

Untangling Text Data Mining

Marti A. Hearst

School of Information Management & Systems

University of California, Berkeley

102 South Hall

Berkeley, CA 94720-4600

<http://www.sims.berkeley.edu/~hearst>

Abstract

The possibilities for data mining from large text collections are virtually untapped. Text expresses a vast, rich range of information, but encodes this information in a form that is difficult to decipher automatically. Perhaps for this reason, there has been little work in text data mining to date, and most people who have talked about it have either conflated it with information access or have not made use of text directly to discover heretofore unknown information.

In this paper I will first define data mining, information access, and corpus-based computational linguistics, and then discuss the relationship of these to text data mining. The intent behind these contrasts is to draw attention to exciting new kinds of problems for computational linguists. I describe examples of what I consider to be real text data mining efforts and briefly outline recent ideas about how to pursue exploratory data analysis over text.

1 Introduction

The nascent field of text data mining (TDM) has the peculiar distinction of having a name and a fair amount of hype but as yet almost no practitioners. I suspect this has happened because people assume TDM is a natural extension of the slightly less nascent field of data mining (DM), also known as knowledge discovery in databases (Fayyad and Uthurusamy, 1999), and information archeology (Brachman et al., 1993). Additionally, there are some disagreements about what actually constitutes data mining. It turns out that “mining” is not a very good metaphor for what people in the field actually do. Mining implies extracting precious nuggets of ore from otherwise worthless rock. If data mining really followed this metaphor, it would mean that people were discovering new

factoids within their inventory databases. However, in practice this is not really the case. Instead, data mining applications tend to be (semi)automated discovery of trends and patterns across very large datasets, usually for the purposes of decision making (Fayyad and Uthurusamy, 1999; Fayyad, 1997). Part of what I wish to argue here is that in the case of text, it can be interesting to take the mining-for-nuggets metaphor seriously.

The various contrasts discussed below are summarized in Table 1.

2 TDM vs. Information Access

It is important to differentiate between text data mining and information access (or information retrieval, as it is more widely known).

The goal of information access is to help users find documents that satisfy their information needs (Baeza-Yates and Ribeiro-Neto, 1999). The standard procedure is akin to looking for needles in a needlestack – the problem isn’t so much that the desired information is not known, but rather that the desired information coexists with many other valid pieces of information. Just because a user is currently interested in NAFTA and not Furbies does not mean that all descriptions of Furbies are worthless. The problem is one of homing in on what is currently of interest to the user.

As noted above, the goal of data mining is to discover or derive new information from data, finding patterns across datasets, and/or separating signal from noise. The fact that an information retrieval system can return a document that contains the information a user requested implies that no new discovery is being made: the information had to have already been known to the author of the text; otherwise the author could not have written it down.

I have observed that many people, when asked about text data mining, assume it should have something to do with “making things easier to find on the web”. For example, the description of the KDD-97 panel on Data Mining and the Web stated:

... Two challenges are predominant for data mining on the Web. The first goal is to help users in finding useful information on the Web and in discovering knowledge about a domain that is represented by a collection of Web-documents. The second goal is to analyse the transactions run in a Web-based system, be it to optimize the system or to find information about the clients using the system.¹

This search-centric view misses the point that we might actually want to treat the information in the web as a large knowledge base from which we can extract new, never-before encountered information (Craven et al., 1998).

On the other hand, the results of certain types of text processing can yield *tools* that indirectly *aid* in the information access process. Examples include text clustering to create thematic overviews of text collections (Cutting et al., 1992; Chalmers and Chitson, 1992; Rennison, 1994; Wise et al., 1995; Lin et al., 1991; Chen et al., 1998), automatically generating term associations to aid in query expansion (Peat and Willett, 1991; Voorhees, 1994; Xu and Croft, 1996), and using co-citation analysis to find general topics within a collection or identify central web pages (White and McCain, 1989; Larson, 1996; Kleinberg, 1998).

Aside from providing tools to aid in the standard information access process, I think text data mining can contribute along another dimension. In future I hope to see information access systems supplemented with tools for exploratory data analysis. Our efforts in this direction are embodied in the LINDI project, described in Section 5 below.

3 TDM and Computational Linguistics

If we extrapolate from data mining (as practiced) on numerical data to data mining from text collections, we discover that there already

¹<http://www.aaai.org/Conferences/KDD/1997/kdd97-schedule.html>

exists a field engaged in text data mining: corpus-based computational linguistics! Empirical computational linguistics computes statistics over large text collections in order to discover useful patterns. These patterns are used to inform algorithms for various subproblems within natural language processing, such as part-of-speech tagging, word sense disambiguation, and bilingual dictionary creation (Armstrong, 1994).

It is certainly of interest to a computational linguist that the words “prices, prescription, and patent” are highly likely to co-occur with the medical sense of “drug” while “abuse, paraphernalia, and illicit” are likely to co-occur with the illegal drug sense of this word (Church and Liberman, 1991). This kind of information can also be used to improve information retrieval algorithms. However, the kinds of patterns found and used in computational linguistics are not likely to be what the general business community hopes for when they use the term text data mining.

Within the computational linguistics framework, efforts in automatic augmentation of existing lexical structures seem to fit the data-mining-as-ore-extraction metaphor. Examples include automatic augmentation of WordNet relations (Fellbaum, 1998) by identifying lexico-syntactic patterns that unambiguously indicate those relations (Hearst, 1998), and automatic acquisition of subcategorization data from large text corpora (Manning, 1993). However, these serve the specific needs of computational linguistics and are not applicable to a broader audience.

4 TDM and Category Metadata

Some researchers have claimed that text categorization should be considered text data mining. Although analogies can be found in the data mining literature (e.g., referring to classification of astronomical phenomena as data mining (Fayyad and Uthurusamy, 1999)), I believe when applied to text categorization this is a misnomer. Text categorization is a boiling down of the specific content of a document into one (or more) of a set of pre-defined labels. This does not lead to discovery of new information; presumably the person who wrote the document knew what it was about. Rather, it produces a

	Finding Patterns	Finding Nuggets	
		Novel	Non-Novel
Non-textual data	standard data mining	?	database queries
Textual data	computational linguistics	real TDM	information retrieval

Table 1: A classification of data mining and text data mining applications.

compact summary of something that is already known.

However, there are two recent areas of inquiry that make use of text categorization and do seem to fit within the conceptual framework of discovery of trends and patterns within textual data for more general purpose usage.

One body of work uses text category labels (associated with Reuters newswire) to find “unexpected patterns” among text articles (Feldman and Dagan, 1995; Dagan et al., 1996; Feldman et al., 1997). The main approach is to compare distributions of category assignments within subsets of the document collection. For instance, distributions of commodities in country C1 are compared against those of country C2 to see if interesting or unexpected trends can be found. Extending this idea, one country’s export trends might be compared against those of a set of countries that are seen as an economic unit (such as the G-7).

Another effort is that of the DARPA Topic Detection and Tracking initiative (Allan et al., 1998). While several of the tasks within this initiative are standard text analysis problems (such as categorization and segmentation), there is an interesting task called On-line New Event Detection, whose input is a stream of news stories in chronological order, and whose output is a yes/no decision for each story, made at the time the story arrives, indicating whether the story is the first reference to a newly occurring event. In other words, the system must detect the first instance of what will become a series of reports on some important topic. Although this can be viewed as a standard classification task (where the class is a binary assignment to the new-event class) it is more in the spirit of data mining, in that the focus is on discovery of the beginning of a new theme or trend.

The reason I consider this examples – using multiple occurrences of text categories to detect trends or patterns – to be “real” data min-

ing is that they use text metadata to tell us something about the world, outside of the text collection itself. (However, since this application uses metadata associated with text documents, rather than the text directly, it is unclear if it should be considered text data mining or standard data mining.) The computational linguistics applications tell us about how to improve language analysis, but they do not discover more widely usable information.

5 Text Data Mining as Exploratory Data Analysis

Another way to view text data mining is as a process of exploratory data analysis (Tukey, 1977; Hoaglin et al., 1983) that leads to the discovery of heretofore unknown information, or to answers for questions for which the answer is not currently known.

Of course, it can be argued that the standard practice of reading textbooks, journal articles and other documents helps researchers in the discovery of new information, since this is an integral part of the research process. However, the idea here is to use text for discovery in a more direct manner. Two examples are described below.

5.1 Using Text to Form Hypotheses about Disease

For more than a decade, Don Swanson has eloquently argued why it is plausible to expect new information to be derivable from text collections: experts can only read a small subset of what is published in their fields and are often unaware of developments in related fields. Thus it should be possible to find useful linkages between information in related literatures, if the authors of those literatures rarely refer to one another’s work. Swanson has shown how chains of causal implication within the medical literature can lead to hypotheses for causes of rare diseases, some of which have received supporting experimental evidence (Swanson, 1987;

Swanson, 1991; Swanson and Smalheiser, 1994; Swanson and Smalheiser, 1997).

For example, when investigating causes of migraine headaches, he extracted various pieces of evidence from titles of articles in the biomedical literature. Some of these clues can be paraphrased as follows:

- stress is associated with migraines
- stress can lead to loss of magnesium
- calcium channel blockers prevent some migraines
- magnesium is a natural calcium channel blocker
- spreading cortical depression (SCD) is implicated in some migraines
- high levels of magnesium inhibit SCD
- migraine patients have high platelet aggregability
- magnesium can suppress platelet aggregability

These clues suggest that magnesium deficiency may play a role in some kinds of migraine headache; a hypothesis which did not exist in the literature at the time Swanson found these links. The hypothesis has to be tested via non-textual means, but the important point is that a new, potentially plausible medical hypothesis was derived from a combination of text fragments and the explorer's medical expertise. (According to Swanson (1991), subsequent study found support for the magnesium-migraine hypothesis (Ramadan et al., 1989).)

This approach has been only partially automated. There is, of course, a potential for combinatorial explosion of potentially valid links. Beferman (1998) has developed a flexible interface and analysis tool for exploring certain kinds of chains of links among lexical relations within WordNet.² However, sophisticated new algorithms are needed for helping in the pruning process, since a good pruning algorithm will want to take into account various kinds of semantic constraints. This may be an interesting area of investigation for computational linguists.

²See <http://www.link.cs.cmu.edu/lexfn>

5.2 Using Text to Uncover Social Impact

Switching to an entirely different domain, consider a recent effort to determine the effects of publicly financed research on industrial advances (Narin et al., 1997). After years of preliminary studies and building special purpose tools, the authors found that the technology industry relies more heavily than ever on government-sponsored research results. The authors explored relationships among patent text and the published research literature, using a procedure which was reported as follows in Broad (1997):

The CHI Research team examined the science references on the front pages of American patents in two recent periods – 1987 and 1988, as well as 1993 and 1994 – looking at all the 397,660 patents issued.

It found 242,000 identifiable science references and zeroed in on those published in the preceding 11 years, which turned out to be 80 percent of them. Searches of computer databases allowed the linking of 109,000 of these references to known journals and authors' addresses. After eliminating redundant citations to the same paper, as well as articles with no known American author, the study had a core collection of 45,000 papers. Armies of aides then fanned out to libraries to look up the papers and examine their closing lines, which often say who financed the research. That detective work revealed an extensive reliance on publicly financed science.

...

Further narrowing its focus, the study set aside patents given to schools and governments and zeroed in on those awarded to industry. For 2,841 patents issued in 1993 and 1994, it examined the peak year of literature references, 1988, and found 5,217 citations to science papers.

Of these, it found that 73.3 percent had been written at public institutions – universities, government labs and other public agencies, both in the United States and abroad.

Thus a heterogeneous mix of operations was required to conduct a complex analyses over large text collections. These operations included:

- 1 Retrieval of articles from a particular collection (patents) within a particular date range.
- 2 Identification of the citation pool (articles cited by the patents).
- 3 Bracketing of this pool by date, creating a new subset of articles.
- 4 Computation of the percentage of articles that remain after bracketing.
- 5 Joining these results with those of other collections to identify the publishers of articles in the pool.
- 6 Elimination of redundant articles.
- 7 Elimination of articles based on an attribute type (author nationality).
- 8 Location of full-text versions of the articles.
- 9 Extraction of a special attribute from the full text (the acknowledgement of funding).
- 10 Classification of this attribute (by institution type).
- 11 Narrowing the set of articles to consider by an attribute (institution type).
- 12 Computation of statistics over one of the attributes (peak year)
- 13 Computation of the percentage of articles for which one attribute has been assigned another attribute type (whose citation attribute has a particular institution attribute).

Because all the data was not available online, much of the work had to be done by hand, and special purpose tools were required to perform the operations.

5.3 The LINDI Project

The objectives of the LINDI project³ are to investigate how researchers can use large text collections in the discovery of new important information, and to build software systems to help support this process. The main tools for discovering new information are of two types: support for issuing sequences of queries and related operations across text collections, and tightly coupled statistical and visualization tools for the examination of associations among concepts that co-occur within the retrieved documents. Both sets of tools make use of attributes associated specifically with text collections and

³LINDI: Linking Information for Novel Discovery and Insight.

their metadata. Thus the broadening, narrowing, and linking of relations seen in the patent example should be tightly integrated with analysis and interpretation tools as needed in the biomedical example.

Following Amant (1996), the interaction paradigm is that of a mixed-initiative balance of control between user and system. The interaction is a cycle in which the system suggests hypotheses and strategies for investigating these hypotheses, and the user either uses or ignores these suggestions and decides on the next move.

We are interested in an important problem in molecular biology, that of automating the discovery of the function of newly sequenced genes (Walker et al., 1998). Human genome researchers perform experiments in which they analyze co-expression of tens of thousands of novel and known genes simultaneously.⁴ Given this huge collection of genetic information, the goal is to determine which of the novel genes are medically interesting, meaning that they are co-expressed with already understood genes which are known to be involved in disease. Our strategy is to explore the biomedical literature, trying to formulate plausible hypotheses about which genes are of interest.

Most information access systems require the user to execute and keep track of tactical moves, often distracting from the thought-intensive aspects of the problem (Bates, 1990). The LINDI interface provides a facility for users to build and so reuse sequences of query operations via a drag-and-drop interface. These allow the user to repeat the same sequence of actions for different queries. In the gene example, this allows the user to specify a sequence of operations to apply to one co-expressed gene, and then iterate this sequence over a list of other co-expressed genes that can be dragged onto the template. (The Visage interface (Derthick et al., 1997) implements this kind of functionality within its information-centric framework.) These include the following operations (see Figure 1):

- Iteration of an operation over the items within a set. (This allows each item retrieved in a previous query to be use as a

⁴A gene g' co-expresses with gene g when both are found to be activated in the same cells at the same time with much more likelihood than chance.

search terms for a new query.)

- Transformation, i.e., applying an operation to an item and returning a transformed item (such as extracting a feature).
- Ranking, i.e., applying an operation to a set of items and returning a (possibly) re-ordered set of items with the same cardinality.
- Selection, i.e., applying an operation to a set of items and returning a (possibly) reordered set of items with the same or smaller cardinality.
- Reduction, i.e., applying an operation to one or more sets of items to yield a singleton result (e.g., to compute percentages and averages).

This system will allow maintenance of several different types of history including history of commands issued, history of strategies employed, and history of hypotheses tested. For the history view, we plan to use a “spreadsheet” layout (Hendry and Harper, 1997) as well as a variation on a “slide sorter” view which Visage uses for presentation creation but not for history retention (Roth et al., 1997).

Since gene function discovery is a new area, there is not yet a known set of exploration strategies. So initially the system must help an expert user generate and record good exploration strategies. The user interface provides a mechanism for recording and modifying sequences of actions. These include facilities that refer to metadata structure, allowing, for example, query terms to be expanded by terms one level above or below them in a subject hierarchy. Once a successful set of strategies has been devised, they can be re-used by other researchers and (with luck) by an automated version of the system. The intent is to build up enough strategies that the system will begin to be used as an assistant or advisor (Amant, 1996), ranking hypotheses according to projected importance and plausibility.

Thus the emphasis of this system is to help automate the tedious parts of the text manipulation process and to integrate underlying computationally-driven text analysis with human-guided decision making within exploratory data analysis over text.

6 Summary

For almost a decade the computational linguistics community has viewed large text collections as a resource to be tapped in order to produce better text analysis algorithms. In this paper, I have attempted to suggest a new emphasis: the use of large online text collections to discover new facts and trends about the world itself. I suggest that to make progress we do not need fully artificial intelligent text analysis; rather, a mixture of computationally-driven and user-guided analysis may open the door to exciting new results.

Acknowledgements. Hao Chen, Ketan Mayer-Patel, and Vijayshankar Raman helped design and did all the implementation of the first LINDI prototype.

References

- J. Allan, J. Carbonell, G. Doddington, J. Yamron, and Y. Yang. 1998. Topic detection and tracking pilot study: Final report. In *Proceedings of the DARPA Broadcast News Transcription and Understanding Workshop*, pages 194–218.
- Robert St. Amant. 1996. *A Mixed-Initiative Planning Approach to Exploratory Data Analysis*. Ph.D. thesis, University of Massachusetts, Amherst.
- Susan Armstrong, editor. 1994. *Using Large Corpora*. MIT Press.
- Ricardo Baeza-Yates and Berthier Ribeiro-Neto. 1999. *Modern Information Retrieval*. Addison-Wesley Longman Publishing Company.
- Marcia J. Bates. 1990. The berry-picking search: User interface design. In Harold Thimbleby, editor, *User Interface Design*. Addison-Wesley.
- Douglas Beeferman. 1998. Lexical discovery with an enriched semantic network. In *Proceedings of the ACL/COLING Workshop on Applications of WordNet in Natural Language Processing Systems*, pages 358–364.
- R. J. Brachman, P. G. Selfridge, L. G. Terveen, B. Altman, A. Borgida, F. Halper, T. Kirk, A. Lazar, D. L. McGuinness, and L. A. Resnick. 1993. Integrated support for data archaeology. *International Journal of Intelligent and Cooperative Information Systems*, 2(2):159–185.
- William J. Broad. 1997. Study finds public science is pillar of industry. In *The New York Times*, May 13.
- Matthew Chalmers and Paul Chitson. 1992. Bead: Exploration in information visualization. In *Proceedings of the 15th Annual International ACM/SIGIR Conference*, pages 330–337, Copenhagen, Denmark.

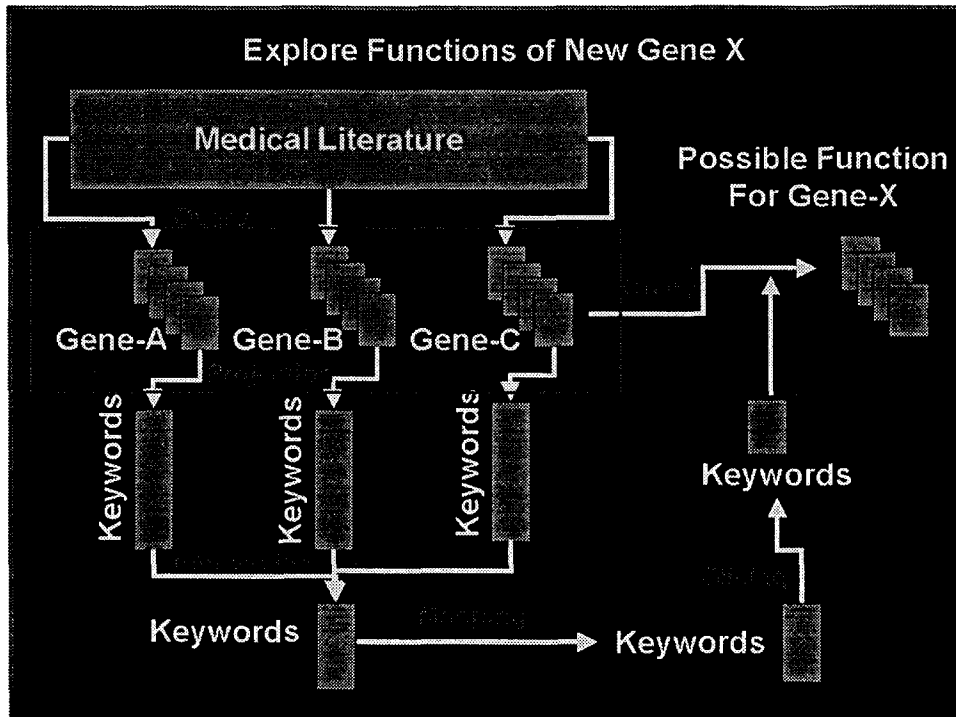


Figure 1: A hypothetical sequence of operations for the exploration of gene function within a biomedical text collection, where the functions of genes A, B, and C are known, and commonalities are sought to hypothesize the function of the unknown gene. The mapping operation imposes a rank ordering on the selected keywords. The final operation is a selection of only those documents that contain at least one of the top-ranked keywords and that contain mentions of all three known genes.

Hsinchen Chen, Andrea L. Houston, Robin R. Sewell, and Bruce R. Schatz. 1998. Internet browsing and searching: User evaluations of category map and concept space techniques. *Journal of the American Society for Information Sciences (JASIS)*, 49(7).

Kenneth W. Church and Mark Y. Liberman. 1991. A status report on the ACL/DCI. In *The Proceedings of the 7th Annual Conference of the UW Centre for the New OED and Text Research: Using Corpora*, pages 84–91, Oxford.

M. Craven, D. DiPasquo, D. Freitag, A. McCallum, T. Mitchell, K. Nigam, and S. Slattery. 1998. Learning to extract symbolic knowledge from the world wide web. In *Proceedings of AAAI*.

Douglass R. Cutting, Jan O. Pedersen, David Karger, and John W. Tukey. 1992. Scatter/Gather: A cluster-based approach to browsing large document collections. In *Proceedings of the 15th Annual International ACM/SIGIR Conference*, pages 318–329, Copenhagen, Denmark.

Ido Dagan, Ronen Feldman, and Haym Hirsh. 1996. Keyword-based browsing and analysis of large document sets. In *Proceedings of the Fifth Annual*

Symposium on Document Analysis and Information Retrieval (SDAIR), Las Vegas, NV.

Mark Derthick, John Kolojechick, and Steven F. Roth. 1997. An interactive visualization environment for data exploration. In *Proceedings of the Third Annual Conference on Knowledge Discovery and Data Mining (KDD)*, Newport Beach.

Usama Fayyad and Ramasamy Uthurusamy. 1999. Data mining and knowledge discovery in databases: Introduction to the special issue. *Communications of the ACM*, 39(11), November.

Usama Fayyad. 1997. Editorial. *Data Mining and Knowledge Discovery*, 1(1).

Ronen Feldman and Ido Dagan. 1995. KDT – knowledge discovery in texts. In *Proceedings of the First Annual Conference on Knowledge Discovery and Data Mining (KDD)*, Montreal.

Ronen Feldman, Will Klosgen, and Amir Zilberstein. 1997. Visualization techniques to explore data mining results for document collections. In *Proceedings of the Third Annual Conference on Knowledge Discovery and Data Mining (KDD)*, Newport Beach.

Christiane Fellbaum, editor. 1998. *WordNet: An*

- Electronic Lexical Database*. MIT Press.
- Marti A. Hearst. 1998. Automated discovery of wordnet relations. In Christiane Fellbaum, editor, *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, MA.
- David G. Hendry and David J. Harper. 1997. An informal information-seeking environment. *Journal of the American Society for Information Science*, 48(11):1036–1048.
- David C. Hoaglin, Frederick Mosteller, and John W. Tukey. 1983. *Understanding Robust and Exploratory Data Analysis*. John Wiley & Sons, Inc.
- Jon Kleinberg. 1998. Authoritative sources in a hyperlinked environment. In *Proceedings of the 9th ACM-SIAM Symposium on Discrete Algorithms*.
- Ray R. Larson. 1996. Bibliometrics of the world wide web: An exploratory analysis of the intellectual structure of cyberspace. In *ASIS '96: Proceedings of the 1996 Annual ASIS Meeting*.
- Xia Lin, Dagobert Soergel, and Gary Marchionini. 1991. A self-organizing semantic map for information retrieval. In *Proceedings of the 14th Annual International ACM/SIGIR Conference*, pages 262–269, Chicago.
- Christopher D. Manning. 1993. Automatic acquisition of a large subcategorization dictionary from corpora. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, pages 235–242, Columbus, OH.
- Francis Narin, Kimberly S. Hamilton, and Dominic Olivastro. 1997. The increasing linkage between us technology and public science. *Research Policy*, 26(3):317–330.
- Helen J. Peat and Peter Willett. 1991. The limitations of term co-occurrence data for query expansion in document retrieval systems. *JASIS*, 42(5):378–383.
- N. M. Ramadan, H. Halvorson, A. Vandelinde, and S.R. Levine. 1989. Low brain magnesium in migraine. *Headache*, 29(7):416–419.
- Earl Rennison. 1994. Galaxy of news: An approach to visualizing and understanding expansive news landscapes. In *Proceedings of UIST 94, ACM Symposium on User Interface Software and Technology*, pages 3–12, New York.
- Steven F. Roth, Mei C. Chuah, Stephan Kerpedjiev, John A. Kolojechick, and Peter Lucas. 1997. Towards an information visualization workspace: Combining multiple means of expression. *Human-Computer Interaction*, 12(1-2):131–185.
- Don R. Swanson and N. R. Smalheiser. 1994. Assessing a gap in the biomedical literature: Magnesium deficiency and neurologic disease. *Neuroscience Research Communications*, 15:1–9.
- Don R. Swanson and N. R. Smalheiser. 1997. An interactive system for finding complementary literatures: a stimulus to scientific discovery. *Artificial Intelligence*, 91:183–203.
- Don R. Swanson. 1987. Two medical literatures that are logically but not bibliographically connected. *JASIS*, 38(4):228–233.
- Don R. Swanson. 1991. Complementary structures in disjoint science literatures. In *Proceedings of the 14th Annual International ACM/SIGIR Conference*, pages 280–289.
- John W. Tukey. 1977. *Exploratory Data Analysis*. Addison-Wesley Publishing Company.
- Ellen M. Voorhees. 1994. Query expansion using lexical-semantic relations. In *Proceedings of the 17th Annual International ACM/SIGIR Conference*, pages 61–69, Dublin, Ireland.
- Michael G. Walker, Walter Volkmuth, Einat Sprinzak, David Hodgson, and Tod Klingler. 1998. Prostate cancer genes identified by genome-scale expression analysis. Technical Report (unnumbered), Incyte Pharmaceuticals, July.
- H. D. White and K. W. McCain. 1989. Bibliometrics. *Annual Review of Information Science and Technology*, 24:119–186.
- James A. Wise, James J. Thomas, Kelly Pennock, David Lantrip, Marc Pottier, and Anne Schur. 1995. Visualizing the non-visual: Spatial analysis and interaction with information from text documents. In *Proceedings of the Information Visualization Symposium 95*, pages 51–58. IEEE Computer Society Press.
- J. Xu and W. B. Croft. 1996. Query expansion using local and global document analysis. In *SIGIR '96: Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 4–11, Zurich.