**A Note About
Information Science
Research**

Gerard Salton

TR 84-633
September 1984

Department of Computer Science
Cornell University
Ithaca, New York 14853

A Note About Information Science Research

Gerard Salton^{*}

Abstract

This note deals with the relationship between information science research and practice. The impression that the field is moribund and that the research output is uniformly inferior is not supported by an examination of the information retrieval literature.

- - - - -

A provocative note by Carl Keren appeared in a recent issue of JASIS under the title "On Information Science". [1] The note raises interesting questions about the relationship between information science research and practice. Keren asks four questions (slightly paraphrased):

- a) do researchers in information science write about research which advances the state-of-the-art?
- b) how much of the research output has really contributed to our body of knowledge?
- c) is information science a name with a recognizable body attached to it? is it a subject whose contents we can define?
- d) is there a lack of feedback between researchers and technologists?

^{*} Department of Computer Science, Cornell University, Ithaca, New York 14853.

Questions such as these are of course not new, and they are reflective of legitimate concerns. In fact, I share some of Mr. Keren's misgivings about the quality and quantity of acceptable research work in our field. But we part company when Keren by implication answers all his questions in the negative, and specifically blames the researchers for what he considers the unsatisfactory state of information science work. In stating his arguments, Keren makes valid points about research based on questionable survey output, about lack of knowledge of the literature, about rediscovery of the obvious, and the like. Unfortunately, the note is also replete with remarks of doubtful relevance and with unsubstantiated claims about the inferior state of information science compared with that of other natural sciences. Taken as a whole, the note presents an anti-research, if not anti-intellectual image, and requires, I believe, some amendments.

Among the less edifying parts of Keren's note is a completely gratuitous quotation of a remark, purported to have been made by Y. Bar-Hillel, to the effect that information scientists who need mathematics to describe their concepts quite probably know little of either information science or mathematics. To which one replies first that the remark has little to do with Keren's concern about the quality or quantity of research in information science, and second that the comment is obviously foolish and should not in fairness to Bar-Hillel have been repeated. (Bar Hillel was an astute observer of our field; but he was partial to sweeping, and sometimes indefensible generalizations which might better be forgotten.) The truth of the matter is that a mathematical approach cannot transform poor information science into good research, and that few people are able to make convincing use of mathematics in the information field. But this does not of course imply that all mathematics is extraneous to information science, or that all information

scientists who use mathematics are inferior.

What about Mr. Keren's main questions? I will beg leave from answering the question about the scope and definition of information science. I just don't find that discussions relating to the definition of disciplines advance the state-of-the-art, and personally I don't have much to contribute. I will then confine myself to the other questions which concern the quality and quantity of research in information science, and the contributions of this research to our body of knowledge and to information science practice. In discussing these questions, Mr. Keren mentions that "I did not research this, nor do I have firm data, but on the strength of my own experience, as a practitioner, I have the impression that only very rarely (does) useful and non-trivial information emanate from the published research "

Well, I do not have any firm data either, but my own assessment generally agrees with Keren's: most of the published research in our field is probably not worth doing and ought to be forgotten. I differ with Keren in that I do not jump to the implied conclusion that therefore our researchers are all inferior and that information science research is useless. Nor do I have any evidence that relative to the number of researchers in information science, the proportion of first-rate work in our field is smaller than it is in other fields. The latter point is made explicitly at some length in Keren's letter.

As I said, I cannot give conclusive evidence either way, but I can cite the literature. What follows is a list of active research topics pursued by people who generally have something to say. All of these topics belong to that part of information science which may loosely be entitled "text analysis, storage, and retrieval", and the literature citations cover the three most recent years (1981-83). I emphasize that the list does not constitute a

ranking of the best people, or the best topics, or the best papers. It is simply a sampling of recent work that offers interesting perspectives to researchers and practitioners alike. For a different view of the information science field in areas not covered here, the reader may wish to turn to other recent surveys. [2,3]

1. Under the vector processing retrieval strategy, both queries and stored documents are represented by sets of terms without Boolean operators. The vector processing model has been criticized because the implicit assumption is that the terms are independent of each other: Bookstein [4] has shown, however, that under various stated assumptions the preferred probabilistic retrieval approach and the vector processing systems are closely related. The vector processing model, whose effectiveness is hard to surpass in practice, has also been studied more recently by Wong and Raghavan. [5]
2. Probabilistic retrieval models have been preferred by researchers because they have a proper theoretical base and allow, at least in theory, the use of arbitrary combinations of dependent terms. Two different, well-known, probabilistic models have recently been compared, and a unified model including elements of both models has been proposed by Robertson, Maron and Cooper. [6,7] An elegant, theoretically optimal probabilistic model based on the maximum entropy principle has been introduced by Cooper. [8,9] Croft has described and evaluated procedures for extending the basic probabilistic model by allowing the use of weighted terms for document representation. [10,11] Yu and co-workers have introduced a generalized probabilistic model in which dependent term pairs, term triples, and higher

order term combinations can be included in decreasing order of the importance of the various term combinations. [12] Finally, Wu and Salton have evaluated some practical methods for actually obtaining the needed probabilistic parameters for the case where term independence is assumed. [13]

3. The probabilistic retrieval approach is theoretically attractive, but until viable methods appear for actually obtaining the probabilistic parameters for dependent term combinations, the simpler vector processing approaches will be used in practice, or the established inverted file procedures based on term list manipulations. In this connection substantial attention has been paid recently to the acceleration and optimization of the similarity computations between queries and documents using inverted file strategies. Smeaton and van Rijsbergen [14] and Murtagh [15] have worked on the optimization of near-neighbor computations usable for query-document comparisons. Noreault and Chatham have applied similar methods to the computation of term-term similarities [16], and Willett has used optimal inverted file strategies to compute the document-document similarities needed for document clustering [17,18] All of these optimized inverted file procedures produce substantial improvements over conventional inverted file methodologies and are directly usable in current operational retrieval environments.
4. The well-known query reformulation process based on relevance feedback uses relevance assessments obtained from retrieval system users in response to earlier information searches. Attar and Frankel have continued working on local (and hence inexpensive) clustering and

relevance feedback procedures. [19] Term relationships obtained from a maximum spanning tree of term-term similarities are used in a relevance feedback environment by van Rijsbergen and co-workers. [20] All of these automatic query expansion and reformulation methods appear to be eminently suitable for practical implementation. [21]

5. Operational retrieval systems are currently based on Boolean query formulations. The conventional Boolean methodology is rudimentary: no term weighting, no output ranking of retrieved items, no automatic query formulation. A great deal of attention has been devoted by researchers to the refinement of Boolean query manipulations. Bookstein has compared different term weighting systems in Boolean query environments. [22] Radecki and Buell and Kraft have studied the fuzzy set retrieval models which permit the use of weighted document identifiers with Boolean queries. [23-26] A different extended Boolean system was introduced by Salton, Fox and Wu which is more general than the fuzzy set system and appears to produce large improvements in retrieval output. This system can be superimposed on the currently operating retrieval methods. Regrettably some mathematical considerations are needed to understand the operations of the extended system. [27]
6. Fancy front-end procedures have been studied by many people and beginnings have been made toward the design of expert systems that facilitate the access to the stored collections. [28-34] Marcus has designed a flexible system that uses a common query language to access a variety of different bibliographic retrieval systems. [29] The

interactive system by Doczkocs is based on flexible term weighting and automatic thesaurus display operations. [30] Frei and Jauslin use graphic terminals and multiple window displays to guide the user through the query formulation process using a hierarchical thesaurus display. [33] Finally, Guida and Tasso propose an expert-type interactive system in which user responses to appropriately phrased questions are utilized as a guide to the search process. [34]

7. Bibliographic citation networks have been used for many years to characterize various disciplines and literature types. Kwok suggests that terms from cited titles be used for document identification purposes [35], and White and Griffith continue the extensive experimentation with citation network manipulations [36,37].
8. The area of text passage retrieval, as opposed to full document retrieval has not yet come into its own, but it offers interesting prospects as a compromise between document retrieval and question-answering. O'Connor [39] has done pioneering work in this area, and Bernstein and Williamson have built an interesting system for retrieval of text passages from a specially constructed handbook on hepatitis diseases. [40]
9. The related area of automatic indexing using linguistic analysis is being seriously revived by various researchers. For example, Dillon uses phrase assignment techniques in various text processing experiments [41-42], and Sparck Jones and Tait describe a sophisticated system for the automatic generation of syntactically (and semantically) acceptable term phrases [43].

10. The interesting area of fast text character matching - for example between query and document texts - is treated by Haskin and Hollaar [44] who propose the construction of special hardware devices for rapid character string processing. Sophisticated simulation methods for the construction of representative bibliographic databases are described by Tague and Nelson. [45] Simulated collections simplify the operations of controlled retrieval evaluation experiments. Finally, a large group of people including, for example, Macleod and Crawford [46], are interested in the design of mixed search and retrieval systems that process formatted numeric data bases as well as text data in an integrated operation.

All of the foregoing research areas are directly relevant to information retrieval practice in that they point the way to the construction of flexible, interactive retrieval systems which should substantially simplify the collection search and retrieval processes for future generations of information seekers. Many other information processing topics might have been given prominence: for example, recent attempts at automatic abstracting [48], word fragment encoding experiments [49], the design of distributed retrieval systems [50], and the area of systems evaluation [51]. For present purposes, the foregoing list suffices to demonstrate that an impressive variety of research topics is being pursued by people who have something to offer.

In the face of the relative wealth of interesting work actually available, Mr. Keren's rhetorical questions about "What have you researchers lately done for us practitioners?" seems to miss the point. There are never any shortcuts in bridging the gap between research and application. That is true in our area as much as in all other intellectual areas of endeavor. It is

necessary to study the literature; it is necessary to have sufficient know-how to discriminate, and to put matters in context. Eventually the pieces will fit together, and the observer can judge the specifics instead of being forced to rely on superficial impressions and generalizations. In our field, as in every other, it is necessary to know the field before being able to contribute.

I see no reason, however, to resign. There has been substantial progress in information system design these past twenty years. No doubt, the changes will accelerate in the years to come. The field will evolve even more rapidly if the practitioners were to stop blaming the research side and asked instead "What have we practitioners lately done for ourselves?".

References

- [1] C. Keren, On Information Science, Journal of the ASIS, Vol. 35, No. 2, March 1984, p. 137.
- [2] M. Kochen, Information Science Research: The Search for the Nature of Information, Journal of the ASIS, Vol. 35, No. 3, May 1984, p. 194-199.
- [3] S. Herner, Brief History of Information Science, Journal of the ASIS, Vol. 35, No. 3, May 1984, p. 194-199.
- [4] A. Bookstein, Explanation and Generalization of Vector Models in Information Retrieval, in Research and Development in Information Retrieval, Lecture Notes in Computer Science, Vol. 146, G. Salton and H.J. Schneider, editors, Springer-Verlag, Berlin, 1983, p. 118-132.
- [5] S.K.M. Wong and V.V. Raghavan, Vector Space Model of Information Retrieval - A Reevaluation, in Research and Development in Information Retrieval, C. J. van Rijsbergen, editor, Cambridge University Press, 1984, p. 167-186.
- [6] S.E. Robertson, M.E. Maron and W.S. Cooper, Probability of Relevance: A Unification of Two Competing Models for Document Retrieval, Information Technology: Research and Development, Vol. 1, No. 1, January 1982, p. 1-22.
- [7] M.E. Maron, Associative Search Techniques Versus Probabilistic Retrieval Models, Journal of the ASIS, Vol. 33, No. 5, September 1982, p. 308-310.
- [8] W.S. Cooper and P. Huizinga, The Maximum Entropy Principle and its Application to the Design of Probabilistic Retrieval Systems, Information Technology: Research and Development, Vol. 1, No. 2, April 1982, p. 99-112.
- [9] W.S. Cooper, Exploiting the Maximum Entropy Principle to Increase Retrieval Effectiveness, Journal of the ASIS, Vol. 34, No. 1, January 1983, p. 31-39.
- [10] W.B. Croft, Document Representation in Probabilistic Models of Information Retrieval, Journal of the ASIS, Vol. 32, No. 6, November 1981, p. 451-457.
- [11] W.B. Croft, Experiments with Representation in a Document Retrieval System, Information Technology: Research and Development, Vol. 2, No. 1, January 1983, p. 1-22.
- [12] C.T. Yu, D. Buckley, K. Lam and G. Salton, A Generalized Term Dependence Model in Information Retrieval, Information Technology: Research and Development, Vol. 2, No. 4, October 1983, p. 129-154.

- [13] H. Wu and G. Salton, The Estimation of Term Relevance Weights Using Relevance Feedback, *Journal of Documentation*, Vol. 37, No. 4, December 1981, p. 194-214.
- [14] A.F. Smeaton and C.J. van Rijsbergen, The Nearest Neighbor Problem in Information Retrieval, *Proceedings of Fourth International Conference on Information Storage and Retrieval, SIGIR Forum*, Vol. 16, No. 1, Summer 1981, p. 83-87.
- [15] F. Murtagh, A Very Fast Exact Nearest Neighbor Algorithm for Use in Information Retrieval, *Information Technology: Research and Development*, Vol. 1, No. 4, October 1982, p. 275-283.
- [16] T. Noreault and R. Chatham, A Procedure for the Estimation of Term Similarity Coefficients, *Information Technology: Research and Development*, Vol. 1, No. 3, July 1982, p. 189-196.
- [17] P. Willett, Document Clustering Using an Inverted File Approach, *Journal of Information Science*, 2, 1980, p. 223-231.
- [18] P. Willett, A Fast Procedure for Calculation of Similarity Coefficients in Automatic Classification, *Information Processing and Management*, Vol. 17, No. 2, 1981, p. 53-60.
- [19] R. Attar and A.S. Frankel, Experiments in Local Metrical Feedback in Full-Text Retrieval Systems, *Information Processing and Managements*, Vol. 17, No. 3, 1981, p. 115-126.
- [20] C.J. van Rijsbergen, D.J. Harper, and M.F. Porter, The Selection of Good Search Terms, *Information Processing and Management*, Vol. 17, No. 2, 1981, p. 77-91.
- [21] P. Pietilainen, Local Feedback and Intelligent Automatic Query Expansion, *Information Processing and Management*, Vol. 19, No. 1, 1983, p. 51-58.
- [22] A. Bookstein, A Comparison of Two Systems of Weighted Boolean Retrieval, *Journal of the ASIS*, Vol. 32, No. 4, July 1981, p. 275-279.
- [23] D.A. Buell and D.H. Kraft, Threshold Values and Boolean Retrieval Values, *Information Processing and Management*, Vol. 17, No. 3, 1981, p. 127-136.
- [24] D.A. Buell and D.H. Kraft, A Model for a Weighted Boolean Retrieval System, *Journal of the ASIS*, Vol. 32, No. 3, May 1981, p. 211-216.
- [25] D.A. Buell, A General Model of Query Processing in Information Retrieval Systems, *Information Processing and Management*, Vol. 17, No. 5, 1981, p. 249-262.
- [26] T. Radecki, Reducing the Perils of Merging Boolean and Weighted Retrieval Systems, *Journal of Documentation*, Vol. 38, No. 3, September 1982, p. 207-211.

- [27] G. Salton, E.A. Fox and H. Wu, Extended Boolean Information Retrieval, Technical Report 82-511, Department of Computer Science, Cornell University, August 1982, also in Communications of the ACM, Vol. 26, No. 11, November 1983, p. 1022-1033.
- [28] C.T. Meadow, T.T. Hewett and E.S. Aversa, A Computer Intermediary for Interactive Data Base Searching, Journal of the ASIS, Vol. 33, No. 5, September 1982, p. 325-332, also Vol. 33, No. 6, November 1982, p. 357-364.
- [29] R.S. Marcus, An Experimental Comparison of the Effectiveness of Computers and Humans as Search Intermediaries, Journal of the ASIS, Vol. 34, No. 6, 1983, p. 381-404.
- [30] T.E. Doczkocs, From Research to Application: The Cite Natural Language Information Retrieval System, in Research and Development in Information Science, Lecture Notes in Computer Science, Vol. 146, G. Salton and H.J. Schneider, editors, Springer-Verlag, Berlin, 1983, p. 251-262.
- [31] M.H. Heine, A Simple Intelligent Front End for Information Retrieval Systems Using Boolean Logic, Information Technology: Research and Development, Vol. 1, No. 4, October 1982, p. 247-260.
- [32] A.S. Pollitt, End User Touch Searching for Cancer Therapy: A Rule Based Approach, Proceedings of Sixth Annual Conference on Research and Development in Information Retrieval, SIGIR Forum, Vol. 17, No. 4, Summer 1983, p. 136-145.
- [33] H.P. Frei and J.F. Jauslin, Graphical Presentation of Information and Services: A User-Oriented Interface, Information Technology: Research and Development, Vol. 2, No. 1, 1983, p. 23-42.
- [34] G. Guida and C. Tasso, IR-NLI: An Expert Natural Language Interface System, Conference on Applied Natural Language Processing, Association for Computational Linguistics, Santa Monica, California, 1983, p. 31-38.
- [35] K.C. Kwok, A Document-Document Similarity Measure Based on Cited Titles and Probability Theory and its Application to Relevance Feedback Retrieval, in Research and Development in Information Retrieval, C.J. van Rijsbergen, editor, Cambridge University Press, 1982.
- [36] H.D. White and B.C. Griffith, Authors as Markers of Intellectual Space: Cocitation in Studies of Science, Technology and Society, Journal of Documentation, Vol. 38, No. 4, December 1982, p. 255-272.
- [37] H.D. White and B.C. Griffith, Author Cocitation: A Literature Measure of Intellectual Structure, Journal of the ASIS, Vol. 32, No. 3, May 1981, p. 163-171.
- [38] E. Noma, Untangling Citation Networks, Information Processing and Management, Vol. 18, No. 2, 1982, p. 43-53.

- [39] J. O'Connor, Citing Statements: Computer Recognition and Use to Improve Retrieval, *Information Processing and Management*, Vol. 18, No. 3, 1982, p. 125-131.
- [40] L.M. Bernstein and R.E. Williamson, Testing of a Natural Language Retrieval System for a Full Text Knowledge Base, *Journal of the ASIS*, Vol. 35, No. 4, July 1984, p. 235-247.
- [41] M. Dillon and A.S. Gray, Fasit: A Fully Automatic Syntactically Based Indexing System, *Journal of the ASIS*, Vol. 34, No. 2, 1983, p. 99-108.
- [42] M. Dillon, Thesaurus-Based Automatic Book Indexing, *Information Processing and Management*, Vol. 18, No. 4, 1982, p. 167-178.
- [43] K. Sparck Jones and J.I. Tait, Automatic Search Term Variant Generation, *Journal of Documentation*, Vol. 40, No. 1, March 1984, p. 50-66.
- [44] R.L. Haskin and L.A. Hollaar, Operational Characteristics of a Hardware-Based Pattern Matches, *ACM Transactions on Database Systems*, Vol. 8, No. 1, March 1983, p. 15-40.
- [45] J. Tague and M. Nelson, Simulation of Bibliographic Databases Using Hyperterms, in *Lecture Notes in Computer Science*, Vol. 146, G. Salton and H.J. Schneider, editors, Springer-Verlag, Berlin, 1983, p. 194-208.
- [46] I.A. Macleod and R.G. Crawford, Document Retrieval as a Database Application, *Information Technology: Research and Development*, Vol. 2, 1983, p. 43-60.
- [47] H.J. Schek, Methods for the Administration of Textual Data in Database Systems, in *Information Retrieval Research*, R.N. Oddy, R.E. Robertson, C.J. van Rijsbergen and P.W. Williams, editors, Butterworths, London, 1981, p. 218-235.
- [48] C.D. Paice, The Automatic Generation of Literature Abstracts: An Approach Based on the Identification of Self-Indicating Phrases, in *Information Retrieval Research*, R.N. Oddy, R.E. Robertson, C.J. van Rijsbergen and P.W. Williams, editors, Butterworths, London, 1981, p. 172-191.
- [49] K. Kropp and G. Walch, A Graph Structured Text Field Index Based on Word Fragments, *Information Processing and Management*, Vol. 17, No. 6, 1981, p. 363-376.
- [50] C. Moulinoux, J.C. Faure, and W. Litwin, Messidor: A Distributed Information Retrieval System, in *Lecture Notes in Computer Science*, Vol. 146, G. Salton and H.J. Schneider, editors, Springer-Verlag, Berlin, 1983, p. 51-61.
- [51] P. Bollmann, Two Axioms for Evaluation Measures in Information Retrieval, in *Research and Development in Information Retrieval*, C.J. van Rijsbergen, editor, Cambridge University Press, 1984, p. 233-246.

**A Note About
Information Science
Research**

Gerard Salton

TR 84-633
September 1984

Department of Computer Science
Cornell University
Ithaca, New York 14853

A Note About Information Science Research

Gerard Salton*

Abstract

This note deals with the relationship between information science research and practice. The impression that the field is moribund and that the research output is uniformly inferior is not supported by an examination of the information retrieval literature.

- - - - -

A provocative note by Carl Keren appeared in a recent issue of JASIS under the title "On Information Science". [1] The note raises interesting questions about the relationship between information science research and practice. Keren asks four questions (slightly paraphrased):

- a) do researchers in information science write about research which advances the state-of-the-art?
- b) how much of the research output has really contributed to our body of knowledge?
- c) is information science a name with a recognizable body attached to it? is it a subject whose contents we can define?
- d) is there a lack of feedback between researchers and technologists?

* Department of Computer Science, Cornell University, Ithaca, New York 14853.

scientists who use mathematics are inferior.

What about Mr. Keren's main questions? I will beg leave from answering the question about the scope and definition of information science. I just don't find that discussions relating to the definition of disciplines advance the state-of-the-art, and personally I don't have much to contribute. I will then confine myself to the other questions which concern the quality and quantity of research in information science, and the contributions of this research to our body of knowledge and to information science practice. In discussing these questions, Mr. Keren mentions that "I did not research this, nor do I have firm data, but on the strength of my own experience, as a practitioner, I have the impression that only very rarely (does) useful and non-trivial information emanate from the published research "

Well, I do not have any firm data either, but my own assessment generally agrees with Keren's: most of the published research in our field is probably not worth doing and ought to be forgotten. I differ with Keren in that I do not jump to the implied conclusion that therefore our researchers are all inferior and that information science research is useless. Nor do I have any evidence that relative to the number of researchers in information science, the proportion of first-rate work in our field is smaller than it is in other fields. The latter point is made explicitly at some length in Keren's letter.

As I said, I cannot give conclusive evidence either way, but I can cite the literature. What follows is a list of active research topics pursued by people who generally have something to say. All of these topics belong to that part of information science which may loosely be entitled "text analysis, storage, and retrieval", and the literature citations cover the three most recent years (1981-83). I emphasize that the list does not constitute a

order term combinations can be included in decreasing order of the importance of the various term combinations. [12] Finally, Wu and Salton have evaluated some practical methods for actually obtaining the needed probabilistic parameters for the case where term independence is assumed. [13]

3. The probabilistic retrieval approach is theoretically attractive, but until viable methods appear for actually obtaining the probabilistic parameters for dependent term combinations, the simpler vector processing approaches will be used in practice, or the established inverted file procedures based on term list manipulations. In this connection substantial attention has been paid recently to the acceleration and optimization of the similarity computations between queries and documents using inverted file strategies. Smeaton and van Rijsbergen [14] and Murtagh [15] have worked on the optimization of near-neighbor computations usable for query-document comparisons. Noreault and Chatham have applied similar methods to the computation of term-term similarities [16], and Willett has used optimal inverted file strategies to compute the document-document similarities needed for document clustering [17,18] All of these optimized inverted file procedures produce substantial improvements over conventional inverted file methodologies and are directly usable in current operational retrieval environments.
4. The well-known query reformulation process based on relevance feedback uses relevance assessments obtained from retrieval system users in response to earlier information searches. Attar and Frankel have continued working on local (and hence inexpensive) clustering and

interactive system by Doczkocs is based on flexible term weighting and automatic thesaurus display operations. [30] Frei and Jauslin use graphic terminals and multiple window displays to guide the user through the query formulation process using a hierarchical thesaurus display. [33] Finally, Guida and Tasso propose an expert-type interactive system in which user responses to appropriately phrased questions are utilized as a guide to the search process. [34]

7. Bibliographic citation networks have been used for many years to characterize various disciplines and literature types. Kwok suggests that terms from cited titles be used for document identification purposes [35], and White and Griffith continue the extensive experimentation with citation network manipulations [36,37].
8. The area of text passage retrieval, as opposed to full document retrieval has not yet come into its own, but it offers interesting prospects as a compromise between document retrieval and question-answering. O'Connor [39] has done pioneering work in this area, and Bernstein and Williamson have built an interesting system for retrieval of text passages from a specially constructed handbook on hepatitis diseases. [40]
9. The related area of automatic indexing using linguistic analysis is being seriously revived by various researchers. For example, Dillon uses phrase assignment techniques in various text processing experiments [41-42], and Sparck Jones and Tait describe a sophisticated system for the automatic generation of syntactically (and semantically) acceptable term phrases [43].

necessary to study the literature; it is necessary to have sufficient know-how to discriminate, and to put matters in context. Eventually the pieces will fit together, and the observer can judge the specifics instead of being forced to rely on superficial impressions and generalizations. In our field, as in every other, it is necessary to know the field before being able to contribute.

I see no reason, however, to resign. There has been substantial progress in information system design these past twenty years. No doubt, the changes will accelerate in the years to come. The field will evolve even more rapidly if the practitioners were to stop blaming the research side and asked instead "What have we practitioners lately done for ourselves?".

- [13] H. Wu and G. Salton, The Estimation of Term Relevance Weights Using Relevance Feedback, *Journal of Documentation*, Vol. 37, No. 4, December 1981, p. 194-214.
- [14] A.F. Smeaton and C.J. van Rijsbergen, The Nearest Neighbor Problem in Information Retrieval, *Proceedings of Fourth International Conference on Information Storage and Retrieval, SIGIR Forum*, Vol. 16, No. 1, Summer 1981, p. 83-87.
- [15] F. Murtagh, A Very Fast Exact Nearest Neighbor Algorithm for Use in Information Retrieval, *Information Technology: Research and Development*, Vol. 1, No. 4, October 1982, p. 275-283.
- [16] T. Noreault and R. Chatham, A Procedure for the Estimation of Term Similarity Coefficients, *Information Technology: Research and Development*, Vol. 1, No. 3, July 1982, p. 189-196.
- [17] P. Willett, Document Clustering Using an Inverted File Approach, *Journal of Information Science*, 2, 1980, p. 223-231.
- [18] P. Willett, A Fast Procedure for Calculation of Similarity Coefficients in Automatic Classification, *Information Processing and Management*, Vol. 17, No. 2, 1981, p. 53-60.
- [19] R. Attar and A.S. Frankel, Experiments in Local Metrical Feedback in Full-Text Retrieval Systems, *Information Processing and Management*, Vol. 17, No. 3, 1981, p. 115-126.
- [20] C.J. van Rijsbergen, D.J. Harper, and M.F. Porter, The Selection of Good Search Terms, *Information Processing and Management*, Vol. 17, No. 2, 1981, p. 77-91.
- [21] P. Pietilainen, Local Feedback and Intelligent Automatic Query Expansion, *Information Processing and Management*, Vol. 19, No. 1, 1983, p. 51-58.
- [22] A. Bookstein, A Comparison of Two Systems of Weighted Boolean Retrieval, *Journal of the ASIS*, Vol. 32, No. 4, July 1981, p. 275-279.
- [23] D.A. Buell and D.H. Kraft, Threshold Values and Boolean Retrieval Values, *Information Processing and Management*, Vol. 17, No. 3, 1981, p. 127-136.
- [24] D.A. Buell and D.H. Kraft, A Model for a Weighted Boolean Retrieval System, *Journal of the ASIS*, Vol. 32, No. 3, May 1981, p. 211-216.
- [25] D.A. Buell, A General Model of Query Processing in Information Retrieval Systems, *Information Processing and Management*, Vol. 17, No. 5, 1981, p. 249-262.
- [26] T. Radecki, Reducing the Perils of Merging Boolean and Weighted Retrieval Systems, *Journal of Documentation*, Vol. 38, No. 3, September 1982, p. 207-211.

- [39] J. O'Connor, Citing Statements: Computer Recognition and Use to Improve Retrieval, *Information Processing and Management*, Vol. 18, No. 3, 1982, p. 125-131.
- [40] L.M. Bernstein and R.E. Williamson, Testing of a Natural Language Retrieval System for a Full Text Knowledge Base, *Journal of the ASIS*, Vol. 35, No. 4, July 1984, p. 235-247.
- [41] M. Dillon and A.S. Gray, Fasit: A Fully Automatic Syntactically Based Indexing System, *Journal of the ASIS*, Vol. 34, No. 2, 1983, p. 99-108.
- [42] M. Dillon, Thesaurus-Based Automatic Book Indexing, *Information Processing and Management*, Vol. 18, No. 4, 1982, p. 167-178.
- [43] K. Sparck Jones and J.I. Tait, Automatic Search Term Variant Generation, *Journal of Documentation*, Vol. 40, No. 1, March 1984, p. 50-66.
- [44] R.L. Haskin and L.A. Hollaar, Operational Characteristics of a Hardware-Based Pattern Matches, *ACM Transactions on Database Systems*, Vol. 8, No. 1, March 1983, p. 15-40.
- [45] J. Tague and M. Nelson, Simulation of Bibliographic Databases Using Hyperterms, in *Lecture Notes in Computer Science*, Vol. 146, G. Salton and H.J. Schneider, editors, Springer-Verlag, Berlin, 1983, p. 194-208.
- [46] I.A. Macleod and R.G. Crawford, Document Retrieval as a Database Application, *Information Technology: Research and Development*, Vol. 2, 1983, p. 43-60.
- [47] H.J. Schek, Methods for the Administration of Textual Data in Database Systems, in *Information Retrieval Research*, R.N. Oddy, R.E. Robertson, C.J. van Rijsbergen and P.W. Williams, editors, Butterworths, London, 1981, p. 218-235.
- [48] C.D. Paice, The Automatic Generation of Literature Abstracts: An Approach Based on the Identification of Self-Indicating Phrases, in *Information Retrieval Research*, R.N. Oddy, R.E. Robertson, C.J. van Rijsbergen and P.W. Williams, editors, Butterworths, London, 1981, p. 172-191.
- [49] K. Kropp and G. Walch, A Graph Structured Text Field Index Based on Word Fragments, *Information Processing and Management*, Vol. 17, No. 6, 1981, p. 363-376.
- [50] C. Moulinoux, J.C. Faure, and W. Litwin, Messidor: A Distributed Information Retrieval System, in *Lecture Notes in Computer Science*, Vol. 146, G. Salton and H.J. Schneider, editors, Springer-Verlag, Berlin, 1983, p. 51-61.
- [51] P. Bollmann, Two Axioms for Evaluation Measures in Information Retrieval, in *Research and Development in Information Retrieval*, C.J. van Rijsbergen, editor, Cambridge University Press, 1984, p. 233-246.