

# Origins and Trajectory of the Long Thesaurus Debate

Stella G. Dextre Clarke

Luke House, West Hendred, Wantage, United Kingdom, OX12 8RR, <stella@lukehouse.org>



Stella Dextre Clarke has until recently been an independent consultant specializing in the design of thesauri and other types of knowledge organization systems. She is probably best known for her work on the national and international standards BS 8723 (Structured Vocabularies for Information Retrieval) and ISO 25964 (Thesauri and Interoperability with other Vocabularies). Her work on standards and on taxonomy development was recognized in 2006 when she won the Tony Kent Strix Award for outstanding achievement in information retrieval. Nowadays she is active as Vice-chair of ISKO-UK and Vice-President of ISKO.

Dextre Clarke, Stella G. 2016. "Origins and Trajectory of the Long Thesaurus Debate." *Knowledge Organization* 43 no. 3: 138-144. 33 references.

**Abstract:** The information retrieval thesaurus emerged in the 1950s, settled down to a more-or-less standard format in the 1970s and has continued to evolve marginally since then. Throughout its whole lifetime, doubts have been expressed about its efficacy with emphasis latterly on cost-effectiveness. Prolonged testing of different styles of index language in the 1970s failed to settle the doubts. The arena occupied by the debate has moved from small isolated databases in the post-war era to diverse situations nowadays with the whole Internet at one extreme and small in-house collections at the other. Sophisticated statistical techniques now dominate the retrieval landscape on the Internet but leave opportunities for the thesaurus and other knowledge organization techniques in niches such as image libraries and corporate intranets. The promise of an ontology-driven semantic web with linked data resources opens another opportunity. Thus much scope remains for research to establish the usefulness of the thesaurus in these places and to inspire its continuing evolution.

Received: 31 January 2016; Accepted: 1 February 2016

Keywords: thesaurus, information retrieval, search, thesauri, indexing, controlled vocabulary

## 1.0 Introduction

In February 2015, the United Kingdom chapter (hereafter ISKO-UK) of the International Society for Knowledge Organization (ISKO) held a debate on the proposition: "This House believes that the traditional thesaurus has no place in modern information retrieval." (A report, including links to presentations and recordings, is available at <<http://www.iskouk.org/content/great-debate>>). The outcome aroused so much interest throughout the ISKO membership that the editor of *Knowledge Organization* agreed to devote the whole of this special issue to the topic. The present article provides background for the topic—an abbreviated history of the debate that has dogged the information retrieval thesaurus from its infancy. The reference list is by no means comprehensive but includes key literature to facilitate further research.

## 2.0 The rise of the information retrieval thesaurus

According to Roberts (1984) as well as Krooks and Lancaster (1993) and Aitchison and Dextre Clarke (2004), the thesaurus arrived on the information retrieval (IR) scene in the 1950s. Roberts provides a detailed, fascinating ac-

count of experiments throughout that decade by pioneers such as Charles Bernier, Peter Luhn, Calvin Mooers, Mortimer Taube and others, culminating in emergence of (Roberts 1984, 281) "the first full-scale, operational in-house retrieval thesaurus" in 1959 for use "to solve pressing practical problems at E. I. Du Pont Nemours and Co., Inc., Wilmington, U.S.A."

Until and around that time, developments were intended to overcome perceived weaknesses in the two established approaches to retrieval, namely classification and indexing. For example, it was recognised (Roberts 1984, Sharp 1967) that synonyms and near-synonyms caused difficulties for efficient use of an index and had to be clustered in some way. According to Roberts (1984, 276, quoting from Bernier), researchers were also groping for "some system of showing relationships between semantemes so that the index user can go from those which he knows to all those which he needs to know for a complete search." The role of the thesaurus was debated even in those times (Vickery 1960)—was it to "control" the search vocabulary or to "expand" it? According to Sharp (1967, 206), V. Slamecka in 1963 recommended a joint role: the thesaurus either "*prescribes* the term to be assigned, and/or it *suggests* the concepts and

terms to be considered instead of, or in addition to, terms thought of by indexers without the aid.” Prophetically, Roberts (1984, 274) notes terminological confusion from the start: “the common label ‘thesaurus’ was already being employed to describe retrieval aids distinct both as to form and purpose. This state of confusion was not to be resolved.”

The next big landmark was the 1967 edition of *TEST* (*Thesaurus of Engineering and Scientific Terms* of the Engineers Joint Council) with an appendix setting out the rules and conventions still recognizable in today’s thesaurus standards. The sequence of national and international standards began with the German DIN 1463 in 1972, shortly followed by ISO 2788 and ANSI Z39.19 in 1974 (all so far being for monolingual thesauri) and in 1985 by ISO 5964 (for multilingual thesauri). Successive updating has led to today’s most widely used standards: in the USA we have ANSI/NISO Z39.19-2005 (National Information Standards Organization, 2005), while internationally ISO 25964 (International Organization for Standardization, 2011) has replaced and brought together the standards for both monolingual and multilingual thesauri.

In the intervening period, probably thousands of thesauri have been produced, consuming vast amounts of the intellectual effort needed for construction and maintenance. The 1993 edition of the European Communities Thesaurus Guide described over 600 structured vocabularies that had appeared in at least one of the official languages of the EU (Eurobrokers 1993), while today’s Taxonomy Warehouse <<http://www.taxonomywarehouse.com/>> lists about 650 vocabularies. The unpublished vocabularies used for in-house applications lie beyond counting.

So much effort! But how well do thesauri function in retrieval? Are they really worth the effort? Despite the lack of computing power at that time and notwithstanding his own contribution to the development of thesauri, in 1961 the IBM researcher H. P. Luhn (1961, 1027) cautioned that “where it is possible to handle natural languages automatically, the advantages of interposing an artificial encoding language, however well standardized, are doubtful.” The key weakness he objected to was that inevitably the selection of key concepts for indexing would be (1022) “strictly based on past knowledge and experience. Aspects of documents which might prove important in the light of future inquiry are therefore liable to be made inaccessible.” This criticism, which applies to any controlled vocabulary, is as valid today as it was in the early days. While thesauri have since improved in many ways, they have not evolved to overcome this particular objection. And so the debate goes on.

### 3.0 Early efforts at evaluation

Luhn’s scepticism, applicable only in automated contexts, had little influence in the days before inexpensive computers spread to everyone’s desktops. Throughout the 1960s, indexing (by humans, not machines) was simply essential. Less obvious was what type of indexing language to use. Starting with the ground-breaking Cranfield tests in 1959-1964 (Dextre Clarke 2001a, Robertson 2008), huge efforts were made to study different characteristics of the vocabularies. This was not easy, because variables such as indexing quality and comprehensiveness have an impact that is hard to separate from the vocabulary characteristics under study. The searcher’s choice of terms and judgements on relevance exert an influence too. Nonetheless, thesaural inclusion of synonyms and lexical variants in the scope of a single search term was investigated and found to deliver clear benefits in retrieval. But the contribution made by other vocabulary attributes was much less clear. E. Michael Keen followed up on the Cranfield work by testing more styles of index language over the next decade, and remained sceptical about the value of using any of them. In 1973 he concluded (33), “no really large differences in retrieval effectiveness and efficiency were observed” and also (34) “acceptance of the view that the index language is a sub-system of minor import and that uncontrolled languages are as efficacious, or more so, than controlled ones, is still mixed.”

In summary the intensive testing in this period was very worthwhile in establishing “recall” and “precision” as the basic measures of retrieval performance but failed to settle other doubts about the value of a controlled vocabulary. It seems there were just too many inter-related variables for each one to be isolated and tested over a statistically valid sample of documents and queries, and the subjectivity of relevance judgements has always added extra doubts.

Summaries of the evaluation findings and key references can be found in Svenonius (1986), Soergel (1994), Dextre Clarke (2001a) and Robertson (2008). Svenonius (1986, 331-2) observed “An experiment sophisticated and large enough to control all of the above variables [concerning effectiveness of controlled vocabularies] has never been conducted and probably never will.” Soergel (1994, 589) pointed out,

Indexing characteristics [in which he included attributes of the indexing vocabulary] and their effects on retrieval are so complex that they largely defy study in artificial test situations. Most experiments fail to account for important interactions among factors as they occur in the real world, and thus give results that mislead more than they enlighten, results that have

little meaning for the assessment or improvement of operational retrieval systems.

The arguments were not dispelled by the gradual spread of computers but continued throughout the 1970s, 1980s and 1990s. A great deal of evidence (Lancaster (1998), Soergel (1994) and Milstead (1994) point to just some of this literature) confirmed that indexing by humans is often inconsistent, even when a controlled vocabulary is used. But the thesaurus retained its band of supporters. Milstead (1994, 579) concluded “there is some research—and many anecdotes—that shows that the best retrieval situation is a mix of human indexing using a controlled vocabulary and the ability to search full text. But there appears to be no work that will assist the index designer in deciding how much expensive human effort is really warranted.” The 4<sup>th</sup> edition of the well-known thesaurus construction manual from Aitchison Gilchrist and Bawden (2000, 5-7) sets out pros and cons of the “natural *versus* controlled language” debate, concluding that “more research is needed” to determine the modifications most likely to enhance thesaurus use in full-text systems.

#### 4.0 A changing search environment

Much of the early debate had been conducted in a scientific arena, as a response to the growth in research literature. For example, the Cranfield experiments were discussed among scholars using science-based methodologies. But the business world had an interest too, and the demand for high-powered retrieval became much more widespread in the nineties. In 1994, the search engine Yahoo! began to popularise the “taxonomy” as an alternative approach to finding information—via directory browse rather than keyword search. Despite Yahoo!’s subsequent loss of market share, the popularity of the taxonomy persists today in many corporate intranets as well as Internet websites and portals. Even back in the twentieth century, the perception of high costs associated with development and maintenance of a thesaurus, and especially with using it to index a large collection, was probably more influential in the decline of the thesaurus than any theory-based arguments about effectiveness in retrieval.

Emergence of the browsable taxonomy stimulated a number of software companies to develop products enabling automatic or semi-automatic categorization of resources and has had several effects on the thesaurus debate:

- The argument about the inconsistency of human indexing has received much publicity from vendors (who can claim their automatic categorization software is 100% consistent—even if it is consistently wrong!)
  - Ranged alongside the taxonomy, the thesaurus has an added competitor-cum-collaborator among controlled vocabularies (including classification schemes and subject heading schemes).
  - The discussion has become quite muddy as many authors and practitioners use the terms “thesaurus,” “taxonomy,” “controlled vocabulary,” etc. quite loosely and interchangeably.
- At the same time, the debate and the decisions about which systems and tools to adopt has extended from researchers and developers to executive staff who may be influenced much more by obvious costs, hearsay and advertising than by the minutiae of scientific argument.
- Since the 1990s, other dramatic changes have transformed the retrieval landscape:
- The development of web browsers popularized the Internet and hence expanded the scope of retrieval from discrete databases to a worldwide network of resources, still growing. While the aim of “controlling” a vocabulary seems feasible enough inside one organization or in any centrally managed facility, in the uncontrolled, external Internet, any such hope would be naïve.
  - Keeping pace with extraordinary growth of computer power (both in capacity and in speed, as well as cost reduction), the annual Text REtrieval Conference (TREC) has enabled evaluation of IR technology on a massive scale. Established in 1992 with powerful backing and ample funding, and still active, its annual cycle of testing has helped software developers refine their algorithms in a wide variety of contexts. Papers from the whole series of conferences can be found at <<http://trec.nist.gov/proceedings/proceedings.html>>.
  - Methods based on algebraic and probabilistic models, supplemented by Google’s famous PageRank algorithm, have been implemented by search engines that are publicly available worldwide. Thus, rapid effective retrieval, driven by statistical techniques without a thesaurus in sight, is available to the lowliest computer user.
  - The semantic web, dream-child of Tim Berners-Lee and unknown until his famous 2001 article in *Scientific American* (Berners-Lee et al. 2001), immediately excited the world’s attention and the prospect has spawned a different approach to information navigation and retrieval. Complementing the statistical methods of the search engines, the semantic web should work through hypertext links. To manipulate these, a computer-based intelligent agent relies on ontologies with rigorously defined relationships between entities. As ontologies rapidly came into fashion, they inspired a fresh look at knowledge organization systems (KOS) such as thesauri.

These 21<sup>st</sup> century developments render much of the debate literature outlined above irrelevant to decisions about today's IR systems. The comparison of natural and controlled language in Aitchison, Gilchrist and Bawden (2000, 5-7), for example, makes no mention of Google (then in its infancy) or the power of statistical approaches to retrieval.

### 5.0 The debate in recent times

In 2008, ISKO UK took a first step towards airing our current round of the debate, by inviting Stephen Robertson to address us on “the state of information retrieval.” His presentation and an audio recording are freely available at <http://www.iskouk.org/content/state-information-retrieval-researchers-view>. Back in the seventies, he was one of the pioneers in the development of probabilistic methods for retrieval (Robertson and Sparck Jones 1976; Robertson, Walker and Beaulieu 1997; Robertson and Sparck Jones 1997), developing key principles and the “BM25” algorithm that fed into the statistical algorithms now used by Google and other search engines and has been heavily involved in many of the TREC experiments. Robertson invited us to view the first 10-15 years of TREC in a simplified way, as a competition between a big range of statistical approaches to retrieval on the one hand and all the other approaches (NLP approaches, ontology-based approaches, KO methods, etc.) on the other. “And if you see it like that, you have to say, statistical approaches won, simply. They were overwhelmingly more successful.” Similarly, if you view the history of search engines as being a competition between statistical methods and others (Yahoo!, directories, etc.), he maintained that the statistical methods come out on top. This argument is hard to defeat while so many of us routinely rely on Google!

A cynic might object that Robertson is a statistician whose judgments are guided by his own interests. To be fair, though, he is also an information scientist who long ago studied KO methods under the guidance of Jason Farradane, that ardent advocate of faceted classification and relational indexing. After highlighting the usefulness of statistics Robertson pointed to the continuing opportunities for KO: “we also need to deal better with smaller systems where extensive machine learning is not feasible e.g. desktop, within-site, specialist.”

While statistical methods undoubtedly dominate web search, retrieval is needed in many other contexts. When Google's enterprise search products are implemented in corporate intranets, many customers are disappointed with the results. This happens because inside a single organization, even a multinational company, the volume of users and searches may be insufficient to fuel the probabilistic

algorithms. Furthermore, the first screen of results, showing the ten most highly ranked items, may not satisfy the searcher who is looking for a particular document that he knows to exist, but simply cannot find. Hence the dissatisfaction reported by White elsewhere in this issue, and a continuing need for KO tools and expertise.

The contexts where statistical methods are not enough include:

- Small and medium-scale in-house collections
- Electronic document and records management systems (EDRMS)
- Knowledge-bases used to hold an organization's store of expertise
- Collections with text in multiple languages
- Bibliographic databases
- Heritage collections already indexed with a controlled vocabulary
- Multimedia resources with little text for the statistics to work on—especially music and still images. (See MacFarlane's paper elsewhere in this issue).

In places such as these, the opportunity for thesauri and other KO techniques continues, although it cannot be guaranteed forever. All workplaces today feel pressure to increase efficiency and reduce costs—especially the cost of intellectual labour. The costs of indexing as well as vocabulary maintenance are hard to conceal; thesaurus advocates have every incentive to prove and demonstrate the cost-effectiveness of their methods.

To add to these niches, interest in an incipient semantic web has brought a new opportunity for thesauri and other KOSs to play a key role, via linked data—preferably linked open data (LOD) so that the whole world has access. For a thesaurus, the idea is to publish it on the Web giving each concept a URI (Unique Resource Identifier) that enables linkage to any other relevant resource. By following these links, users can navigate directly from one place to the next. For example, Baca and Gill (2015) describe the encoding as LOD of three major multilingual KOSs—the Getty's *Art & Architecture Thesaurus*, the *Union List of Artist Names* and the *Thesaurus of Geographic Names*. These vocabularies are used worldwide by museums and other cultural heritage collections to index or catalogue their holdings. That means anyone searching for works by a particular artist, or for a type of artefact, can easily be routed through to the repositories where the items are held and possibly displayed. Portals such as Europeana (<http://www.europeana.eu/>) and ARTstor (<http://www.artstor.org/>) provide the search and navigation interfaces that set up the routing for participating institutions. Elsewhere in this issue Tudhope and Binding spell out more of the opportunities, while as a case study Kempf and Neubert de-

scribe how exploitation of Web technologies has revolutionized the retrieval capabilities of the STW Thesaurus for Economics. Could it be that a family of SKOS (Simple Knowledge Organization System)-enabled thesauri like STW will soon form an interlinked structure among the Wikipedia family as envisioned by Garcia-Marco?

The emerging semantic web has led also to much interest in ontology development with a knock-on effect on the study of thesauri. People nowadays ask why thesauri do not have rigorously specified relationships in the style of ontologies. A superficial answer is straightforward—little need for them was perceived until the twenty-first century. Until about 20 years ago, the only purpose of the associative links (RT) in a thesaurus was to help the indexer or searcher navigate the thesaurus and think of more terms to use “instead of or as well as” the ones first thought of (International Organization for Standardization 2011; Dextre Clarke 2001b; Lancaster 1986; Roberts 1984). Since the early days of thesaurus development, editors and researchers (Willetts 1975, Aitchison, Gilchrist and Bawden 2000, 62-66; Dextre Clarke 2001b, 47) have been aware of the multiplicity of relationship types eligible for the name “associative.” But since the associative links were not used “directly” in retrieval, designation of the type of RT was not among the features tested in the Cranfield experiments, or (so far as I can establish) in any subsequent IR evaluation. (While relational indexing was among the variants evaluated by Keen (1973), the relationships in question were those found syntagmatically in documents and queries, rather than the paradigmatic relationships captured by the indexing vocabulary.) Nowadays, however, there may be potential for retrieval software to exploit particular relationship types, which could in turn justify the effort needed to identify them.

In view of this potential for establishing and exploiting differentiated relationships between concepts, numerous initiatives are under way. For example, relationships such as “has disease,” “has pest” and “controls” have been added to AGROVOC, the long-established thesaurus of the Food and Agriculture Organization of the United Nations (Lauser 2001, Soergel et al. 2004, Caracciolo and Keizer 2015). Since the relationships appropriate to different domains, and the uses to which they may be put, are quite varied, such efforts do not yet appear in the thesaurus standards, and it is hard to generalize about their cost-effectiveness.

Could it be that thesauri are about to evolve in the direction of ontologies? This would make them costlier than ever, since the establishment of ontology-style relationships demands greater rigour from people with excellent knowledge of the subject domain in question. After the barrage of criticism that thesauri are not cost-effective, it would be ironic to find their salvation resting

on increased costs! Ironic, but not impossible, if the extra expenditure brings much greater effectiveness. Elsewhere in this issue Hjørland, concerned at the way in which Google and other search engines have overtaken the thesaurus, suggests that retrieval performance would be enhanced by inclusion of relationship types appropriate to the domain of the thesaurus. It is to be hoped that more evaluations, including cost evaluations, will be published to fuel a continuing debate.

## 6.0 Conclusions

Clearly our debate is by no means concluded. But it does need to move on from the bare question of “whether” the thesaurus has a place to more constructive questions such as:

- “What or where” is the place of the thesaurus in modern IR?
- How should the thesaurus “evolve” to meet the needs of modern IR?
- “How best to integrate” a thesaurus with other tools and technologies in IR applications (not forgetting options to enable a user-friendly interactive interface for expanding or refining the search)?
- “What software” and other tools are needed to create, maintain and manipulate the evolving thesaurus?
- Is there a place for “new styles of KOS” combining thesaural features with those of other types of KOSs?
- How can we test and demonstrate the “cost-effectiveness” of thesauri and other KOSs?

An exploration of these questions offers plenty of scope for research as well as the practice of KO.

## References

- Aitchison, Jean and Stella Dextre Clarke. 2004. “The Thesaurus: A Historical Viewpoint, with a Look to the Future.” *Cataloging & Classification Quarterly* 37 nos. 3/4: 5-21.
- Aitchison, Jean, Alan Gilchrist and David Bawden. 2000. *Thesaurus Construction and Use: A Practical Manual*. 4th ed. London: Aslib.
- American National Standards Institute. 1974. *American National Standard Guidelines for Thesaurus Structure, Construction and Use*. New York: American National Standards Institute; ANSI Z39.19-1974.
- Baca, Murtha and Melissa Gill. 2015. “Encoding Multilingual Knowledge systems in the Digital Age: The Getty Vocabularies.” *Knowledge Organization* 42: 232-43.

- Berners-Lee, Tim, James Hendler and Ora Lassila. 2001. "The Semantic Web." *Scientific American* May 2001 (0501). <http://www.sciam.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21>
- Caracciolo, Caterina and Johannes Keizer. 2014. "What KOS Can Do, With the Proper Tools Available. About AGROVOC, edited in VocBench and used in the AGRIS Web Application." In *Knowledge Organization—Making a Difference: ISKO UK Biennial Conference, Jul 13-14 2014*, ed. Stella Dextre Clarke. London. <http://www.iskouk.org/sites/default/files/CaraccioloSlidesISKO-UK2015.pdf>
- Dextre Clarke, Stella G. 2001a. "Organising Access to Information by Subject." In *Handbook of Information Management (previously Handbook of Special Librarianship and Information Work)*, 8th edition, ed. Alison Scammell. London: Aslib; pp. 72-110.
- Dextre Clarke, Stella G. 2001b. "Thesaural Relationships." In *Relationships in the Organization of Knowledge*, Carol A. Bean and Rebecca Green, eds. Dordrecht: Kluwer, pp. 37-52.
- Eurobrokers, on behalf of the Commission of the European Communities. 1992. *Thesaurus Guide: Analytical Directory of Selected Vocabularies for Information Retrieval*, 2nd ed. Luxembourg: European Communities.
- International Organization for Standardization. 2011. ISO 25964-1: *Information and Documentation—Thesauri and interoperability with other vocabularies - Part 1: Thesauri for information retrieval*. Geneva: International Organization for Standardization.
- International Organization for Standardization. 2013. ISO 25964-2: *Information and Documentation—Thesauri and Interoperability with other Vocabularies - Part 2: Interoperability with other Vocabularies*. Geneva: International Organization for Standardization.
- International Organization for Standardization. 1974. ISO 2788-1974. *Documentation—Guidelines for the Establishment and Development of Monolingual Thesauri*, 1st ed. Geneva: International Organization for Standardization.
- International Organization for Standardization. 1985. ISO 5964-1985. *Documentation—Guidelines for the Establishment and Development of Multilingual Thesauri*. Geneva: International Organization for Standardization.
- Keen, E. Michael. 1973. The Aberystwyth index languages test. *Journal of Documentation* 29:1-35.
- Krooks David A. and F.W. Lancaster. 1993. "The Evolution of Guidelines for Thesaurus Construction." *Libri* 43:326-42.
- Lancaster, F W. 1998. *Indexing and Abstracting in Theory and Practice*. London: LA Publishing.
- Lancaster, F W. 1986. *Vocabulary Control for Information Retrieval*. Arlington, Virginia: Information Resources Press.
- Lauser, Boris. 2001. "From Thesauri to Ontologies: A Short Case Study in the Food Safety Area in How Ontologies are More Powerful than Thesauri—From Thesauri to RDFS to OWL." In *ECDL 2001 OAI Workshop. Agricultural Information and Knowledge Management Papers*. <ftp://ftp.fao.org/docrep/fao/008/af236e/af236e00.pdf>
- Luhn, H. P. 1961. "The Automatic Derivation of Information Retrieval Encodements from Machine-Readable Texts." In *Information Retrieval and Machine Translation*, Kent, Allen., Editor.. New York: Interscience Publication, pp. 1021-28.
- Milstead, Jessica L. 1994. "Needs for Research in Indexing." *Journal of the American Society for Information Science*. 45:577-82.
- National Information Standards Organization. 2005. *ANSI/NISO Z39.19-2005: Guidelines for the Construction, Format, and Management of Monolingual Controlled Vocabularies*. Maryland, USA: National Information Standards Organization.
- Roberts, Norman. 1984. "The Pre-History of the Information Retrieval Thesaurus." *Journal of Documentation* 40:271-85.
- Robertson, S. E. and K. Sparck Jones. 1976. "Relevance Weighting of Search Terms." *Journal of the American Society for Information Science* 27:129-46.
- Robertson, S. E. and K. Sparck Jones. 1997. *Simple, Proven Approaches to Text Retrieval*. Technical Report 356. Cambridge, UK: Computer Laboratory, University of Cambridge.
- Robertson, S., S. Walker and M. Beaulieu. 1997. "Laboratory Experiments with Okapi: Participation in the TREC Programme." *Journal of Documentation* 53:20-34.
- Robertson, Stephen. 2008. "On the History of Evaluation in IR." *Journal of Information Science* 34:439-56.
- Sharp, J. R. 1967. "Information Retrieval." In *Handbook of Special Librarianship and Information Work*, 3rd ed, Wilfred Ashworth, editor. London: Aslib, pp. 141-232.
- Soergel, Dagobert. 1994. "Indexing and Retrieval Performance: The Logical Evidence." *Journal of the American Society for Information Science* 45:589-99.
- Soergel, Dagobert, Boris Lauser, Anita Liang, Frehiwot Fisseha, Johannes Keizer and Stephen Katz. 2004. "Reengineering Thesauri for New Applications: The AGROVOC Example." *Journal of Digital Information* 4 no. 4: Article No. 257. <http://journals.tdl.org/jodi/article/view/jodi-126/111>
- Svenonius, Elaine. 1986. "Unanswered Questions in the Design of Controlled Vocabularies." *Journal of the American Society for Information Science* 37:331-40.
- Thesaurus of Engineering and Scientific Terms (TEST)*. 1967. New York: Engineers Joint Council and U.S. Department of Defense.

Vickery, B. C. 1960. "Thesaurus: A New Word in Documentation." *Journal of Documentation* 6:181-89.

Willetts, Margaret. 1975. "An Investigation of the Nature of the Relation between Terms in Thesauri." *Journal of Documentation* 31:158-84.