

THE POSSIBLE EFFECT OF ABSTRACTING GUIDELINES ON RETRIEVAL PERFORMANCE OF FREE-TEXT SEARCHING†

RAYA FIDEL

Graduate School of Library and Information Science, University of Washington, FM-30,
Seattle, Washington 98195

(Received 13 December 1985; in revised form 19 February 1986)

Abstract—A survey of abstracting guidelines used by producers of bibliographic databases examined abstracting instructions which aim to enhance free-text retrieval. Editors consider content of abstracts and their language as a primary factor in retrieval enhancement. Requirements about the type of abstracts, i.e. informative or indicative, and about their length are not affected by the capability of free-text retrieval. While instructions about content and language of abstracts indicate that better control and coordination results in better retrieval performance, it is not clear as yet in what way type and length of abstracts are affecting retrieval.

INTRODUCTION

Abstracts are important facilitators of free-text searching in most online bibliographic databases. Ever since abstracting and indexing services provided bibliographic databases in a printed form, however, the *functions* assigned to indexing have been very different from those assigned to abstracting. Indexing facilitated efficient retrieval of information, and abstracting facilitated judgments about the relevance of the retrieved information. Abstracts have even been viewed as substitutes for original documents.

Researchers in the design and evaluation of information systems have been most interested in the performance of free-text searching. Indexing is a costly, labor-intensive process, and index languages are expensive to construct. Further, several theoretical issues relating to their construction remain unresolved. The capability to retrieve documents without the aid of an artificially constructed index language seems, therefore, extremely attractive. Beginning with the Cranfield studies[1], numerous experiments have tried to determine if free-text searching indeed outperforms searching with the aid of a controlled index language. Results are contradictory, and the debate continues as to whether or not future systems should dispose of controlled vocabularies and be restricted to the free-text searching capability, e.g. [2,3,4, and 5].

The availability of bibliographic databases online changed the role of abstracts. The abstracts of most publicly available bibliographic databases can be searched in the free-text mode which allows users to search online for the occurrence of any terms they think appropriate. The system then retrieves citations to all documents whose abstracts contain the requested terms or combinations of them. This process takes place without the use of a specific index language. In other words, abstracts have become an important *enhancement* for information retrieval.

Despite this new role, vital questions concerning the construction of abstracts to facilitate efficient retrieval are still unanswered. Active research on abstracting ceased in the mid-1960s, after the importance of abstracts for judging document relevance and for use as substitutes was established by a series of tests[6], and a comprehensive survey of 130 abstracting policies produced criteria to be used in the evaluation of abstracts[7]. While these studies are milestones in abstracting research and their results are still highly valuable, they do not examine the appropriateness of abstracts for free-text retrieval.

†This study was supported by the Graduate School of the University of Washington.

Among practitioners, designers of information systems are providing abstracts for free-text searching, and system users heavily employ this capability. Several abstracting and indexing services have investigated the use of free-text searching to determine how to design their index languages and abstracting policies, e.g. [8 and 9]. In addition, online searchers attempt to identify efficient strategies for free-text searching of abstracts[10].

In view of the extensive experience in free-text searching gained in recent years, the time seems ripe for a systematic investigation of characteristics of abstracts that are essential for successful retrieval.

To analyze the criteria vital for acceptability of abstracts for free-text searching, I surveyed abstracting policies used by database producers to determine the degree to which existing abstracting and indexing services already incorporate such criteria into their abstracting. The analysis of data revealed characteristics of abstracts that are important to free-text retrieval; it also suggests some hypotheses about their effects on retrieval performance that other researchers may want to test.

ABSTRACTING GUIDELINES

A description of the survey and a detailed description of instructions pertinent for free-text retrieval is given elsewhere[11]. Figure 1 is a summary of these instructions as collected from abstracting guidelines of 57 databases.

THE CONTENT OF ABSTRACTS

General Statements

- * Use "important" concepts and terms (e.g., those which will enhance free-text retrieval, those for which a document gives enough information, or key words).

Index terms

- * Coordinate concepts used in abstracts with assigned descriptors.
 - (a) Assign concepts in abstract that are identical to descriptors.
 - (b) Assign concepts in abstracts that complement descriptors (e.g., relevant terms that are missed in descriptor indexing and in titles, terms that are more specific than descriptors, or a particular type of terms that is important to the subject area, such as geographic names).
 - (c) Assign concepts in abstracts that both complement and are identical to descriptors.
- * Enhance indexing independent of any index language used.

Check Lists

- * Follow a list of retrieval-related elements that should be included in abstracts.
- * Forms of check lists:
 - (a) Categories that should be included in abstracts (e.g., materials, properties and processes) and the conditions under which they should be included (e.g., only when they are discussed elaborately, or whenever mentioned).
 - (b) Specific and particular guidelines (e.g., "whenever dealing with a new product, mention the company name").

THE LANGUAGE OF ABSTRACTS

Use of Author Language

- * Use author language.
- * Do not use author language.
 - (a) Use standardized and concrete terms specific to a subject area.
- * Use both author language and synonyms.

Relationship to Index Language Used

- * Coordinate terms in abstracts with descriptors.
- * Complement descriptors with terms in abstract (e.g., use synonyms or more specific terms).
- * Use specific and well accepted terms for particular categories (such as materials, processes, and products).

Practices to Avoid

- * Do not use the negative (e.g., use sick instead of not healthy)
- * Do not list terms which have a common last word as a series (such as upper, middle, and working class).

Word Forms

- * Follow local language practices (e.g., change American spelling for British databases).
- * Always spell out terms in certain categories (such as processes, materials, products).
- * when a term and a descriptor are the same, record the term in the form used by the descriptor.
- * Express terms both in their abbreviated form and in their complete form.

Fig. 1. A summary presentation of abstracting guidelines for free-text searching.

Two findings of the survey should be mentioned before abstracting instructions are summarized.

First, of the 123 databases examined which provide abstracts for free-text searching and which have some sort of abstracting guidelines, only 57 (46%) include free-text instructions. Moreover, among the 103 databases that use descriptor indexing, only 40 (39%) provide instructions for free-text retrieval. It appears, therefore, that a considerable number of editors do not see the need for free-text instructions, and others may have realized this need but have not developed such instructions yet.

Second, the survey found that there is a disagreement among editors of abstracting policies in selecting the best way to enhance retrieval; as can be seen in Fig. 1, some of the instructions collected are contradictory. It is clear that only an analysis of the effect of each instruction on retrieval performance can determine which instructions are the "right" ones.

In their study of general abstracting guidelines, Borko and Chatman addressed three issues: the function, the content, and the form of abstracts[7]. Interestingly, only the latter two are considered by database producers in their specific instructions for computerized databases.

The first issue, the function of abstracts, determines if abstracts should be informative or indicative: whether they should condense the information in documents or report on what documents are about. While editors of computerized databases usually instruct abstractors about which type of abstract to compose, none of them presume that it relates to free-text retrieval. In other words, the nature of abstracts, whether informative or indicative, is still determined only by the purpose and use of abstracts and not by any factors that relate to information retrieval.

When general abstracting policies address the second issue, the content of abstracts, they indicate the *information* that should be included in abstracts, such as study method or research results. Editors who are interested in retrieval with abstracts additionally designate which *concepts* should be included in abstracts. While editors generally agree about the nature of concepts to include, they present a variety of opinions, some contradictory, about the specific criteria to use for concept selection.

Concepts to be included in abstracts for retrieval enhancement may be selected in coordination with assigned descriptors or may be added independently of any other indexing. When editors elect to coordinate concepts in abstracts with descriptor assignment, they can give abstractors clear instructions about whether these concepts should complement descriptors or duplicate them. If both trends are desired, specific guidance should be given as to when each approach should be applied. Regardless of whether the selection of concepts for abstracts is coordinated with indexing, editors may find it useful to supply abstractors with a checklist of categories or specific elements to be included in abstracts. It would also be advantageous to designate the conditions under which each category should be included.

The third issue, the form of abstracts, relates primarily to writing style and length of abstracts. Surprisingly, in all the abstracting manuals reviewed, no mention is made of possible relationships between the length of abstracts and efficiency of retrieval. Writing style and the language of abstracts, however, are related to retrieval: a number of editors issue policies about the way in which concepts should be represented to enhance retrieval. Unfortunately, some instructions make it difficult to maintain an acceptable writing style, with brevity and clarity of writing most affected.

The language of abstracts may or may not closely follow the language of authors. If it is not tied to the original document, editors must decide whether or not to coordinate terms in abstracts with index language. When coordination is preferred, editors should decide whether terms selected for abstracts should duplicate or complement index terms. For some databases, editors may select categories for which a particular form is desirable. While certain language usages can be ruled out, editors do not agree about whether consistency or variety in word forms should be required.

The central consideration for machine-searchable abstracts is whether their nature, content and form indeed affect retrieval. Currently, we have almost no evidence of such a relationship, but our experience in online searching and the knowledge we gained through retrieval tests indicate that abstracting policies probably affect retrieval performance. This

experience can help us to develop hypotheses about the nature of this effect, and, more importantly, guide research by pointing out which issues are significant and which are marginal to retrieval performance. Let us then examine each topic in abstracting policy and consider its effect on retrieval.

TYPES OF ABSTRACTS

Informative abstracts present condensed versions of original documents, while indicative abstracts report what documents are about. For instance, a sentence such as: "health care facilities in urban areas are more modern, more cost-effective, but less user-oriented than those in rural areas" in an informative abstract, would probably correspond in an indicative abstract to: "health care facilities in urban and rural areas are compared." If one considers terms and phrases in abstracts as access points, one can see that sentences in informative abstracts provide more access points than sentences in indicative ones. Using terminology from research in indexing, one could say that informative abstracts provide higher exhaustivity in indexing because they represent a larger number of aspects than indicative abstracts.

Experience in retrieval tests shows that high exhaustivity results in high recall but low precision[12]. In other words, informative abstracts will result in higher recall than indicative abstracts which will result in higher precision. This notion is plausible as demonstrated by the following example. Suppose users are looking for information about cost effectiveness of health care facilities. Our sample document will be retrieved with an informative abstract but might be rejected based on an indicative one. On the users' end, those who desire high recall for this request, who would like to see any document on the topic, would regret missing this document; thus, an informative abstract suits them best. Users who require high precision, who are interested only in documents that thoroughly discuss the cost effectiveness of health care facilities, on the other hand, would prefer the document to be rejected; for them, an indicative abstract is a better choice.

In addition, examination of the difference between the *nature* of access points provided by these two types of abstracts shows that indicative abstracts furnish access points that can be best described as major descriptors: only descriptors that represent the major topics of a document are included. In contrast, informative abstracts provide indexing that includes *both* major and minor descriptors. Although not rigorously tested, it is plausible that searching only on major descriptors increases precision while searching on all descriptors increases recall. This reinforces our previous assumption that informative abstracts are best for high recall requirements, while indicative ones satisfy high precision requests.

Looking at the specificity of terms included in abstracts, on the other hand, one finds that access points provided by indicative abstracts are frequently more general than their equivalent access points in informative abstracts. The notion among information scientists, however, is that specificity increases precision and may increase recall[12], depending on the terms used. In addition, it is also accepted that indexing with general terms may increase recall when searchers attempt to increase recall by entering broader terms. This leads us to the conclusion that informative abstracts, because they include highly specific terms, result in better precision, and under certain circumstances indicative ones may support high recall searches.

The straightforward way to resolve this contradiction would be to test a collection in which each document is represented by both types of abstract and compare retrieval performance for various requests. Results, however, may depend also on the nature of test requests. Soergel, for example, shows that exhaustivity increases recall for some requests but decreases recall and precision for others[13]. It would be important, therefore, to examine a stratified sample of requests. No typology for such stratification is available as yet, but one may safely hypothesize that informative abstracts result in better precision and better recall for requests that are specific and/or well-defined, while indicative abstracts provide better precision for general and/or ill-defined questions.

Resolving this issue is in the interest of both editors and searchers. When choosing a type of abstract, editors can rely on information beyond the nature of the material

abstracted and the expected use of abstracts. Based on the expected nature of requests most frequently searched on their database, they can determine which type of abstract is better for retrieval. Searchers too can use such findings to support database selection for individual requests. Experienced searchers who are familiar with a number of databases in a particular field sometimes "know" which database-specific abstracts are best for a specific strategy. Explicit relationships between types of abstracts and retrieval performance would help them base such decisions on empirical evidence.

LENGTH OF ABSTRACTS

As is the case with the nature of abstracts, editors of abstracting policies do not seem to assume a relationship between the length of abstracts and retrieval performance. In other words, no particular length is indicated as appropriate for computerized text retrieval.

The question of length is very similar to the question of abstract type. Clearly, an abstract of 300 words includes more access points, and probably more concepts, than an abstract of 100 words written to represent the same document. We do not know, however, whether concepts in long abstracts are also more specific than those included in short ones. Thus, it is not clear whether long abstracts provide for high or low recall or whether short ones enhance or impair precision.

As in the previous example, editors, researchers, as well as online searchers could benefit from revealing the relationship between length of abstracts and retrieval performance. In particular, editors who tailor the length of abstracts to fit resources or allocated space may have the opportunity to systematically examine the trade-offs between investment in organizing material for retrieval and in retrieval performance.

In conclusion, it seems that the type of an abstract and its length indeed affect retrieval performance. At present, however, we do not know what form this effect takes. It is not surprising, therefore, that editors of abstracting policies do not use criteria that relate to free-text retrieval in their decision about types and length of abstracts. Studies should be conducted to spell out this effect so editors improve service efficiency.

THE CONTENT OF ABSTRACTS

For abstracts to be valuable for online retrieval, they must include concepts that are central to documents. Thus, the general requirement that such concepts be included is merely an instruction to index documents with terms in abstracts. The effect on retrieval performance is determined by the manner in which this natural language indexing is related to controlled vocabulary indexing and on the specific policy that is used.

Unlike type and length of abstracts, which have as yet only an obscure connection with online retrieval, the effect of content and language of abstracts on retrieval is direct and observable. As a result, we are in a better position to recommend approaches to content and language that can improve retrieval performance.

The first people to experience the effect of content and language of abstracts are online searchers. When both controlled vocabulary terms and natural language terms are available as access points, they have greater control over retrieval performance than when only one mode is provided because they are free to choose a retrieval mode that will best suit the request at hand. Moreover, experienced online searchers are well-versed in the trade-offs between these two modes of searching and use them extensively to improve search results[14]. Therefore, whenever one analyzes the effect of the content of abstracts on retrieval performance, one has to consider moves made by searchers in online searching.†

It is also not surprising that content and language of abstracts are the only issues that editors consider when they formulate policies for free-text retrieval. The experience of online searchers is accessible to them (in fact, most of these editors are probably experienced in

†It is important to remember again that for an abstracting policy to optimize retrieval it has to be known to searchers. Editors, therefore, should make their policies known to online searchers. When they know what terms are selected for abstracts, searchers will be able to avoid a trial-and-error period in which they try to detect abstracting policies in order to increase their efficiency and productivity.

online searching), and the effect of these issues on retrieval can be taken into consideration. Most interestingly, our cumulative experience in information retrieval in general, and in online searching in particular, is almost sufficient at present to determine how to select the content and language of abstracts to enhance retrieval performance as illustrated in the following discussion.

As shown in the summary of abstracting guidelines (Fig. 1), concept selection may or may not be coordinated with descriptor assignment. When concept selection is not related to controlled vocabulary indexing, it is difficult to predict the effect of the content of abstracts on retrieval performance. All we can say is that if abstracting instructions are for specific terms, and/or for an exhaustive document representation, precision and recall will be affected as with specific and exhaustive indexing *if* searchers use free-text retrieval. But, as long as we do not know the criteria for specificity and exhaustivity of descriptor assignment, we cannot determine whether searching abstracts can indeed improve retrieval.

Therefore, to allow searchers maximum latitude in improving search results, the selection of concepts to be included in abstracts *should* be coordinated with descriptor indexing. Moreover, it is clear that concepts included in abstracts should *complement*, rather than duplicate, concepts selected for descriptor indexing. To include in abstracts, say, pertinent concepts that are not in the title or descriptor fields, language equivalents of subject codes, or concepts that are more specific than the descriptor enhanced, can improve *both* precision and recall.

To show the advantage of complementing descriptors let us consider, for instance, a search on a request about mini-pigs in a database that has the descriptor *Swine*, but not *mini-pig*. If precision is at stake, searchers can combine, using the AND operator, the descriptor *Swine* with the free-text term *mini-pigs*. This would retrieve all documents that are actually *about* some kind of a pig because they were indexed under the descriptor, and that specifically deal with mini-pigs. If recall is most important, searchers can group together all documents indexed under the descriptor *Swine* with all documents in which the term *mini-pigs* occurs in their abstracts, whether or not they are indexed under the descriptor. In both cases retrieval performance is improved because it is assumed that the term used in abstracts complements the descriptor.

To further improve performance, a checklist of categories to be included in abstracts may be used (Fig. 1). When searchers are assured that, say, materials and processes, or company names are always included in abstracts, they can use this information to increase both precision and recall. The interesting point here is that the number of categories included in abstracts will not affect retrieval one way or the other as long as searchers know which categories are involved. This seems to contradict commonly held ideas among information scientists: the larger the number of categories that must be included in abstracts the more exhaustive the assignment of index terms, and thus, recall will be higher but precision will be very low. This idea does not hold, however, when searchers know the specific categories because then they can formulate strategies that will compensate for these deficiencies.

THE LANGUAGE OF ABSTRACTS

The issue of what kind of language to select for abstracts presents some unresolved questions but can also serve as an example of the principle that the more planning and caution are invested in abstract writing, the better the retrieval will be. The main concern here is that terms used in abstracts correspond as closely as possible to terms used by users when searching for a particular concept. Because we know very little about which terms are used by users under what circumstances, most ideas expressed in abstracting policies of how to best select promising terms (e.g. searchers use authors' language) are merely suggestions for further research.

One way to overcome our inability to predict which terms will be entered by users is to represent concepts with more than one term. Thus, the use of synonymous terms to represent concepts certainly improves retrieval. Unfortunately, this option cannot be recommended because it impairs writing quality. Readers of abstracts may find it confusing when a concept, which is usually also a significant one, is represented by a number of terms within

one paragraph. Therefore, the main source of disagreement between editors of abstracting policies is whether they should be consistent and find a way to select the "best" term, or should introduce diversity to term selection, so that a relevant document will be retrieved even if different terms are used in searching.

In addition, those who support consistency in term selection suggest several methods for this purpose: use of author language, use of thesaurus terms whenever possible, or use of specialized authority lists or sets of rules for the representation of concepts from well identified categories such as company names or of process types.

The problem here is two-fold. First, it is not clear that consistency is indeed a desirable attribute for efficient retrieval. While there is some evidence that consistency in descriptor indexing improves retrieval[15], there is no evidence that consistent use of terms contributes to better retrieval. In fact, the counterargument, that diversity may improve retrieval, is equally plausible. This assumption, therefore, needs to be tested. Second, even if we find that consistency is important for retrieval, we have no evidence that either of the methods suggested will indeed select the "best" terms.

We do not know whether author language should be retained because we do not know whether users usually use author language. One study found that the vocabulary suggested by experts in a particular area for a thesaurus construction project was essentially the same as the vocabulary collected from the literature[16], but we do not know whether these experts will use the same terminology when they search the literature. It seems, however, that this issue is not of paramount importance. We will be better off after we study terms used by users in their requests for information. But whether the same terms are used also by authors is of no consequence for human abstracting, although it might be of significance to automated abstracting.

The other suggestion, that abstract terms duplicate thesaurus terms whenever possible, brings us back to the issue of complementarity vs. duplication between abstract and thesaurus terms. As explained previously, the main contribution of abstracts to the enhancement of retrieval is their ability to complement descriptor indexing.

This leaves us with the only plausible method at this time to achieve consistency: the requirement that selected categories be represented in standardized terms. This approach is particularly useful for categories that are not represented in a controlled vocabulary such as company or geographic names. Long experience in information retrieval has proved that authority lists or rules for a standardized expression of these terms improve retrieval.

Another consistency issue is the selection of word forms. Here again, there is no evidence that a particular form is a preferred one, or whether consistency is desired at all. There are some obvious instances, such as the use of British spelling for British users, but the prevailing assumption that word forms should be consistent must still be tested.

Lastly, all language usages that editors suggest to avoid are indeed impediments for retrieval. Writing abstracts without arcane language, negative phrases, or series of terms where the common one is only written once, definitely improve retrieval. Observing these suggestions, however, requires additional attention from abstractors.

To summarize, even though the issue of whether or not consistency improves retrieval is unresolved, it is clear that control over concepts included in abstracts and over their language can only help in retrieval. The more systematic the control, the easier it is for searchers to improve search results, so long as they are informed about the manner in which concepts and terms are selected.

SUMMARY

While research on abstracting has not addressed issues relating to the acceptability of abstracts for free-text retrieval, retrieval tests and our experience in online searching can shed some light on the issue and reveal criteria that are central to free-text retrieval.

Editors of abstracting policies issue special instructions regarding the content of abstracts and their language. These instructions are primarily based on our experience with online searching, where abstracts are actually used for free-text retrieval. The effect of both the content and the language of abstracts on retrieval performance is relatively direct and

observable. Instructions for these two issues can be detailed and specific. Experience has led most editors to require that concepts included in abstracts complement descriptor indexing, and that the language used in abstracts be controlled.

The effect of other issues in abstracting policies on retrieval performance cannot be directly derived from our experience. Although it is plausible to assume that the type of an abstract and its length indeed affect retrieval performance, we do not know what form this effect takes. It is not surprising, therefore, that abstracting guidelines pertinent to free-text retrieval do not address these issues. Type and length of abstracts, however, are among the first, and certainly most central, decisions in any abstracting policy. It is important, therefore, to uncover their impact on retrieval. When testing this impact, we will benefit most if we identify a match between types of requests on the one hand, and types of abstracts and their length on the other.

Acknowledgements—I wish to thank Kathy Waldera for her assistance with data collection and contacts with database producers, and Oren Sreebny for analyzing abstracting guidelines.

REFERENCES

1. Cleverdon, C. W. *Report on the Testing and Analysis of an Investigation into the Comparative Efficiency of Indexing Systems*. Cranfield, England: College of Aeronautics; 1962.
2. Cleverdon, C. W. Optimizing Convenient Online Access to Bibliographic Databases. *Information Services & Use*, 4(1/2):37-47; 1984.
3. Lancaster, F. W. Trends in Subject Indexing from 1957 to 2000. In (Taylor, P. J., ed.) *New Trends in Documentation and Information*. London: Aslib; 1980.
4. Henzler, R. G. Free or Controlled Vocabularies: Some Statistical User-Oriented Evaluations of Biomedical Information Systems. *International Classification*, 5(1):21-26; 1978.
5. Dubois, C. P. R. The Use of Thesauri in Online Retrieval. *Journal of Information Science*, 8(2):63-66; 1984.
6. Resnick, A. Relative Effectiveness of Document Titles and Abstracts for Determining Relevance of Documents. *Science*, 134(3484):1004-1005; 1961.
7. Borko, H., Chatman, S. Criteria for Acceptable Abstracts: A Survey of Abstracters' Instructions. *American Documentation*, 14(2):149-160; 1963.
8. Carrow, D., Nugent, J. Comparison of Free-Text and Index Search Abilities in an Operating Information System. *Proceedings of the American Society for Information Science Annual Meeting*, 14(2):2-e8; 1977.
9. Markey, K., Atherton, P., Newton, C. An Analysis of Controlled Vocabulary and Free Text Search Statements in On-Line Searches. *Online Review*, 4(3):225-236; 1980.
10. Wagers, R. Effective Searching in Database Abstracts. *Online*, 7(5):60-77; 1983.
11. Fidel, R. Writing Abstracts for Free-Text Retrieval. *Journal of Documentation*, 42(1):11-21; 1986.
12. Williamson, N. J. Subject Access in the On-Line Environment. *Advances in Librarianship*, 13:49-97; 1984.
13. Soergel, D. *Organizing Information: Principles of Data Base and Retrieval Systems*. Orlando, FL: Academic Press; 1985.
14. Fidel, R. Online Searching Styles: A Case-Study-Based Model of Searching Behavior. *Journal of the American Society for Information Science*, 35(4):211-221; 1984.
15. Leonard, L. E. *Inter-Indexer Consistency and Retrieval Effectiveness: Measurement of Relationships*. Ph. D. Dissertation. University of Illinois at Urbana-Champaign; 1975.
16. Kim, C., Kim, S. D. Consensus vs. Frequency: An Empirical Investigation of the Theories for Identifying Descriptors in Designing Retrieval Thesauri. *Information Processing and Management*, 13(4):253-258; 1977.