

Abstracts and Abstracting in Knowledge Discovery

MARIA PINTO AND F.W. LANCASTER

María Pinto, Departamento de Biblioteconomía y Documentación, Universidad de Granada,
18071 Granada, Spain.

F.W. Lancaster, Graduate School of Library and Information Science, University of Illinois

ABSTRACT

VARIOUS LEVELS OF CRITERIA FOR JUDGING the quality of abstracts and abstracting are presented. Requiriments for abstracts to be read by humans are compared with requeriments for those to be searched by computer. It is concluded that the wide availability of complete text in electronic form does not reduce the value of abstracts for information retrieval activities even in such more sophisticated applications as Knowledge discovery.

INTRODUCTION

Abstracts were first to be read by humans, providing concise summaries or descriptions of published items suitable for inclusion in printed indexing services or in scholarly journals along with the articles to which they relate. When computers sarterd to have a impact on information retrieval in the 1960s, abstracts became important as human-readable output form electronic databases. Later, as storage and processing cost declined, they began to assume a new role- that of computer-searchable surrogates for larger bodies of text.

Today, of course, it is economically feasible to store vast of text in computer-searchable form. Nevertheless, this has not made abstracts redundant. They remain useful summaries to be read by humans. Furthermore, if recall and precision are purposes because the searching of full text will frequently cause an unacceptable level of irrelevancy. Several investigations (e.g. Tenopir,1985) have shown that searching abstracts may be more effective or more cost-effective than searching of full text, while Salton (1971) found that, while full text gave better overall results than abstracts in the type of automatic proccessing employed in his SMART retrieval systems, the differences were not great and the abstracts allowed more cost-effective processing.

On the surface, one might assume that Knowledge discovery operations would be most likely to succeed when the complete text of items is processed. This is not necessarily so because full text can generate so many spurious relationships that significant and useful associations will be virtually impossible to recognize. Abstracts may still have great value in knowledge discovery activities as they do in many others.

This article will review various criteria by which the quality of abstracts may be judged. It will then discuss which criteria apply most clearly to the value of abstracts in Knowledge discovery applications.

QUALITY IN GENERAL

The world quality occurs frequently in every day life and, in this general setting stand for an idea that, while not necessarily exact, seems readily understood. On the other hand, in the more formal and restricted application—such as science, technology, commerce, and education—much less agreement exist on what “quality” really means and how the quality on something is to be measured and expressed. This is less true, of course, when applied to things that are all concrete. The quality of many manufactured products can be precisely quantified. This results from the fact that they must conform to standards that are strictly enforceable and are precisely quantifiable—e.g., steel either meets a standard relating to its composition or it does not. In the manufacturing situation, then, “quality control” is not a nebulous idea— it relates to the extent to which products meet the required standards. Despite it being an imprecise idea in many contexts, it is obvious that the last decade or so has brought a great increase in concern for “quality” in virtually all areas of human endeavor. The growth of the literature on the subject is a tangible manifestation of this objective. Nevertheless, it is somewhat misleading to speak of quality as though it were a single idea. Instead, one may recognize various levels or perspectives, as illustrated in Table 1. At the one extreme, there is the abstract or transcendental idea of quality, one that is static, absolute, and existing only in philosophical and metaphysical speculation. At the other extreme is the “user” perspective, which is personal and even, perhaps, idiosyncratic. It is also dynamic and “relative”—in the sense that it often involves a comparison and the choice of one among several alternatives. Frequently the choice will be made on the basis of cost, which could be a cost in monetary form or in terms of time and convenience.

Table 1.

VARIOUS POSSIBLE LEVELS OR PERSPECTIVES RELATING TO QUALITY

<i>Perspective</i>	<i>Basis of judgment</i>	<i>Characteristics</i>
ABSTRACT	Philosophy Static	Absolute Speculation
ORGANIZATIONAL		
PROCESS	Standards, regulations, norms	Some processes may be strictly regulated; others not
PRODUCT	Standards	For manufactured products,

		may be objective and enforceable
SERVICE	Standards or norms	More subjective than objective; rarely enforceable
USER/CUSTOMER	Cost	Dynamic
	Value	Relative
	Personal value system	

Between these extremes, we have other levels or perspectives, identified in the table as being "organizational." Quality related to products varies greatly with type of product. For the many products that must be manufactured to conform to standards, quality can be considered close to absolute, at least relative to the standards, but not completely so since most manufacturing standards accept a range of values, albeit a very narrow one in many cases. Intellectual products, such as various forms of publication, are less susceptible to true standardization. At least, this is true of their content. The container (paper, binding, and so on) can be standardized.

The process perspective is heterogeneous. Some processes can be standardized. In fact, in some cases, processes may be subjected to absolute regulation—e.g., concerning cleanliness, safety, and other health-related issues. Again, intellectual processes are not as susceptible to regulation or standardization.

The service perspective falls midway between the product perspective and the user perspective. Services can rarely be judged in absolute terms. Although some aspects of service can be quantified—e.g., number of seats per reader, number of students per instructor—the standards are rarely completely enforceable so they tend to be normative values rather than true standards, and some services (e.g., associated with organized religion or with certain social agencies) seem not susceptible to evaluation against any type of standard.

Nevertheless, approaches to the enforcement of quality within service agencies have become increasingly sophisticated in the last several years, culminating in adoption of the principles of total quality management (TQM), which include emphasis on customer satisfaction and on continuous improvement.

QUALITY IN INFORMATION SERVICES SETTINGS

Since information tends to be intangible, it is quite difficult to obtain agreement on appropriate measures of quality for most elements of information service. All of the various perspectives represented in Table 1, except for the purely philosophical, can apply in the information service environment. Quality can be considered in tangible terms for many aspects

of information products but can be quite elusive elsewhere, especially in both the service perspective and the user perspective. Take, for example, the case of an electronic database. Quantifiable measures of quality can be applied when the database is considered as a *product*—i.e., its coverage of the literature within its scope, the average number of access points per item, up-to-dateness, and so on. Retrospective search and current awareness services derived from use of the database present more difficult problems. While certain measures of service quality can be objective and quantified (e.g., average time elapsing from demand to delivery of response), the more important measures, such as those of recall and precision, are both subjective and difficult to apply. When the user perspective is considered here, of course, the situation becomes even more subjective. For example, a database search can retrieve many items that match a user's stated request or stored interest profile but may still be judged of little value by the user, because the actual information needed did not appear in the search results, because the items retrieved were already known to the user, because he considered them as insignificant contributions to the subject, or for some other reason that might be quite idiosyncratic. Moreover, if the user has to pay for the service, he may apply a purely cost-effectiveness measure to judge the quality of the search results—i.e., the cost per useful item retrieved.

The process perspective on quality is not as nebulous as the user perspective, but it is still an area in which it is difficult to apply true standards. This is because many of the processes are intellectual. While certain applications can be standardized (e.g., form of name in catalog entries), others, such as subject indexing, are not susceptible to standardization except in very trivial aspects. Quality concerns applied to another intellectual process, abstracting, is the focus of our present discussion.

QUALITY CONSIDERATIONS APPLIED TO ABSTRACTING

From a psycholinguistic perspective, abstracting is more ambitious and complex than indexing: not only must the text of documents be analyzed in some detail but text (the abstract) must also be produced. This text must be coherent syntactically and semantically and, at the same time, be a reasonable summary of the original document. Abstracting is the most difficult of all operations normally applied in a document processing environment because, today at least, an abstract must act as both content description and retrieval tool. Fidel (1986) has shown that these two uses may not be completely compatible.

A possible model of the abstracting process is presented in Figure 1. In actual fact, four levels of processing are represented. The goals are defined by the service or journal producing the abstracts and may be embodied or reflected in guidelines for the abstractors. The individual abstractor observes the goals by following these guidelines. The two processes, "content interpretation/selection" and "content transformation," are directly equivalent to the conceptual analysis and translation stages of subject indexing (Lancaster, 1998). The former

is concerned with understanding what is discussed in the original text and deciding which elements should be included in the abstract, while the latter is concerned with the composition of the abstract—i.e., how the selected elements are to be presented in the text of the abstract.

The process headed "checking" is the process directly related to quality. It has several possible dimensions: the individual abstractor may impose his/her own review of quality before submitting the abstract for further processing, the abstractor's work may later be checked by an editor or senior abstractor before publication, and readers may apply their own quality checks relating to the intelligibility of the abstract and its value in predicting the relevance of the original item to their own interests.

Figure 1 suggests that the quality of the abstract is largely determined by the quality of the knowledge base of the abstractor. The knowledge base incorporates both linguistic knowledge (ability to interpret the language of texts in the subject area dealt with) and nonlinguistic knowledge: understanding of the subject matter, of the needs and interests of the audience served, and of the guidelines under which the abstractor is to operate.

Despite the fact that their application in retrieval (as substitutes for or complements to sets of index terms) makes them more important now than ever before, especially in the Internet environment (Wheatley & Armstrong, 1997), there exist no generally accepted measures of the quality of abstracts. Of course, many writers have identified their desirable attributes. Borko and Bernier (1975), for example, regard abstracting as a form of writing that has a unique style (it is not a "natural" form); abstracts must be brief, accurate, and clearly written. Unlike Cremmins

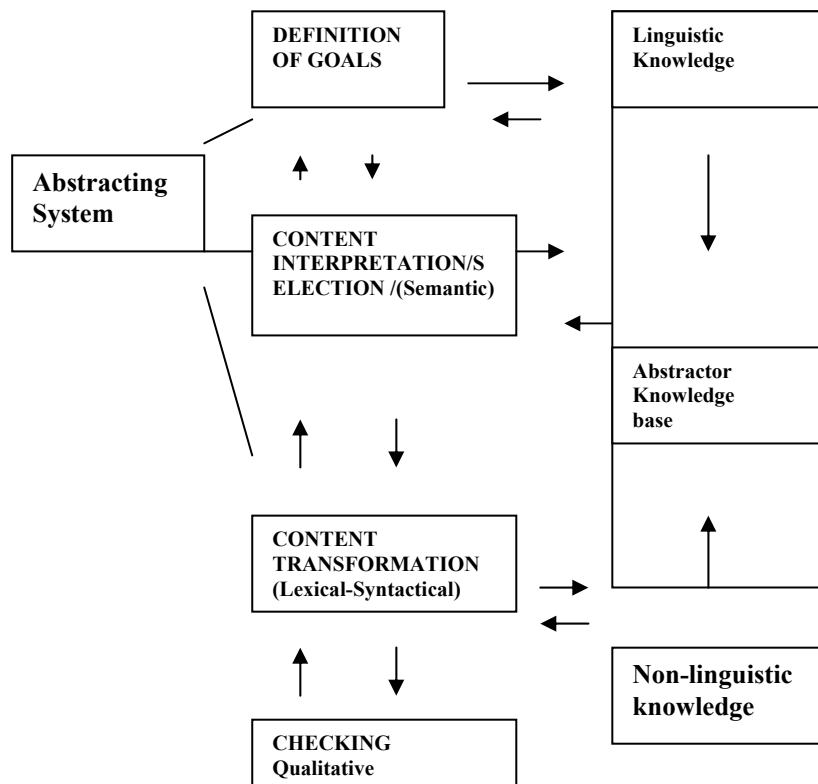




Figure 1. Integrated Model of the Abstracting Process. _____

(1996), they do not claim that they must have "elegance." Lancaster (1998) suggests two broad criteria for judging quality, are the major points of the article covered and are they represented accurately, succinctly, and unambiguously? The latest English-language standard (National Information Standards Organization, 1997), while it gives guidance on style, makes no attempt to provide criteria that can be used to assess quality. Other writers (e.g., Brown & Day, 1983) have focused on the art of text summarization or on the skills needed by a good abstractor (e.g., see Endres-Niggemeyer, Maier, & Sigel, 1995).

Interest in the evaluation of abstracts can be traced back to at least the late 1950s. For example, Edmundson et al. (1959) proposed several criteria: comparison with an "ideal" abstract, the retrievability of a document by the abstract, and the extent to which the abstract could be used to answer test questions as well as the use of intuitive subjective judgment. Payne, Munger, and Altman (1962) also suggested a test of the value of abstracts in answering questions, as well as a measure of the amount of text reduction achieved in an abstract, and the use of a consistency test in which the similarity of different abstracts, prepared from the same document, is compared. Vinsonhaler (1966) recommended use of a seven-point scale to determine the similarity between an abstract and the document it relates to; also proposed was a more conventional approach, one of predictive validity—the extent to which abstracts are able to correctly predict the relevance of documents.

Mathis (1972) offered a numerical value, known as the "data coefficient" (DC), for the evaluation, expressed by a formula that incorporates a data retention factor and a length retention factor. The value of the DC is increased by reducing the number of words in the abstract, by increasing the number of concepts ("data elements") represented, or both.

Several of these approaches have been applied over the years. The most favored is a test of the ability of an abstract to predict the relevance of a document to a particular information need. Investigators who have applied this to abstracts, or to extracts derived by computer, include Rath, Resnick, and Savage (1961); Resnick (1961); Kent et al. (1967); Dym (1967); Shirey and Kurfeerst (1967); Saracevic (1969); Marcus, Benenfeld and Kugel (1971); Thompson (1973); and Keen (1976).

Hartley, Sydes, and Blurton (1996) provide an example of a study in which abstracts are judged on their ability to answer various questions; in this case, they were comparing "structured" abstracts with unstructured ones. Salton et al. (1997) used a variation of the

similarity approach: the extent to which an automatically-derived extract resembles one derived by humans.

Other approaches have assessed the "readability" of abstracts using standard readability formulas, comprehension measures, or both. Examples can be found in the work of Dronberger and Kowitz (1975), King (1976), Tenopir and Jacso (1993), and Hartley (1994). More recently, Wheatley and Armstrong (1997) studied readability of a variety of abstracts drawn from Internet sources.

A more "linguistic" approach was used by Salager-Meyer (1991), who analyzed a sample of medical abstracts from this perspective, finding almost half to be "poorly structured" (i.e., having discursal deficiency). Since "discursal deficiency" can include such things as conceptual scatter (e.g., results reported in different places in the abstract), as well as omission of an important element (e.g., purpose of research) from the abstract, the author implies that abstracts flawed in this way will be less effective in conveying information. Elsewhere, Pinto (1992, 1994, 1995) has dealt in detail with the process of text summarization from the viewpoint of linguistic structure.

It is clear that the various quality criteria proposed or used in the past look at abstracts/abstracting from different perspectives. In fact, virtually all perspectives represented in Table 1 can apply to abstracts or abstracting, as shown in Table 2.

The process perspective deals primarily with attributes of cognitive representation. Here analogies can be drawn between the process of abstracting and the process of indexing (Lancaster, 1998). The exhaustivity of the abstract relates to its breadth of coverage. In essence, it is a measure of the extent to which all of the themes of the original text are represented in the abstract. Clearly, an abstract is unlikely to include all the content of the original text (unless it is completely trivial) so the exhaustivity of the abstract can be considered as the extent to which all of

Table 2.

ATTRIBUTES OF QUALITY ASSOCIATED WITH DIFFERENT PERSPECTIVES

ON ABSTRACT AND ABSTRACTING

<u>Process perspective</u>	<u>Service Perspective</u>
Exhaustivity	Customer satisfaction Accuracy
Cost-effectiveness	
Readability	
Cohesion/coherence	<u>User perspective</u>
Cost	Cost
	Value
<u>Product perspective</u>	<u>Process/product perspective</u>

Consistency

Density

Brevity

Cost

Cost

the themes (ideas, conclusions, or whatever) judged important are covered in the abstract. This implies that some group of people, presumably specialists in the subject area dealt with, can agree on what is important in the original and what is not.

In an ideal situation, of course, an abstract should be tailored to the needs of a particular audience. This is most obvious in the case of one written for an in-house bulletin prepared, for example, to serve a particular company or research organization. In this case, an exhaustive abstract would be one that covers all the themes of the original that are of potential interest to the limited community. In an extreme case, this might be a single theme—e.g., results of applying a particular drug extracted from a medical article discussing multiple approaches to the treatment of some disease. Clearly, the writer of such an abstract must have a good knowledge of the needs and interests of the target community as well as familiarity with the subject matter dealt with. The more heterogeneous the interests of the audience served, the less likely one is to reach agreement on which themes to include in the abstract and which not: difficult in the case of general mission-oriented abstracts (e.g., serving the needs of an entire industry), more difficult still in the case of abstracts intended to serve the needs of an entire discipline.

Accuracy refers to the extent to which the abstract correctly represents the original text. A theme covered in the abstract could be an inaccurate representation of the original because of an intellectual error (the abstractor misinterprets the text) or an error of carelessness (the abstractor records incorrectly—e.g., gives a wrong numerical value). The former should be relatively rare but could occur if the abstractor is not fully familiar with the subject matter or if the original text is somewhat obscure. A special case would be the situation of an abstractor dealing with a language in which he is not completely fluent. Accuracy errors of the second type would be attributable to personal characteristics of the abstractor (ability to concentrate, ability to transcribe correctly), including qualities that could vary considerably from one day to the next, and to working conditions. Most significant of the latter would be pressures associated with required productivity, where an abstractor may be required to produce a specified number of abstracts in a particular time period. Of course, once the abstract has been printed and distributed, it would be impossible to determine whether an error of this type was attributable to the abstractor or was introduced at some later stage of the production process.

The readability of an abstract is determined by the ability of the abstractor to express himself clearly, concisely, and unambiguously, by the rules or guidelines under which he operates, and by the format of the abstract (e.g., some claim that abstracts structured into

paragraphs with topical headings are easier to comprehend). To the extent that general tests of the readability of text (e.g., the Flesch Reading Ease formula) or of comprehension (e.g., cloze criteria) are applicable to abstracts, readability can be an objective measure and one that can be quantified.

Cohesion/coherence is related to readability but is not identical with it. These properties relate to connectivity between different parts of a text. Extracts prepared by computer (selecting sentences on the basis of statistical, positional, or linguistic criteria) will frequently be lacking in these properties, even though the total extract may be a satisfactory representation of the principal themes of the original text. Salager-Meyer (1991) is perhaps the only author to apply such linguistic criteria to humanly prepared abstracts. A major measure used was that of conceptual scatter—the extent to which related elements (e.g., results) are separated in an abstract. Since structured abstracts (see Haynes, 1993; Hartley, 1994; Hartley, Sydes, & Blurton, 1996) are formatted into paragraphs with pre-established subheads (e.g., methods, results), they are less likely to exhibit such conceptual scatter. Factors affecting cohesion/coherence are the same as those affecting readability.

The product perspective (see Table 2) relates to the technical adequacy of the abstract. The idea of *consistency in abstracting* is similar to consistency in subject indexing. It refers to the degree to which two individuals produce abstracts that are similar to each other (inter abstractor consistency) or the degree to which the same individual agrees with himself when abstracting a document on different occasions (infra-abstractor consistency). In the indexing situation, a distinction can be made between consistency in conceptual analysis and consistency in the translation of the conceptual analysis into a particular vocabulary (e.g., terms drawn from a thesaurus). Consistency in abstracting, however, applies only at the conceptual level since it is unrealistic to expect different individuals to use exactly the same words or grammatical constructions. Presumably, consistency will be greatest when abstractors work to precise rules as to what to include and what not. For obvious reasons, structured abstracts should be more consistent than others.

In abstracting, just as in indexing, consistency is not the same as quality (Cooper, 1969). Nevertheless, if two abstractors (or indexers) consistently produce similar results, while a third agrees little with the other two, one is generally inclined to believe that the consistent abstracting (indexing) will be "better." Salton, Singhal, Mitra, and Buckley (1997) justify their automatic procedures for selecting and linking pieces of text on the grounds that the summary thus produced is as likely to agree with a humanly-produced summary as one humanly-produced summary is to agree with another. In translating from one language to another also, consistency (similarities) has been suggested as an indicator of quality (Brew & Thompson, 1994).

Brevity is an obviously desirable attribute of a good abstract, and it is susceptible to exact measurement. Moreover, length is one of the few attributes that the published standards can and do address precisely, at least in terms of a recommended range in number of words. Nevertheless, brevity should always be secondary to other considerations such as exhaustivity and accuracy. Moreover, absolute standards make little sense since several factors would influence the brevity: length, complexity or diversity of the original, type of abstract (indicative, informative, critical), and accessibility of the original (one could argue that materials less physically or intellectually accessible—e.g., published in obscure sources or unfamiliar languages—should be abstracted more fully).

Cost can be related to abstracts at different levels: the intellectual cost, of creating an abstract, the cost per abstract of producing a printed publication, the cost per abstract in distribution (e.g., as part of a current awareness service), and so on. Factors affecting cost differ from level to level. For example, abstract length has a major effect on the cost of producing a printed publication but much less effect on the inclusion of an abstract in an electronic database. Cost of writing the abstract in the first place depends most obviously on who the writer is, how much he/she is paid, and who is paying. The cost of abstracting can be looked at from several different perspectives. For example, use of author-generated abstracts is economical for database producers. From the much broader (society) perspective, however, they are very expensive since the time of such authors as research scientists can be considered to be so valuable that it is perhaps better spent on other things.

Carried to its logical conclusion, of course, one could argue that the greatest cost associated with abstracting is the cost of the time spent by people in reading the abstracts (thus the importance of such factors as brevity and readability) and in taking actions based upon them (thus the importance of such factors as accuracy and exhaustivity). Cost, then, is a multifaceted attribute when related to abstracts and abstracting. For this reason, it appears within all the perspectives illustrated in Table 2.

Density is a measure that relates the attribute of exhaustivity to that of brevity. It thus, in a sense, combines the process and product perspectives. Given that the abstract includes everything that should be included— all the topics of potential interest to the intended audience—the briefer the abstract the better providing, of course, that other requirements, such as readability, are not significantly degraded. Density, then, refers to the amount of information content provided by an abstract of a certain length. The density of an abstract can be considered related to its entropy—the extent to which uncertainty about the original document is reduced for the reader of the abstract. Standard tests of the relevance predictability of abstracts address this issue.

The data coefficient proposed and tested by Mathis (1972) was a precise measure of density, defined by the equation $DC = C/L$ —i.e., the data coefficient (DC) is the "data

retention factor," C, divided by the 'length retention factor," L. The C value is the measure of exhaustivity as defined earlier in this discussion, while the L value is the number of words in the abstract divided by the number in the original. Clearly, the DC of an abstract improves as either exhaustivity or brevity increase.

While the process and product perspectives consider abstracts as entities in their own right, the service perspective is obviously concerned with their application. Providers of abstracts, whether publishers and editors of scholarly journals or producers of secondary databases in printed or electronic form, are presumably concerned with offering a product that the majority of their customers (journal readers, database users) will find acceptable. Customer satisfaction will most obviously be associated with the process and product parameters discussed earlier, perhaps most closely to accuracy, readability, and exhaustivity. Clearly, the providers will also be concerned with production and distribution costs so, ultimately, "quality" becomes a matter of cost-effectiveness—i.e., customer satisfaction at least cost.

As mentioned earlier, the user perspective on quality will tend to be subjective, relative, dynamic and, perhaps, idiosyncratic. Users of abstracts will be likely to judge their quality in practical and pragmatic terms. They are unlikely to demand elegance but they will expect readability. Ultimately, they will judge abstracts and abstracting services in terms of costs and value to themselves. Taking the user's own time into account, the predictive validity of the abstract is of paramount importance. That is, users will be unhappy with a service whose abstracts frequently cause the incurring of costs associated with obtaining complete texts that turn out to be irrelevant. Nor will they be satisfied with one that frequently fails to lead them to sources that they would judge valuable if seen in full form.

CURRENT METHODS

The automatic processing of text has increased considerably over the years as computing power has increased, computing and storage costs have decreased, and more and more text has become available in electronic form, largely as a byproduct of various forms of publishing. The development of the Internet and the World Wide Web, which makes vast quantities of text accessible to huge numbers of users, has made text search the norm rather than the exception. As might be expected from all of this, interest in automatic text processing methods has increased very greatly in the 1990s, in the research community as well as in government and commercial sectors. Current approaches to the processing of text, for information retrieval and related purposes, are well portrayed in the proceedings of a series of conferences. Most important among these have been the Text Retrieval Conferences (TREC) organized by the (U. S.) National Institute of Standards and Technology (Sparck Jones, 1995; Harman, 1997), the Message Understanding Conferences (MUC), the Conferences on Applied Natural Language Processing, and the International Conferences on Document Analysis and Recognition. The

TREC and MUC conferences are particularly important for their methodology: all participating research groups must apply their text processing procedures to some common pre-established tasks, allowing performance comparisons across the methods.

Current methods of text processing for information-retrieval-like purposes go beyond text search, automatic indexing and automatic extracting procedures (all of which have existed, to some extent at least, since the late 1950s), now including such activities as text linkage, text augmentation, and text generation. Nevertheless, while current approaches may achieve rather better results, they do not differ much in principle from those first introduced forty to fifty years ago, even though they may be given different names ("text summarization" in place of abstracting/extracting, "text categorization" in place of indexing/classification, and so on) and may be more sophisticated in some respects (e.g., not just extracting text but putting the extracts into a pre-established template). While some current approaches claim to apply techniques drawn from artificial intelligence research, and the term "intelligent text processing" is sometimes used to refer to procedures of this type (see, for example, Jacobs, 1992), it is doubtful that any can be considered to exhibit true intelligence (Lancaster & Smith, 1999).

KNOWLEDGE DISCOVERY

The great majority of the criteria of quality proposed and used in the past apply most obviously to abstracts intended to be read by humans. As mentioned earlier, if abstracts are intended primarily as useful document surrogates for search purposes, the quality criteria become somewhat different. Unfortunately, a good abstract for search purposes is unlikely to be good for a human reader. Indeed, an abstract prepared solely for computer searching; such as the telegraphic abstracts of the semantic code system (Perry & Kent, 1958), may not be readable by humans at all, and abstracts prepared primarily for search purposes, such as the mini-abstracts proposed by Lunin (1967), may be somewhat difficult for humans to comprehend.

For retrieval purposes, and especially in knowledge discovery tasks, exhaustivity and accuracy are extremely important, and the other attributes in Table 2 diminish in significance. In fact, for abstracts intended solely for search purposes, such criteria as readability and coherence/cohesion are not important at all, while other attributes are applicable in opposite ways. Most obviously, brevity is not necessarily desirable since the retrievability of an abstract will be directly related to its length (i.e., number of access points provided). Nevertheless, for reasons mentioned before, there is likely to be an optimum length for effective search and discovery operations. The data retention factor proposed by Mathis (1972) seems a particularly appropriate criterion in knowledge discovery applications since it relates length to completeness of content coverage. Also undesirable for knowledge discovery purposes is internal consistency because redundancy improves retrievability. That is, if a particular idea is expressed in different ways in an abstract (no synonym control), this

increases the probability that the text will match an expression selected by a particular searcher or that meaningful relationships between related ideas will be revealed.

CONCLUSION

Text surrogates for larger bodies of text, whether one refers to them as "abstracts," "summaries," or some other term, have proved extremely useful in a wide variety of information processing applications for very many years. The increasing application of computers to text processing has not reduced their value (although criteria for judging their quality may have changed somewhat), and one has no reason to suppose that their value diminishes as more critical or sophisticated operations, including those of knowledge discovery, are applied to the text.

REFERENCES

- Borko, H., & Bernier, C. L. (1975). *Abstracting concepts and methods*. New York: Academic Press.
- Brew, C., & Thompson, H. S. (1994). Automatic evaluation of computer generated text: A progress report on the Text Eval project. In *Proceedings of the Human language Technology Workshop (March 8-U, 1994)* (pp.108-113). San Francisco, CA: Morgan Kaufmann.
- Brown, A. L., & Day, J. D. (1983). Macrorules for summarizing texts: The development of expertise. *Journal of Verbal Learning and Verbal Behavior*, 22(1), 1-14.
- Cooper, W. S. (1969). Is inter-indexer consistency a hobgoblin? *American Documentation*, 20(3), 268-278.
- Cremmins, E. T. (1996). *The art of abstracting*. 2d ed. Arlington, VA: Information Resources Press.
- Dronberger, G. B., & Kowitz, G. T. (1975). Abstract readability as a factor in information systems. *Journal of the American Society for Information Science*, 26(2), 108-111.
- Dym, E. D. (1967). Relevance predictability: I. Investigation, background and procedures. In A. Kent, O. E. Taulbee, J. Belzer, & G. D. Goldstein (Eds.), *Electronic handling of information: Testing and evaluation* (pp. 175-185). Washington, DC: Thompson Book Co.
- Edmundson, H. P.; Oswald, V. A., Jr.; & Wyllys, R. E. (1959). *Automatic indexing and abstracting of the contents of documents*. Los Angeles, CA: Planning Research Corporation.
- Endres-Niggemeyer, B.; Maier, E.; & Sigel, A. (1995). How to implement a naturalistic model of abstracting: Four core working steps of an expert abstractor. *Information Processing of Management*, 37(5), 631-674.
- Fidel, R. (1986). Writing abstracts for free-text searching. *Journal of the American Society for Information Science*, 42(1), 11-21.
- Harman, D. (1997). The TREC conferences. In K. Sparck Jones & P. Willett (Eds.), *Readings in information retrieval* (pp. 247-256). San Francisco, CA: Morgan Kaufmann.
- Hartley, J. (1994). Three ways to improve the clarity of journal abstracts. *British Journal of Educational Psychology*, 64(1), 331-343.

Hartley, J., & Sydes, M. (1996). Which layout do you prefer? An analysis of readers' preferences for different typographic layouts of structured abstracts. *Journal of Information Science*, 22(1), 27-37.

Hartley, J.; Sydes, M.; & Blurton, A. (1996). Obtaining information accurately and quickly: Are structured abstracts more efficient? *Journal of Information Science*, 22(5), 349-356.

Haynes, R. B. (1993). More informative abstracts: Current status and evaluation. *Journal of Clinical Epidemiology*, 46, 595-597.

Jacobs, P. S. (Ed.). (1992). *Text-based intelligent systems: Current research and practice in information extraction and retrieval* Hillsdale, NJ: Lawrence Erlbaum.

Keen, E. M. (1976). A retrieval comparison of six published indexes in the field of library and information science. *Unesco Bulletin for Libraries*, 30(1), 26-36.

Kent, A.; Belzer, J.; Kurfeerst, M.; Dym, E. D.; Shirey, D. L.; & Bose, A. (1967). Relevance predictability in information retrieval systems. *Methods of Information in Medicine*, 6(2), 45-51.

King, R. (1976). A comparison of the readability of abstracts with their source document*. *Journal of the American Society for Information Science*, 27(2), 118-121.

Lancaster, F. W. (1998). *Indexing and abstracting in theory and practice*, 2d ed. Urbana-Champaign: University of Illinois, Graduate School of Library and Information Science.

Lancaster, F. W., & Smith, L. C. (In press). *Intelligent technologies in library and information service applications: A realistic appraisal* Medford, NJ: Information Today.

Lunin, L. (1967). The development of a machine-searchable index-abstract and its application to biomedical literature. In B. Flood (Ed.), *Three Drexel information science-research studies* (pp. 47-134). Philadelphia, PA: Drexel Press.

Marcus, R. S.; Benenfeld, A.R.; & Kugel, P. (1971). The user interface for the Intrex retrieval system. In D. E. Walker (Ed.), *Interactive bibliographic search: The user/computer interface* (pp. 159-201). Montvale, NJ: AFIPS Press.

Mathis, B. A. (1972). *Techniques for the evaluation and improvement of computer-produced abstracts*. Columbus: Ohio State University, Computer and Information Science Research Center (OSU-CISRC-TR-72-15. PB 214 675).

National Information Standards Organization. (1997). *Guidelines for abstracts*. Bethesda, MD: NISO.

Payne, D.; Munger, S.J.; & Altman, J. W. (1962). *A textual abstracting technique: A preliminary development and evaluation support*. Pittsburgh, PA: American Institutes for Research (2vols. AD 285081-285082).

Perry, J. W., & Kent, A. (1958). *Tools for machine literature searching*. New York: Interscience Publishers Inc.

Pinto, M. (1995). Documentary abstracting: Toward a methodological model. *Journal of the American Society for Information Science*, 46(3), 225-234.

Pinto, M. (1994). Interdisciplinary approaches to the concept and practice of written text documentary content analysis (WTDCA). *Journal of Documentation*, 50(2), 111-133.

Pinto, M. (1992). *El resumen documental: Principios y métodos*. Madrid: Fundación Germán Sánchez Ruipérez.

Rath, G.J.; Resnick, A.; & Savage, T. R. (1961). Comparison of four types of lexical indicators of content. *American Documentation*, 12(2), 126-130.

Resnick, A. (1961). Relative effectiveness of document titles and abstracts for determining relevance of documents. *Science*, 134(3484), 1004-1006.

Salager-Meyer, F. (1991). Medical English abstracts: How well are they structured? *Journal of the American Society for Information Science*, 42(7), 528-531.

Salton, G. (Ed.). (1971). *The SMART retrieval system: Experiments in automatic document processing*. Englewood Cliffs, NJ: Prentice-Hall.

Salton, G.; Singhal, A.; Mitra, M.; & Buckley, C. (1997). Automatic text structuring and summarization. *Information Processing (Sⁱ Management*, 33(2), 193-207.

Saracevic, T. (1969). Comparative effects of titles, abstracts and full texts on relevance judgements. *Proceedings of the American Society for Information Science*, 6, 293-299.

Shirey, D. L., & Kurfeerst, M. (1967). Relevance predictability: II. Data reduction. In A. Kent; Taulbee, O. E.; Belzer, J.; Goldstein, G. D. (Eds.), *Electronic handling of information: Testing and evaluation* (pp. 187-198). Washington, DC: Thompson Book Co.

Sparck Jones, K. (1995). Reflections on TREC. *Information Processing Ssf Management*, 31(3), 291-314.

Tenopir, C. (1985). Full text database retrieval performance. *Online Review*, 9(2), 149-164.

Tenopir, C., & Jacs6, P. (1993). Quality of abstracts. *Online*, 17(3), 44-55.

Thompson, C. W. N. (1973). The functions of abstracts in the initial screening of technical documents by the user. *Journal of the American Society for Information Science*, 24(4), 270-276.

Vinsonhaler, J. F. (1966). Some behavioral indices of the validity of document abstracts. *Information Storage and Retrieval*, 3(1), 1-11.

Wheatley, A., & Armstrong, C.J. (1997). Metadata, recall, and abstracts: Can abstracts ever be reliable indicators of document value? *Aslib Proceedings*, 49(8), 206-213.

