SYMPOSIUM

Cognitive Science Meets Cognitive Engineering

Organizer & Chair

Richard Catrambone
School of Psychology
Georgia Institute of Technology
Atlanta, GA 30332-0170
richard.catrambone@psych.gatech.edu

Participants

Susan T. Dumais
Bellcore Information Science Research Group
445 South St. - rm 2L-371
Morristown, NJ 07960
std@bellcore.com

Jay Elkerton
NYNEX
500 Westchester Avenue
White Pains, NY 10604
je@nynexst.com

Bonnie E. John
Departments of Computer Science and Psychology and
The Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh PA 15213
Bonnie_John@cs.cmu.edu

Michael G. Shafto
Aerospace Human Factors Research Division
NASA-Ames Research Center
Moffett Field, CA 94035-1000
shafto@eos.arc.nasa.gov

Discussant

Peter Pirolli
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
pirolli@parc.xerox.com

Goal of Symposium

The intent of this symposium is to provide a forum for discussion about what it is like to try to "do" cognitive science in settings involving real products, real customers, and other constraints in which the bottom line is not necessarily to produce a pristine test of a theory but rather, to compete in the marketplace or perhaps save a life. Cognitive science is full of interesting and complex theories. These theories though are often difficult to translate into action in settings that require rapid responses to design and implementation questions. Thus, in the face of real-world constraints, cognitive science can seem of little use.

Cognitive scientists have to overcome the resistance to their ideas felt by interface designers, engineers, and other people who create artifacts upon which cognitive science could potentially have a large impact during various stages of development. In order to overcome this resistance, there must be demonstrations that cognitive scientists can be productive team members, that cognitive science can be used successfully under workplace constraints, and that it can develop tools that can be used by people who are not cognitive scientists.

Despite the challenging picture and daunting tasks outlined above, cognitive science *has* been successfully applied in work settings and the examples below show that it can make useful contributions.

Retrieving Information from External Sources

Susan T. Dumais
Bellcore

People have surprising difficulty in retrieving relevant information from external sources. Anyone who has looked for something in the index of a book, the Yellow Pages, or in computer databases has experienced this problem. You know that the information is there, but you can't think of the right words to find it! Although much research in cognitive psychology has been devoted to detailed descriptions of how people retrieve information from their own memories, much less work has addressed the question of retrieving information from external sources such as other people, books, or computer databases. One problem that is immediately evident in attempting to retrieve information from external sources is the tremendous mismatch between the searcher's language and that of the target information.

We call this the "vocabulary mismatch" problem. In brief, the words that a searcher uses to find an object are often different than the words used by the author or indexer to describe that object. For a wide variety of domains, we found that two people (e.g., a searcher and an indexer) generated the same main descriptor for objects only 10-20% of the time (Furnas, Landauer, Gomez & Dumais, 1987). Comparable poor agreement has been found in studies of interindexer consistency (Tarr & Borko, 1974) and in the generation of search terms by expert searchers (Fidel, 1985). This problem is not evident in traditional models of human memory because memory probes are in the same language as the memory representation - thus, by definition, there is no mismatch.

Vocabulary mismatch is a major barrier to the effective use of computer databases because most information retrieval tools are "word based." That is, they depend on matching words in users' queries with words in the objects of interest. Because of the tremendous diversity in the words that people use to describe essentially the same idea or concept (synonymy), relevant objects will be missed. Blair and Maron (1985), for example, found that up to 80% of known relevant materials were missed even by expert searchers. One way of thinking about the retrieval failures of word-based systems is that words are assumed to be independent. A query about "cars" is no more likely to retrieve an article about "automobiles" than it is one about "elephants," if neither author used precisely the term "cars." This is clearly untrue of human memory and seems undesirable in computer systems.

We have developed an automatic statistical method called Latent Semantic Indexing (LSI) which models the inter-relationships among terms and exploits them to improve retrieval. We use singular-value decomposition (SVD), a technique closely related to eigenvector decomposition and factor analysis, to model associative structure (see Deerwester, Dumais, Landauer, Furnas & Harshman, 1990 for details). A large matrix of terms to document association data is decomposed into a set of k=100-300 orthogonal factors that best approximate the original matrix. The idea is that these k largest independent linear components capture the major associational structure in the matrix and remove the noise or variability in word usage. The result of the analysis is a k-dimensional "semantic space" in which each term and document is represented. Terms which are used in similar contexts will be near each other in this space even if they never co-occur in a document. What this means from a user's perspective is that documents can be similar to a query even if they share no terms in common.

For a wide range of test databases LSI was 20-30% better than comparable word-based retrieval method in discriminating relevant from irrelevant documents (Deerwester et al., 1990; Dumais, 1991). Retrieval performance can be further improved by iterative retrieval and successive refinement of the original query through methods like relevance feedback. We have recently extended the LSI analysis to large diverse corpora (three gigabytes of text, one million articles from news wires, newspapers, computer science articles, U.S. patents, technical memos, etc.) with comparably positive results. The computational tractability of LSI was critical for success in this problem. Unlike manual thesaurus construction or most AI knowledge-based methods, the LSI statistical approach is completely automatic and thus widely applicable even to large, diverse collections. In addition, LSI has been successfully applied for a variety of other problems including: information filtering, cross language retrieval, retrieving noisy pen input, word sense disambiguation, finding synonyms, and text coherence analyses.

An analysis of the difficulties people have in retrieving information from external sources pointed to a "vocabulary mismatch" problem that has received little attention in cognitive science. We developed an associative learning model (LSI) which takes advantage of the inter-relationship among terms to allow people to more effective retrieving information from computer databases. We hope to examine the question of which "contexts" are most useful for learning, and to incorporate information about prior probabilities to further improve people's ability to retrieve information from external sources.

References

Blair, D.C. and Maron, M.E. (1985). An evaluation of retrieval effectiveness for a full-text document-retrieval system. *Communications of the ACM, 28(3)*, 289-299.

Deerwester, S., Dumais, S. T., Landauer, T. K., Furnas, G. W. and Harshman, R. A. (1990). Indexing by latent semantic analysis. *Journal of the Society for Information Science, 41(6)*, 391-407.

Dumais, S. T. (1991). Improving the retrieval of information from external sources. *Behavior Research Methods, Instruments and Computers, 23(2)*, 229-236.

Fidel, R. (1985). Individual variability in online searching behavior. In C.A. Parkhurst (Ed.). *Proceedings of the ASIS 48th Annual Meeting*, Vol. 22, 69-72.

Furnas, G.W., Landauer, T.K., Gomez, L.M. and Dumais, S.T. (1987). The vocabulary problem in human system communication. *Communications of the ACM, 30(11)*, 964-971.

Tarr, D. and Borko, H. (1974). Factors influencing inter-indexer consistency. In *Proceedings of the ASIS 37th Annual Meeting*, Vol. 11, 50-55.