Simulating text understanding for educational applications with Latent Semantic Analysis: Introduction to LSA.

Thomas K Landauer
University of Colorado at Boulder
And
Joseph Psotka[1]
US Army Research Institute

The following seven articles describe new educational tools that rely on a unique capability of computer programs to abstract knowledge relationships from vast quantities of text, and using this, to determine the similarity of knowledge expressed in two or more texts. The tools can make comparisons among instructional sources and expository student writing, and use the results to assess the quality of the student compositions, and guide and even tutor students to revise and improve their compositions. The pedagogical motive is simple: students need to learn - and learn how to learn - in a manner that allows them to express their knowledge in discourse, and they need to learn the important skills of verbal knowledge expression. This is particularly important for ill-structured knowledge domains, such as history, the social sciences, and everyday practical knowledge, where answers are not clearly prescribed, but remain ambiguous and constantly changing.

For too long education has been too much tied to assessment and tutorial feedback techniques that demand only choice among a narrowly prescribed set of alternative answers, the solution of formulaic problems, or the provision of a few words or discursively simplified phrases. It is not that such methods are bad or unreliable as assessment and motivational tools; it is that they are insufficient for teaching, learning and measuring the full range of study and knowledge-application skills that competent adults need. Moreover, reading and studying with only those assessments as goal-evaluation techniques encourages oversimplification, removes essential conflicts and sources of debate with Procrustean brute force, and invokes less effective efforts at deep and generalizable understanding. The reason for over-dependence on these methods is equally simple. The only alternative has been human scoring and commenting of student compositions, where the students have been encouraged to deal with issues that stretch their understanding. Expert human reading is undoubtedly still the best instrument for grading and critiquing these compositions, and it is the gold standard against which the methods employed in the following articles have been compared. However, it is too labor-intensive, demands too much professional effort, and so is too expensive to be used enough for a sufficient amount of practice and tutorial feedback to improve these difficult skills. Thus, the availability of a computational method to do some of the job, even if less well, offers a way to greatly increase time on task for acquiring these important skills.

[1] This work was supported by the US Army Research Institute and the Research and Advanced Concepts Office. The paper reflects the opinions of the authors and not official Army policy.

With the powerful new capabilities of these tools, the speed and flexibility of computer delivery makes it possible to engineer innovative interactive conversational learning activities that are responsive, stimulating, and enjoyable. Because of the greater-than-human capacity of a computer-based system to communicate with the simultaneous activities of many students doing many different things, the application of computer text understanding can potentially enhance interactive learning and performance environments in new and powerful ways. Indeed, we believe that the applications reported here only scratch the surface of what can be done for learning and assessment with computer text understanding, even with the first-generation technologies on which they currently depend.

<u>An informal description of Latent Semantic Analysis</u>. In all the following articles, the method used for computer text understanding-or, more precisely, computer-based comparison of the semantic content of texts-is a machine-learning technique called Latent Semantic Analysis (LSA). While LSA does not provide a model of all aspects of human language processing, it acquires and represents human word and passage meaning well enough to replace human text comprehension for many practical purposes. The formal LSA model relies on sophisticated mathematical and computer methods, computations that we believe correspond in effect to what the human mind does, that have been presented elsewhere (Deerwester at al., 1990; Landauer & Dumais, 1997, Landauer, Foltz and Laham, 1998). However, a non-technical description will show what LSA assumes and how it works.

Latent Semantic Analysis expresses an old idea about word meaning, that words occupy positions in a "semantic space" and their meaning is the relation of each word to all the others. A psychological version of the assumption underlying LSA is that people start by associating perceptual objects and experiences, including words, that are encountered near each other in time. However, the implicit cognitive processing involved goes far beyond piece-wise association. It takes all the myriad local relations and rearranges them to fit them together into a single consistent map, a semantic space that represents how each object, event or word is related to each other. LSA as mathematical computer model gets its experience by being fed a large body of electronic text. First, a computer version of associative learning establishes a quantitative relation between each unique word type (e.g., the word *computer* wherever it appears) and every paragraph in which it appears (say this one). This step is essential, and the particular way it is done matters. However, the real power of LSA comes from a subsequent process that takes these millions of separate links, and computes how all the myriad interlocking, consistent and inconsistent relations among them can best be reconciled. LSA then factors these links into a common semantic space in which each word and any passage in the language has its own place.

LSA does this in a way that is analogous to how cartographers once mapped the surface of the Earth. They started with rough estimates of distances between pairs of points: sightings from mountain tops, foot-travel days, sailing times. They put all of this piece-wise data together by placing the points-towns, river junctions, hills-onto a single picture in a way that preserved all the measurements as well as possible. If successful, they not only improved many of the distance estimates, but got an enormous added benefit in being able to read off the infinite number of point-to-point distances that had never been measured.

This kind of mapping works because of a simple fact of geometry. A structure of points in which each is connected to at least two others so as to form a closed collection of interlocking triangles is rigid, so any missing paths are strictly defined. However, this only works if you have assumed the correct shape - the right dimensionality - of the surface that the points are mapped onto. For example, mapping the spherical surface of the Earth onto a flat paper creates major distortions in relative size, and distance.

Latent Semantic Analysis in essence applies the same principle to mapping the meaning of words and passages into an abstract semantic space. In physical mapping, two towns could be quite close, but because they are on different sides of a canyon their direct distance may never have been measured. However, it is easily determined from the map. Similarly, two words may be quite alike in meaning but rarely used in the same context, for example because they are synonyms. However, when LSA places them in semantic space, it can bring them as close together (or far apart) as they deserve.

In LSA, the best number of dimensions usually turns out to be somewhere between 50 and 1,500, not the two or three of physical maps. We believe that the high dimensionality needed by LSA is a product of the way the brain is structured combined with the statistical structure of experience.

Dimensionality might be given additional meaning in non-mathematical form if it is considered in terms of your travels for work and pleasure in a city. Although the city can be fitted onto a two-dimensional map; your travels are more meaningfully described in more complex dimensions. For instance, when you visit your doctor, you may have to go physically to a place near your grocery store, but the two visits are worlds apart in purpose and in the average time interval between them. So, for instance, you are much more likely to shuttle between your kitchen and the grocery store, than between the kitchen and the doctor. The visit to the doctor is more closely related to a trip to the pharmacy, to another specialist, to a hospital, and so on. These trips may not only recur more often close to each other in time, but the objects found at these locations may also be moved along with you, for example, from the pharmacy to your medicine cabinet. In this way your travels can create many different groupings and more complex dimensions out of the simple two- dimensional structure of the physical space.

<u>A somewhat more formal, but still non-mathematical introduction to LSA</u>. The assumption underlying LSA is that similarities and differences in the meanings of words can be largely induced from similarities and differences in the discourse contexts in which they do and do not occur. In turn, similarities in the meanings of verbal passages can be largely induced from the combination (in the mathematical sense) of words that they contain. This assumption in turn implies that the usually dominant determiner of verbal meaning is the choice of words

and combinations of words into utterances, and that for many purposes the order of words in passages can be ignored in estimating similarity of meaning with little loss of accuracy. The question of whether these assumptions are correct has been asked, and answered in the affirmative, by evaluating the ability of models based on them to simulate a wide range of human verbal phenomena. These successes have encouraged us to hypothesize that the assumptions also characterize the mechanisms by which humans induce verbal meanings from discourse. We have found that important properties of language, such as the problem of "the poverty of the stimulus" in the phenomenally rapid acquisition of vocabulary by children, and many of the apparent indeterminacies of reference noted, for example, by Wittgenstein, are clarified by these assumptions in LSA (Landauer and Dumais, 1997; Landauer, Foltz and Laham, 1998; Kintsch, 1998, in press). However, the adequacy or completeness of LSA as a theory of human verbal meaning acquisition is not the focal issue in the papers in this special edition; we are interested primarily in the fact that a model built on these assumptions can acquire a good enough representation of meaning to usefully mimic human comprehension for application to instruction and assessment- although, in fact, the ability to do so does add evidence in support of the theory.

Turning these basic assumptions into a computational system requires a model in which word representations are a computable function of a set of observed linguistic contexts, and the representation of a linguistic context is a computable function of the set of words of which it consists, plus an effective algorithm by which such representations can be conjointly constructed from corpora of natural language usage. In the current instantiation of LSA, a linear function relates word and passage meaning representations, and a linear factorization technique is used to construct the representations from ordinary text as high-dimensional vectors. In LSA, a passage is the sum of the vectors of its contained words, and a word is the mean of the vectors of the passages in which it is observed. The constructive computation is the matrix-algebraic method Singular Value Decomposition (SVD) (Eckhart and Young, 1936; Golub and Van Loan, 1989), followed by dimension reduction. (It may be worth noting that other unsupervised learning techniques, for example principal components analysis or Hebbian neural-net training, might replace SVD in LSA, and that the current LSA algorithm includes pre- and post- SVD data transformations that have significant influences on its results.) In applying LSA, we have usually used the cosine of the angle - the similarity of direction - between vectors in the semantic space as the measure of similarity of meaning, and the length of vectors as a measure of the quantity or intensity of meaning in the given direction. Empirically, cosines and vector lengths have given the best results in simulations (Rehder et al., 1998; Laham, 199?); however, other measures on the vector space could be used. For details of LSA theory and implementation, see (Berry, 1992; Deerwester et al., 1990; Landauer & Dumais, 1997.)

The semantic space that LSA derives depends critically on the text corpus on which it is computed. We have tried to use training corpora that are representative in both size and content of the verbal sources from which the humans to be simulated would have learned to understand verbal meanings. This has proven to be possible and satisfactory in many cases. For example we have used a systematically sampled corpus representative of the text that one typical American student would have read up to first year of college to simulate a typical literate adult, and we have used textbooks for survey courses to represent an individual student's knowledge. Nonetheless, there are limitations to the corpus resources that are currently available and capable of being handled by the computations. For one thing, text collections do not include the voluminous exposure to oral sources that students have had. This is not as serious as might at first appear, because most rare words are encountered only in print, and both the common words of oral usage and the more specialized words used in educational and occupational settings tend to appear often enough in the text that students read to be adequately learned by LSA. Nonetheless, the difference undoubtedly limits the accuracy of our simulations to some extent. In addition, for many questions and applications, it would be useful to simulate the varying semantic spaces of a population of individuals rather than just one. An empirical corpus of tens or hundreds of individual life-time readings would be extremely hard to collect, and even a statistical sampling approximation would require a very much larger representative corpus than currently available. Another problem arises when there is a need to simulate not individuals' lexical knowledge, but that of a whole community of scholars or a whole cultural knowledge source. So, for example, if we wished to compare the content of a Ph.D. comprehensive exam essay on the history of antibiotics in medicine with all that could be known about the topic, we might want to base the LSA semantic space on a corpus many times as large and varied as any one person could possibly have read.

Such a corpus would also be extremely hard to collect, and impossible to analyze by SVD at high enough dimensionality with current hardware and software. Nonetheless, that LSA representations of human meaning can be good enough for many educational purposes has been demonstrated in an extensive set of studies. Applications to education described in the following articles, and others like them, are the most relevant. However, LSA derived word and passage meanings have also been tested in many other ways. For example, after training on appropriate text corpora, LSA: passed standard multiple choice vocabulary tests at literate-human levels; passed college course multiple choice exams at student levels; accurately mimicked psycholinguistic phenomena, including lexical decision priming, semantic category judgments, similarity based logical reasoning errors, and decision times for numerical inequality judgements; accurately measured the coherence and comprehensibility of text; and correctly reflected relations among synonyms, antonyms, singulars and plurals, simple and compound words, and multiple word senses (Foltz, Kintsch & Landauer, 1998; Kintsch, 1998, in press; Landauer, 1999; Landauer & Dumais, 1997; Landauer, Foltz and Laham, 1998).

We have also conducted some experiments on educational applications that are not covered in the papers in this volume. For example, in one that led to what we call the Goldilocks Principle, college and medical students wrote short essays on the anatomy and function of the heart. We then randomly assigned texts at different levels of sophistication, and tested what was learned. Either conceptually too sophisticated or too simple texts, as measured by LSA relations between essays and texts, produced about one standard deviation less learning than "just right" texts (Wolfe et al., 1998; Rehder et al., 1998).

We believe that the educational applications of LSA to date make a strong showing of promise, and that many more innovative opportunities await the ingenuity of educational technologists. While we must caution that the application of LSA requires deep familiarity with its underpinnings, and its appropriate, correct and practical implementation - as well as its limitations -we invite others to experiment with it. To this end, we maintain a web site at http://LSA.colorado.edu. The site has further information, references, and demonstrations, as well as facilities with which similarities between words and passages based on semantic spaces derived from several corpora (none or more of which may be appropriate for a particular case) can be obtained for trial use.

Overview of articles. The articles in the present volume provide a cross section of experimental uses of LSA: for formative and summative assessment of student compositions for middle school, high school, and undergraduates in college; a Web-based learning environment that provides feedback to students about the adequacy of summaries that they are writing; tutors for hardware, operating systems, and computer literacy; a system for matching jobs to personnel with the right knowledge and experience; and a system for comparing the mental models of experts and novices. Each article provides empirical findings and exploratory suggestions for additional uses and implementations.

The paper by Eileen Kintsch and others describes the *Summary Street* and *State the Essence* systems that use LSA iteratively for evaluation and componential feedback. *Summary Street* has been created for students who are studying difficult, new material. it has been tested so far with sixth-grade students learning to write summaries of materials they read, to help and motivate them to rewrite their compositions in guided iterative revision cycles. The feedback has shifted from providing students a rating of how well they have covered the central points in a text – whether they have “stated the essence” – to a presentation format that encourages rewriting. In both versions the feedback consists in pointing out spelling errors and informing the writer whether or not the summary covers the main topics and is the proper length. In addition, writers can ask for feedback on possible redundant and irrelevant sentences, but rather than listing problematic sections and sentences, the new version of the summarization tool displays some of the feedback in a graphic form and offers better guidance for dealing with the problems.

The goals for the summarization tool were to give students extended practice in writing and revising summaries while reducing teachers' load of reviewing and grading successive drafts, and to help motivate students to work hard and independently by providing immediate and individualized feedback on how to revise their writing. Although empirical data to demonstrate that the intervention was effective are sparse, the overall convergence of evidence was supportive. For these summaries, once again LSA was demonstrated to be as effective in assessing the quality of the summaries as a teacher: the correlation between the teacher grade and the LSA cosine was $r = 0.64$. The correlation between a second scorer and the teacher was r = 0.69. Thus, LSA scores are quite comparable to how an experienced teacher rates these summaries. Furthermore, LSA does almost as well as teachers at determining the source of knowledge for a given sentence, a fact that can be useful in designing future versions of the system. However, the most dramatic and important finding is also the easiest to confirm: that students, with this environment's help, persisted in revising their writing many more times than those who only worked with a word processor, and were able to do so without the time-consuming intervention of teachers. This possibility of more extensive opportunities for students to create and revise summaries is the banner headline to promise, not immediate mastery of these very difficult skills, but a long term and sustained improvement that could provide significant advantages to all students.

Peter W. Foltz, Sara Caukwell and Scott Kendall continue the theme of supporting writing and revision but move it to an undergraduate level. They discuss a classroom trial of a web-based essay grader and critiquer that demonstrated the promise of increasing the writing and particularly the number of revisions students made in an undergraduate psycholinguistics course taught in 1997 and 1998. Latent Semantic Analysis can be used to provide both overall scores of essays and more specific feedback about what individual pieces of information are missing from an essay. This second feature is particularly useful for encouraging revisions because it can be used to inform students where to look for additional information in their texts or notes, or to get them to think more completely about the missing concepts before writing revisions. The program encouraged students to submit essays and receive estimates of the overall grade of the essay, and to take advantage of feedback about information that appeared to be missing from their essays. Students could revise their essays and resubmit until they were satisfied with their grades. As anticipated, the results indicated that average scores improved considerably, by a whole letter grade, during the course of the revisions.

Arthur C. Graesser, Peter Wiemer-Hastings, Katja Wiemer-Hastings, Derek Harter, Natalie Person , and the Tutoring Research Group at the University of Memphis describe their work on "AutoTutor", an automated computer tutor that assists students in learning about hardware, operating systems, and the Internet in an introductory computer literacy course. Latent semantic analysis (LSA) is a component of the mechanism that evaluates the quality of student contributions in a tutorial dialog. LSA's evaluations of college students' answers to deep reasoning questions are equivalent to the evaluations provided by intermediate experts of computer literacy, but not as high as more accomplished experts in computer science.

LSA is also capable of discriminating different classes of student ability (good, vague, erroneous, versus non-responding students), by matching their answers to anticipated good answers stored in English, as opposed to LISP code or other structured code. Because it is stored in English, the anticipated answers can be authored by a teacher or other individual who is not an expert programmer. One of the salient benefits of using LSA for the pattern match operations is that the content of a curriculum script can be written in English. Overall, these evaluations of LSA's capabilities in this report reflect its very useful capability of assessing students' knowledge of complex real world issues.

The article by Peter Wiemer-Hastings and Arthur C. Graesser describes a system for editing and revising compositions that uses LSA in a unique way. In combination with other techniques that focus on grammatical and syntactic issues, LSA provides the environment, called Select-a-Kibitzer, with the power to address a wide-range of meaning-oriented composition issues, including coherence, purpose, topic, and overall quality. Providing feedback about the semantic content of a composition requires great skill, and the evidence that LSA can perform this high level function relatively satisfactorily at a college level, suggests great promise for this technology. One interesting technique pioneered in this work is to determine the best prototype sentences for each topic by choosing the text with the highest average cosine with the other texts ON the topic. With these sets of topics and prototype sentences in hand, Select-a-Kibitzer can give feedback in a variety of ways. This diversity of responding makes it more flexible and interesting, and effective.

The article by Darrell Laham, Winston Bennett, and Thomas Landauer describes a system for analysis of course and curriculum content, matching individuals to tasks and training materials, and recombining training segments for new tasks. Military organizations are increasingly faced with rapid changes in technology and missions, and need constantly changing mixes of competencies and skill. Assembling personnel with the right knowledge and experience for a task is especially difficult when there are few experts, unfamiliar devices, redefined goals, and short lead times for training and deployment. Large civilian organizations face similar challenges in adapting to international competition, new technologies, and organizational re-alignments. When too few adequately trained personnel are available for suddenly critical tasks, organizations need the ability (a) to identify existing personnel who could perform the task with the least training, and (b) to create new training courses quickly by assembling components of old ones. New LSA-based agent software helps to identify required job knowledge, determine which members of the workforce have that knowledge, pinpoint needed retraining content, and maximize training and retraining efficiency. The LSA-based technology extracts, represents and matches information about people, occupations, and experience contained in textual databases.

To demonstrate and evaluate the system, Laham et al. analyzed the tasks and personnel in three Air Force occupations. They measured the match of each airman to each task and estimated how well each airman could replace another. They also demonstrated the potential to match knowledge subcomponents needed for new systems with ones contained in training materials and those possessed by individual airmen. The research provides results that demonstrate that LSA can successfully characterize tasks, occupations and personnel and measure the overlap in content between instructional courses covering the full range of tasks performed in many different occupations. This research shows the potential for LSA-based methods to identify ways in which occupations might be reorganized to increase training efficiency, improve division of labor efficiencies, or redefine specialties to produce personnel capable of a wider set of tasks and easier reassignment. The natural language query design intrinsic to LSA eliminates the known problems inherent in keyword matching of field-restricted databases.

Finally, the paper by Jared T. Freeman, Bryan T. Thompson, and Marvin S. Cohen describes how Latent Semantic Indexing (LSI) was used to model expert knowledge of a complex domain using very short texts written by journeymen and experts alone, rather than from a broad knowledge base of published narratives in existing books. The analysis focused on arguments made by Navy officers as they defended and critiqued assessments of events in an Anti-Air Warfare scenario. The modeling technique consisted of three steps: LSI was used to create a highly dimensioned argument space; the LSI factor space was reduced using Multi - Dimensional Scaling (MDS) to five dimensions; and these dimensions were then analysed to define a "mental model" of AAW issues in the scenario. Experts and journeymen reliably differed in their frequency of use of dimensions of the model, thus providing some evidence of the model's validity. Automated classification of the arguments into the model elements achieved 84% accuracy, indicating that LSI may help us to diagnose individuals' mental models. Overall, the findings suggest that LSI may prove useful in knowledge engineering to create instructional materials in domains where few, authoritative written materials exist, and it may help us to assess student's mental models in those domains. LSI may be a useful tool for automating aspects of knowledge engineering, modeling expertise, and diagnosing knowledge deficiencies.

<u>LSA and the Creation of Instructional Knowledge Bases</u>. The promise of Intelligent Tutoring Systems has been promulgated for several decades now (e.g. Shute and Psotka, 1996) but has consistently failed to be realized, largely because of the enormous efforts required to code the various knowledge bases for subject matter content, instructional strategies, student assessment, misconceptions, etc. that might be required for a good tutor. Although LSA has not cracked this nut completely by any means, it offers for the first time a new approach, another avenue, for achieving these goals. Instead of hand tailoring simulation objects or first order predicates, it appears to be good enough to have people simply write their best textual description of the problem and its solutions, including those inevitable admonitions about what NOT to do. With a sufficient number of these protocols from a range of novices to experts (cf. Freeman, et. al, 2000) , it should be possible for LSA to construct a consensual representation for each of the required knowledge bases. LSA as a knowledge extraction tool for engineering ITS knowledge bases has only been explored in a few preliminary ways (cf. Graesser et al., 2000) but compared to traditional knowledge engineering, it is a much simpler task to ask people to write a few paragraphs on aspects of the topic of their expertise and interest.

<u>A closing note on LSA's psychological realism</u>. For an audience interested in learning, a final comment on the status of LSA as a theory of human learning and cognition seems in order. Certainly, LSA is not a complete, or completely accurate, model of human language acquisition and understanding. It lacks input from any source other than text on which to learn word meanings, and as a

simulation system ignores the dynamic modifications of meaning that are carried by the order of words-their sentential grammar and syntax (but see Kintsch, in press, for progress on this front). However, even with these handicaps, LSA computes similarities among words and passages that usually correspond closely to those exhibited by people.

Now a critic might say, with some justification, that all LSA is doing in measuring the similarity of two utterances is assessing the degree to which they use the same words or words of similar meaning. This might seem trivial, but it makes an extremely important point. To represent the meaning of a passage for a human, it appears very nearly sufficient to have a good way to represent word meanings mathematically, then simply add them together ignoring their order. Who would have thought it? Then again, would it really be such a surprise to find that verbal meaning is carried primarily in the combinations of words one chooses?

The ability to assess the similarity of meaning of words and passages computationally in a manner that mirrors human comprehension to a high level of approximation offers unlimited possibilities for useful applications. The following articles exploit this opportunity in the creation of novel educational tools that do things that have not been done before, or at least not done with nearly the level of ease, accuracy, and generality that LSA provides. However, our point here is to stress that there are scientific implications in the other direction. That because these tools can succeed in carrying out educational activities that if done by a person would be described as dependent on understanding language, the system too must be doing a good deal of language understanding. And since we don't know how humans understand language, a computational model that can do much of the job provides useful clues as to how we might be doing it.

REFERENCES

Berry, M. W. (1992). Large scale singular value computations. International Journal of Supercomputer Applications, 6(1), 13-49.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing By Latent Semantic Analysis. Journal of the American Society For Information Science, 41(6), 391-407.

Eckhart, C., and Young, G. (1936) The approximation of one matrix by another of lower rank. Psychometrica 1:211-218.

Foltz, Peter W., Caukwell, S. and Kendall, S. Supporting content-based feedback in online writing evaluation with LSA. Interactive Learning Environments. This volume.

Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998) The measurement of textual coherence with Latent Semantic Analysis. Discourse Processes, 25, 285-308.

Graesser, A. C., Wiemer-Hastings, P., Wiemer-Hastings, K., Harter, D., Person, N., and the Tutoring Research Group. (in press). Using Latent Semantic Analysis to Evaluate the Contributions of Students in AutoTutor. Interactive Learning Environments. This volume.

Golub, G. H., and Van Loan, C. F. (1989) Matrix Computations. The Johns Hopkins University Press, 2nd edition.

Kintsch, E., Steinhart, D., Lamb, R., Matthews, C. & the LSA Research Group (1999). Developing summarization skills through the use of LSA-based feedback. Interactive Learning Environments. This volume.

Kintsch, W. (1998) Comprehension: A paradigm for cognition. Cambridge, England. Cambridge University Press.

Kintsch, W. (in press) A computational theory of metaphor comprehension. Psychonomic Bulletin and Review.

Laham, D., Bennett, W., and Landauer, T.K. (in press) An LSA-based software tool for matching jobs, people, and instruction. Interactive Learning Environments. This volume.

Landauer, T, K, (1998) Learning and representing meaning: The Latent Semantic Analysis theory. Current directions in Psychological Science. 7, 161-164

Landauer, T. K., and Dumais, S. T. (1997) A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. Psychological Review 104 (2):211-240.

Landauer, T. K., Foltz, P. & Laham, R. D. (1998) An introduction to Latent semantic analysis. Discourse Processes, 25, 259-284.

Rehder, B., Schreiner, M. E., Wolfe, M. B. W., Laham, D., Landauer, T. K. & Kintsch, W. (1998) Using Latent Semantic Analysis to assess knowledge: Some technical considerations. Discourse Processes, 25, 337-354.

Shute, V. J., and Psotka, J. (1996) Intelligent Tutoring Systems: Past, Present, and Future. In D. Jonassen, (Ed.), **Handbook of Research In Educational Communications and Technology**. New York: Simon and Schuster Macmillan, pp. 570 - 600.

Stahl, G., Kintsch, E., Kintsch, W., Laham, D., Landauer, T. K., DePaula, R., Schreiner, M.E. & Steinhart, D. (1999) An LSA Summarization Environment: State the Essence (In preparation).

Wolfe, M. B. W., Schreiner, M. E., Rehder, B., Laham, D., Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998) Learning from text: Matching Readers and texts by Latent Semantic Analysis. Discourse Processes, 25, 309-336.