or not the specific descriptor *undergraduate students* or the broader descriptor *students* exist [Z7D1]. Obviously, this is a specific use of the generic search: it can be carried out only for multiwords phrases and when a part of the phrase is generic by nature.

Inclusive searching might be favored when searching databases which specifically recommend it and provide commands that perform such searching automatically. In these databases, a single command retrieves all the

citations with descriptors that are narrower than the descriptor entered.

Searchers can elect to *increase precision* by limiting a search to descriptors only [Z8], or by limiting it to major descriptors [Z9]. The first option ensures that the articles whose citations are retrieved indeed deal with the subject matter, rather than merely mention it [Z8R1]. The second is used to reduce the number of citations retrieved [Z9R1], or to make sure that a concept is central to the articles whose citations are retrieved [Z9R2].

Additional means to increase precision are to introduce role indicators [Z10], to specify document type [Z11], and to use textword synonyms in the title field in combination with descriptors [Z12]. The last option is considered by some to be a quick way to extract a subset that includes citations that are highly relevant from an already relevant set [Z12R1]. For example, one may extract a highly relevant subset from the set retrieved with the descriptor *students* by adding the requirement that the term *students* appears in the titles of the articles as well. In addition, searchers who do not trust the indexing of a particular database might choose this option [Z12D1].

A Single-Meaning Term That is Not Mapped to a Descriptor

When a request term cannot be mapped to a descriptor, the most direct option is to enter the term as a textword key [L]. Searchers, however, have other choices: They can enter a textword to probe indexing, or they can try and enter the term as a descriptor anyway when they assume they had difficulties in locating a descriptor. It is important, therefore, to examine the reasons for entering a textword directly without trying the other options.

A number of request-specific conditions may encourage a searcher to enter a textword directly. A searcher may do so if he or she believes that most specific retrieval is desired [LR1], or if the term itself is specific and well defined, that is, a term that is "ideal" for free-text searching [LR2]. The latter argument was frequently advanced by searchers when the term was a multiword phrase and it was possible to use word-proximity operators.

Further, searchers may detect a textword during the online session and add it because it appears in titles or abstracts of relevant citations or because it is commonly used in the literature. The searcher may add it as a new concept, using the AND operator, to increase precision [LR3], or OR it with other terms in an existing concept to increase recall (e.g., names of particular examinations) [LR4]. In addition, searchers may enter textwords if the use of related descriptors results in a poor retrieval [LR5].

The nature of the controlled vocabulary for a database is also an important factor in the selection of textwords. A searcher may enter a textword key di-

rectly, rather than probe indexing, believing that the term would not be a descriptor [LD1]. This would happen when: a thesaurus excludes a specific type of term such as geographic names or other proper names; the concept belongs to a subject area that is not covered by the thesaurus; or the thesaurus is outdated and, therefore, would not include terms that represent "new" concepts. Further, searchers who do not trust the thesaurus' vocabulary or the indexing in a database may prefer to enter textwords directly [LD2].

Some searchers have adopted general guidelines that they apply whenever a term is not mapped to a descriptor. These include: if a term represents a concept accurately there is no need to probe indexing [LS1]; searching with textwords is best for high recall [LS2]; and terms that have been suggested by users can be entered as textwords with no further probing [LS3].

Searchers who do not hold to such guidelines, on the other hand, would enter textwords only to probe indexing, hoping to find descriptors that were assigned to relevant citations [M].

In some cases, searchers may use a textword key to search for a single-meaning term that cannot be mapped to a descriptor in a particular way: They require that it occurs in a field other than the common ones, such as the journal title field [N]. Suppose a user is interested only in the psychological aspects of students taking final examinations, and suppose that the term "psychology" cannot be mapped to a descriptor. Searchers may predict that searching for the occurrence of "psychology" 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 "psych" in the titles, or articles that were published in sources whose titles include this stem.

After unsuccessful attempts to find a descriptor, searchers may enter a request term as a descriptor, even believing it does not appear in the thesaurus [O]. They would choose this option either because they assume that the term might have been added to the thesaurus without their knowledge (for instance, before the supplements have been published) [OD1], or because the term is a descriptor in another database [OD2].

It is Not Known If a Term is Mapped to a Descriptor

When searchers elect not to check the thesaurus for a descriptor, they may: enter textwords directly [P]; use textwords to probe indexing [Q]; or enter as a descriptor a term that *might* be a descriptor [R].

Entering Textwords Directly. For some requests, searchers believe it is best to enter textwords without checking the thesaurus. They select this option when: they decide to enter the terms while they are online and have no time to examine the thesaurus [PR1]; the search is of the "quick-and-dirty" variety, or they are "just fishing" [PR2]; or the term is used to eliminate

irrelevant citations [PR3]. For example, a searcher may eliminate all the citations that include the term *Ph.D.* in their titles or abstracts from the set about students' behavior (using the AND NOT operator), because the user is not interested in examinations leading to this degree.

The unavailability of thesauri and their unsatisfactory quality, as well as the number of database to be searched, also lead searchers to enter a textword without looking for descriptors. Searchers would do so if: they do not trust the thesaurus and the indexing in a database [PD1]; they have decided to search a number of databases for one request, a decision they may make before or during the actual online session [PD2]; the thesaurus is not available to them [PD3]; or they think that they are familiar with the thesaurus and are convinced that it would not contain an adequate descriptor [PD4]. When they decide to change databases during a terminal session, searchers may enter a search statement that was constructed for the first database, including both descriptors and textwords, to be searched in the second database without checking its thesaurus [PD2].

Some searchers follow general guidelines which favor searching with textwords only. They may prefer to use terms that have been suggested by the user because they believe that the use of these terms results in more relevant citations [PS1]. Or, they believe that textwords are better for recall [PS2].

Use of Textwords to Probe Indexing. Searchers enter textwords to probe indexing because the thesaurus is not available to them [QD1], because when a concept is not completely clear to them they may not be sure which descriptor to use [QR1], or because they generally prefer to start with textwords and only then check for descriptors [QS1].

Entering as a Descriptor a Term That Might be a Descriptor. Searchers may enter a term as a descriptor, when they add the term to the query formulation during the online session and they feel time is too precious to check the thesaurus [RR1]. They may resort to this option also when they perform a multidatabase search [RR2].

Lastly, if terms are descriptors in another database [RD1], or if the thesaurus is not available [RD2], searchers may enter descriptors without checking the thesaurus, as they would do if they "knew" that it was a descriptor or thought it should be [RD3].

Discussion

The selection routine demonstrates that online searchers indeed use rules to support their selection of search keys. While an individual searcher usually applies these rules in an intuitive way rather than following a prescribed set of directions, these rules can be presented in a formal model. Further, this presentation indicates that intermediary expert systems for the selection of search keys could be constructed, and that the

case study method is a useful means to acquire the knowledge that is necessary to build such expert systems.

Intermediary Expert Systems

The selection routine is pertinent to the construction of intermediary expert systems that advise users in the selection of search keys. In particular, it can address request and user characteristics that are beyond the topical description of a search.

Although various techniques have been used to develop user models (Daniels, 1986), it is not clear what user characteristics are important for the success of an information-retrieval encounter. For example, can the age, profession, or geographic location of a user help an intermediary expert system decide on a search strategy? Paice emphasizes the significance of user and request characteristics when he observes that unlike other expert systems, in intermediary systems user interaction plays a central role, and the main concern is, therefore, *what* questions to ask and *when* to ask them (Paice, 1986). The selection routine actually uncovers these questions and suggests a sequence for their display.

The routine shows, for instance, that while online searchers do not take the age of a user into consideration, they may use their knowledge about the user to determine whether he or she prefers high precision or high recall, or whether the user wants the citations to include the exact terms used in the request. It also points out that for some terms or databases the selection of search keys is limited to one option, regardless of request characteristics, while for other terms or databases those characteristics play an essential role in decisions about search-key selection.

For example, if an expert system maps a common term to a descriptor for a request with two concepts, the descriptor should be used, and there is no need at this point to question the user. In contrast, if a single-meaning term is mapped to a descriptor broader in meaning, the system should inquire about recall and precision requirements, display records of some articles indexed with the broader descriptor and ask the user to assess their relevance, or inquire whether the concept that is represented by a broader descriptor is central to the request or is it used only for fine-tuning. Based on user's responses, the system can decide whether to use textwords, to AND textwords with descriptors, or whether to enter the descriptor alone.

The routine by itself is not sufficient for an expert system to perform like a searcher; additional knowledge in the form of semantic dictionaries or networks, thesauri, and rules are required to accomplish this task. Semantic dictionaries should be constructed to resolve a variety of linguistic and terminological issues, such as how to determine whether a term is common or single-meaning. Current thesauri should be developed to include an expanded lead-in vocabulary: terms that are

not descriptors but are listed in the thesaurus and lead to descriptors through "see" references. In addition to references from synonyms and quasi-synonyms to descriptors, such lead-in vocabulary should include mechanisms to make information embedded in phrases explicit. For example, a thesaurus should lead a user from a concept such as *attitudes towards themselves* to terms or descriptors such as *self-image* or *self-esteem*.

Rules should be added to the selection routine to refine its decision-making capabilities. For example, when a single-meaning term is mapped to a broader descriptor and high precision is required, two options are available: use textwords, or use textwords in combination with descriptors (using the AND operator). The second option provides higher precision than the first one. A rule is needed then to help determine which option to select. Such a rule may state, for instance, that the user be questioned about the level of precision required, or it may state that for requests with three or more concepts the first option be selected, but for requests with two concepts the second option be selected.

Such semantic dictionaries, thesauri, and rules do not exist as yet, and more research is needed to construct them. Once developed, though, the selection routine could be integrated into the knowledge base of a system to create a powerful intermediary expert system.

The Case Study Method

The applicability of the case study method to the extraction of knowledge from multiple experts is evidenced by the successful generation of a formal model that describes the selection of search keys. The use of this method in this study led to two conclusions: (1) the method of controlled comparison can successfully resolve conflicting evidence; and (2) observation and analysis of a relatively small number of searchers is sufficient to create a model that describes their searching behavior in formal terms.

The method of controlled comparison is used to explain observations that are seemingly contradictory. For example, according to the selection routine, searchers have two options when a single-meaning term is mapped to a descriptor through partial match: they can enter the descriptor, or they can use textwords for an inclusive search. These two options are polar opposites. The reasons provided by searchers to explain their choice, however, uncovered additional factors that played a role: Concern for recall may encourage searchers to use a textword key in an inclusive mode, if possible, or it may direct searchers to enter the descriptor if it is only ORed with other search keys, or if it was spotted in the indexing of a relevant citation. Thus, request requirements and indexing were discovered to be factors that affect the selection of search keys.

The original selection routine was based on the observation of the searching behavior of eight searchers. The

observation of the present study's 39 searchers resulted in only one major modification. Though only two additional options were discovered, a new condition was added to the original selection routine: a searcher does not know if a term is mapped to a descriptor. This condition was not spotted in the observations for the original routine because that study was limited to medical librarians who never searched a request without consulting a thesaurus. This new condition was uncovered immediately by observing the first searcher who was selected from a nonmedical subject area.

The experience derived from using the case study method shows, therefore, that limiting the sample of searchers to be observed by factors such as subject area or environment prevents the creation of a general model of searching behavior. On the other hand, if one takes into account the variety that exists among searchers, the observation of a relatively small number of searchers is sufficient for the creation of formal models that describe their searching behavior.

Acknowledgments

The study team on this project consisted of Nancy Phelps, Michael Crandall, Cynthia Altick Cunningham, and Kathleen McCrory. Their contribution to the study was substantial and highly appreciated. The online searchers who participated in this study were extremely cooperative spirits. Their collaboration made the study possible. My special thanks to Dagobert Soergel and Elaine Svenonius who reviewed earlier versions of this report. Their insightful comments and suggestions provided an inestimable contribution.

References

- Croft, W. B. (1987). Approaches to intelligent information retrieval. *Information Processing & Management*, 23, 249-254.
- Daniels, P. J. (1986). Cognitive models in information retrieval—an evaluative review. *Journal of Documentation*, 42, 272-304.
- Diesing, P. (1971). *Patterns of discovery in the social sciences*. Chicago, IL: Aldine-Atherton.
- Fidel, R. (1984). The case study method: a case study. *Library and Information Science Research*, 6, 273-283.
- Fidel, R. (1986). Towards expert system for the selection of search keys. *Journal of the American Society for Information Science*, 37, 37-44.
- Fidel, R. (1988). *Extracting knowledge for intermediary expert systems: the selection of search keys*. Syracuse, NY: ERIC. (ED 314 059)
- Kirby, M., & Miller, N. (1986). MEDLINE searching in Colleague: reasons for failure or success of untrained end users. *Medical Reference Services Quarterly*, 5, 17-34.
- Paice, C. (1986). Expert systems for information retrieval? *Aslib Proceedings*, 38, 343-353.
- Saracevic, T., Kantor, P., Chamis, A.Y., & Trivison, D. (1988). A study of information seeking and retrieving. I. Background and methodology. *Journal of the American Society for Information Science*, 39, 161-176.
- Sewel, W., & Teitelbaum, S. (1986). Observations of end-user online searching behavior over eleven years. *Journal of the American Society for Information Science*, 37, 234-245.

Smith, P. J., Shute, S. J., Chignell, M. H., & Krawczak, D. (1987). *The role of the human factors engineer in designing the interface to a knowledge based system*. Columbus, OH: Department of Industrial and Systems Engineering, The Ohio State University. (Report ISE-174).

Vickery, A., Brooks, H., & Robinson, B. (1987). A reference and referral system using expert system techniques. *Journal of Documentation*, 43, 1-23.