LATENT SEMANTIC ANALYSIS FOR CAREER FIELD ANALYSIS AND INFORMATION OPERATIONS

Darrell Laham
Knowledge Analysis Technologies
Boulder, Colorado

Winston Bennett
Air Force Research Laboratory
Mesa, Arizona

Marcia Derr
Knowledge Analysis Technologies
Boulder, Colorado

ABSTRACT

This paper reviews two current Air Force Research Laboratory / Human Effectiveness Directorate (AFRL/HEA) efforts that are maturing Latent Semantic Analysis (LSA) tools for the Air Force. The first effort is developing new LSA-based agent software that helps decision makers to identify required job knowledge, determine which members of the workforce have the knowledge, pinpoint needed retraining content, and maximize training and retraining efficiency. Modern organizations are increasingly faced with rapid changes in technology and missions and need constantly changing mixes of competencies and skills. 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. LSA is being used to analyze course content and materials from current training pipelines and to identify appropriate places in alternative structures where that content can be reused. This saves time for training developers since the preexisting content has already been validated as a part of its earlier application.

AFRL/HEA's second research effort involves a demonstration of a combined speech-to-text and LSA-based software agent for embedding automatic, continuous, and cumulative analysis of verbal interactions in individual and team operational environments. The agent will systematically parse and evaluate verbal communication to identify critical information and content required of many of today's AF operators. LSA is promising new technology that has significant potential for assisting operators in the performance of their tasks because it can "listen" and in almost real-time evaluate free-form verbal communication from a variety of sources and match content to stored language dictionaries. One application of this technology being explored is tracking and scoring the tactical communications that occur between the members of a four-ship air combat flight and their weapons director to identify areas of training need and as an additional tool for assessing the efficacy of DMT scenarios and missions.

ABOUT THE AUTHORS

Darrell Laham is a cofounder of Knowledge Analysis Technologies and serves as its Chief Technology Officer. He holds a Ph.D. in Psychology and Cognitive Science from the University of Colorado at Boulder. Dr. Laham has been the chief architect and project leader on several previous and ongoing projects in education, training, and workforce management for the Army, Navy, Air Force, DARPA, NASA, and NIST, as well as commercial applications of KAT's educational technologies.

Winston Bennett Jr. is a Senior Research Psychologist and team leader for training systems technology and performance assessment at the Air Force Research Laboratory, Human Effectiveness Directorate, Warfighter Training Research Division in Mesa, AZ. He received his Ph.D. in Industrial/Organizational Psychology from Texas A&M University in 1995. Dr. Bennett has been an active member of I/ITSEC since 1994 and has served on two conference committees.

Marcia Derr, who holds a Ph.D. in Computer Science from Stanford University, has published in the areas of natural language processing and database technology. Prior to her current consulting position with Knowledge Analysis Technologies, she worked on a variety of exploratory software projects at Bell Laboratories and at U S WEST Advanced Technologies.

Projects funded by SBIR Contract Number F41624-99-C-5003 Air Force Research Laboratory, Mesa Branch.

LATENT SEMANTIC ANALYSIS FOR CAREER FIELD ANALYSIS AND INFORMATION OPERATIONS

Darrell Laham
Knowledge Analysis Technologies
Boulder, Colorado

Winston Bennett
Air Force Research Laboratory
Mesa, Arizona

Marcia Derr
Knowledge Analysis Technologies
Boulder, Colorado

INTRODUCTION

There are currently two major efforts underway that are maturing Latent Semantic Analysis (LSA) tools for the Air Force.

CareerMap—Matching People to Training and Jobs

The first effort is developing new LSA-based agent software, called CareerMap, that helps decision makers to identify required job knowledge, determine which members of the workforce have the knowledge, pinpoint needed retraining content, and maximize training and retraining efficiency. Air Force Research Laboratory / Human Effectiveness Directorate's (AFRL/HEA) testbed application is formal schoolhouse training for Navigators and Electronic Weapons Officers as they transition into Air Combat Officers. LSA is being used in conjunction with other AFRL/HEA-developed course content modeling methods to evaluate alternative NAV-EWO training pipelines. LSA is being used to analyze course content and materials that are being used in the current pipeline and to identify appropriate places in the alternative structures where that content can be reused. This saves time for training developers since the preexisting content has already been validated as a part of its earlier application. Moreover, gaps in the content for the new training structure become readily apparent with this type of analysis.

Automated Mission Communications Analysis

AFRL/HEA's second research effort is being done in collaboration with the Crew Systems Interface Division (AFRL/HEC) at Wright Patterson AFB and involves a demonstration of a combined speech-to-text and LSA-based intelligent software agent for embedding automatic, continuous, and cumulative analysis of verbal interactions in individual and team operational environments. At the present time, it is impossible to systematically parse and evaluate verbal communication to identify critical information and content required of many of today's AF operators. LSA is promising new technology that has significant potential for assisting operators in the performance of their tasks because it can "listen" and in almost real-time evaluate free-form verbal communication from a variety of sources and match content to stored language dictionaries.

One application of this technology being explored is tracking and scoring the tactical communications that occur between the members of a four-ship air combat flight and their weapons director to identify areas of training need and as an additional tool for assessing the efficacy of DMT scenarios and missions. Similarly, we envision the combined technologies being useful in providing an embedded assistant to help track and evaluate incoming communication and to highlight or otherwise "flag" pertinent information and changes in content that may be of importance to operators and other personnel.

LATENT SEMANTIC ANALYSIS (LSA)

LSA is a method for automatically extracting and representing knowledge in massive databases of relevant electronic text (Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990). It was developed through ten years of basic and applied research supported by Bell Communications Research, DARPA, ONR, ARI, NASA, AFRL, the McDonnell Foundation and others. LSA has been extensively validated in both controlled experiments and field tests (Landauer & Dumais, 1997; Landauer, Foltz, and Laham, 1998; Landauer, 1998).

Automated Analysis of Meaning

As a psychological theory of the acquisition, induction, and representation of knowledge, LSA research has provided new insights on how people learn the meanings of words. LSA is instantiated as a mathematical system for computational modeling of cognitive processes. As a tool, LSA is used as an artificial intelligence (machine learning) system useful in various educational and industrial applications.

LSA provides a method for determining the similarity of meaning of words and passages by analysis of large text corpora such as domain knowledge libraries, writing samples, e-mail files, course materials, and job and training historical records. After processing a large sample of machine-readable language, LSA represents the words used in it, and any set of these words such as a sentence, paragraph, or essay either taken from the original corpus or new, as points in a very high (e.g. 300) dimensional semantic space . LSA is closely related to neural net models, but is based on singular value decomposition, a mathematical matrix decomposition technique closely akin to factor analysis that is applicable to text corpora approaching the volume of relevant language experienced by people.

Word and passage meaning representations derived by LSA have been found capable of simulating a variety of human cognitive phenomena, ranging from developmental acquisition of recognition vocabulary to word-categorization, sentence-word semantic priming, discourse comprehension, and judgments of essay quality. In many applications LSA judgments of similarity agree well with human judgments.

CAREERMAP

The Problem

Modern 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. 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.

Current solution methods for such problems require large investments of expert labor and are often either unacceptably slow or insufficiently effective. For example, determining whether a particular person's background for a particular assignment requires some time-consuming training component, or allows it to be omitted, would probably necessitate intensive general training research as well as extensive individual questioning. Creating an efficiently individualized training course would be just as difficult. While a well developed engineering art, in the best case course design takes many months of specialized analysis and trial.

Thus, more effective methods are desired for characterizing, locating, and training personnel who can optimally perform the set of duties required by any new mission. This calls for information technologies that can:

- **(a)** represent knowledge and skills,
- **(b)** identify people with all or parts of the knowledge and task experience required by a mission—wherever and in whatever occupation they are currently,
- **(c)** determine precisely what, if any, retraining each person needs in order to perform which new duties,
- **(d)** reduce the effort required to create new training programs, and
- **(e)** minimize the time required for training and retraining.

The objective of this research was to develop and test the practical capability of LSA in application to these problems. The new personnel data mining application of LSA exploits the explicit and implicit knowledge that already exists in extensive textual computer files of systems documentation, training and test materials, task analyses, and service records (Laham, Bennett, and Landauer, 2000).

Description of LSA Representations

The biggest advantage of LSA knowledge representation for the present purpose is that different types of data objects (e.g. occupations, job tasks, personnel, training materials) can all exist as vectors within the same semantic space and can therefore be directly compared to each other in meaningful ways. People can easily make holistic judgments of similarity between a task to be performed and a set of people who might be called upon to perform the task. However, the structure in which this information is usually stored in computer files, i.e. in relational databases, has precluded the possibility of automated judgments of this sort.

Traditional database structures are very brittle in that search and retrieval are overly dependent on specific data field choices (e.g. zip code field, job title field) and on exact keyword matching. While in many cases exact matching on highly structured data is desirable (e.g. find the names of all people in zip code 30405), in many other cases the choices can be overly restrictive and/or ambiguous (find all the people who list their job title as 'Doctor'). In the latter example, those people who listed their title as Medical Doctor, Physician, Surgeon, General Practice Doctor, and other medical specialties would not match the query and would be inappropriately excluded.

The LSA Solution

Figure 1 shows a 2-dimensional representation of 2 LSA objects, a Job and a Candidate. In actual LSA representations, it requires 100-500 orthogonal dimensions to characterize an object this 2-D representation is for illustrative purposes only. Unlike traditional factor analysis, where the dimensions have been interpreted and named, LSA dimensions are not *in and of themselves* meaningful. Both objects are seen as points in the Semantic Space having some score on the X dimension and an independent score on the Y dimension. In most of the work reported in this paper, each of the objects has scores for 300 orthogonal dimensions, thus each object is represented as a vector of 300 numbers, rather than a vector of 2 numbers. In actual LSA, the meaning of an object is determined when the full set of dimension scores is used in the comparison of the object to other objects. In most cases LSA uses the cosine of the angle between objects as the measure of similarity.

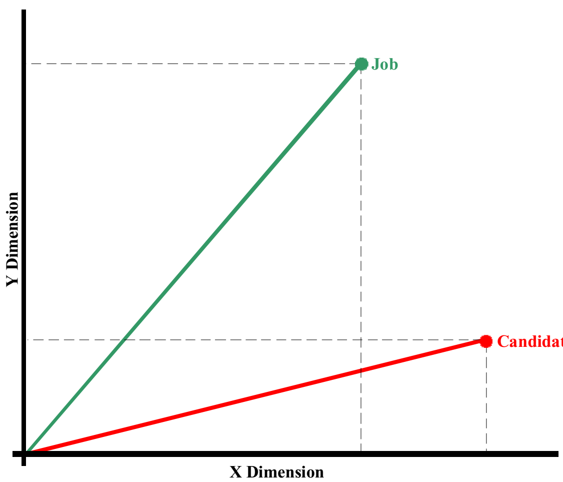

Figure 1. A 2-D representation of 2 LSA objects.

In Figure 2 the capacity to make similarity comparisons between different data objects is illustrated. To determine which of two candidates is best suited for a job, the LSA system would use the smaller of the two angles 1 and 2. Cosines range between 1 and -1 for all possible angles an angle of 0 degrees has a cosine of 1 (the vectors lie on top of each other), an angle of 90 degrees has a cosine of 0 and an angle of 180 degrees has a cosine of -1. The smaller the angle, the higher the cosine, and the more similar the two objects are considered. In the Figure 2 case, Candidate 1 is a better choice for the Job.

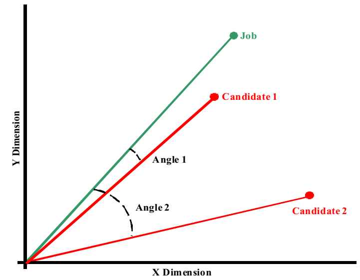

Figure 2. Two candidates for a job.

The CareerMap Application

CareerMap (http://www.careermap.org) is a demonstration of the Intelligent Search Agent which can aid in the performance of mission critical training management tasks. It is internet based, so it works with any browser. The demonstration system is populated with sample data from the Air Force, Department of Labor Occupational Network (O*Net), job listings from the Office of Personnel Management and resumes from Yahoo and other sources. In the demonstration system, one can create a text query to retrieve and rank any type of data object known to the system (e.g. AFSC descriptions or training records). (See Figure 3).

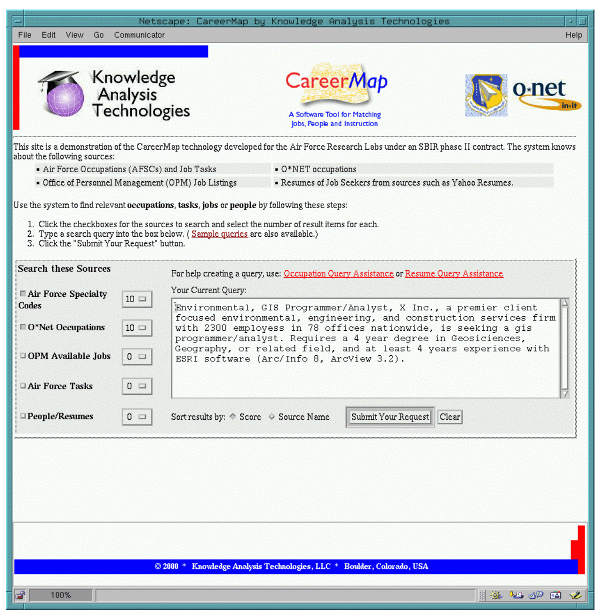

Figure 3. The CareerMap Application

The query will return the most similar objects in the system. The user can review the results and, if desired, expand the query using retrieved data (Figure 4).

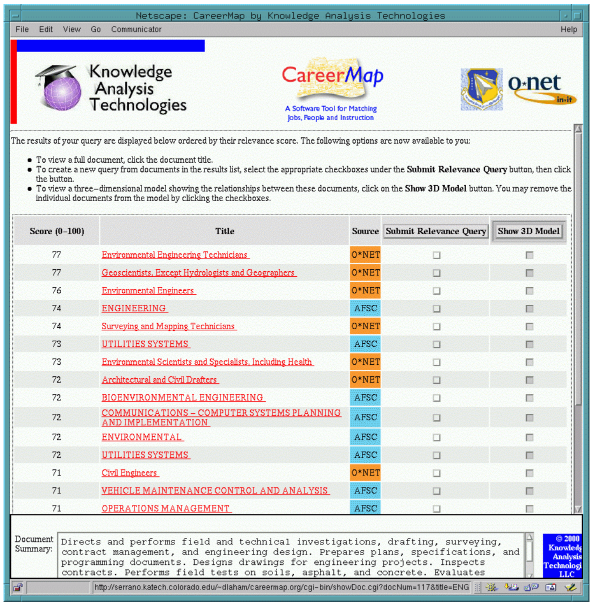

Figure 4. CareerMap Object Return List

An interesting data analysis feature of CareerMap employs visualization tools to help the user discover the relationships between the objects returned by the search (see Figure 5). The return list of objects are those which are nearest the query in semantic space. The visualization tools allows the interrelationships between the objects to be seen.

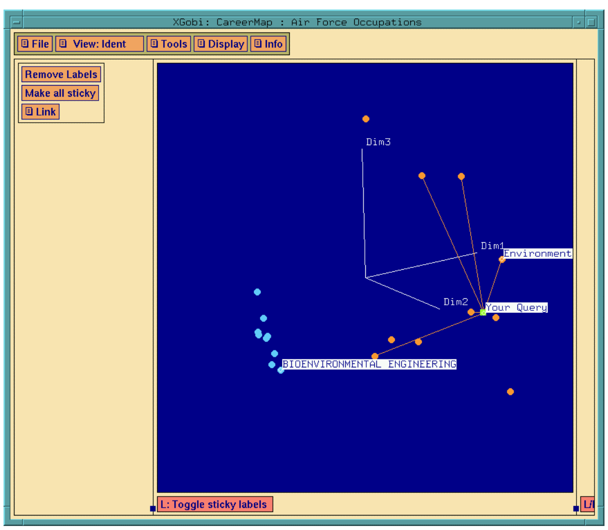

Figure 5. CareerMap 3D data visualization tools

Potential Applications of the Method

Job placement or occupation assignment. Practical applications to job assignment were most directly illustrated by the research just reported. The simplest case is direct replacement of one airman with another. For this, a query takes the form of the to-be-replaced airman's identification number, and the *k* most similar airmen known to the system (potentially all those in the Air Force plus others where relevant) are returned and listed in terms of their overall task-experience pattern—the closeness of their points in the joint semantic space representing tasks, occupations and airmen. Their complete service records can then be displayed and compared. If it is desired to add a new member to a work group, the descriptions of those tasks that are most in need of additional help can be entered as the query and the system will list in order those airmen whose total experience is most like the new job requirements. Note that in performing this match, LSA goes beyond simply counting the number of tasks in common between the wanted list and the service record, instead factoring in previous experience (and, later, training) in occupations and tasks that are similar but not identical to those in need of performance. Thus, it would be quite possible, in the absence of any airman who had done any of the prescribed tasks, to nevertheless find one or more candidates who had done similar work, the estimate of similarity having been automatically induced by LSA from the entire corpus of data without human intervention.

The technique could be used to add people to perform new jobs, by adding to the query a free-form description of the tasks involved. Because LSA captures semantic and conceptual similarity of verbal expressions, it will correctly match *ad hoc* task descriptions with official task definitions and job descriptions. The system can also form a representation of the overall mix of tasks required by a group by combining representations of the knowledge possessed by all its present members. In case of downsizing, the system would make it possible to find a set of personnel to transfer out of a group that would either leave it most like its previous composition, or desirably modified, again without relying on a crude counting operation or intuition.

The opportunity and manner of application for selecting airmen for missions, for example expeditionary war-fighting missions with unique challenges, is relatively straightforward. Given a careful verbal description of the mission, including all the tasks to be performed, the equipment, weapons, devices, procedures, numbers of airmen needed in each role, and perhaps even factors such as locale, terrain and likely weather and other challenges, the LSA matching technique would rank airmen for suitability to each task on the basis of the totality of their previous task and occupational experience, along with, if available, relevant (as determined by LSA) test scores and performance ratings.

Curriculum overlap analyses. The Air Force (like other military and civilian organizations) offers hundreds of specialty training courses, many of which overlap substantially in content, many of which may contain content no longer relevant to tasks currently in demand, and some which are missing content made desirable by changes in technology, missions or staffing. In many cases it would be desirable to combine, condense, or modify courses. Teaching unnecessary numbers of courses or redundant components in multiple courses is expensive in instructional staff and facilities, and even more expensive in wasted student time and resources. Teaching material that is sub-optimally matched to work requirements, either by being superfluous, redundant, or by failing to equip airmen with the best skill sets for all the tasks it would be desirable for them to be able to perform, is probably even more expensive in the long run.

To rationalize the content and organization of content for multiple training programs, a method is needed by which the overlap in course content can be easily assessed. Presently such analyses are performed, if at all, by highly labor intensive efforts by subject matter experts and training specialists. We have already demonstrated that LSA can do this kind of analysis automatically to a quite useful degree. Our studies were based on analysis only of course examination items, but appeared to give a great deal of useful information about course overlap.

LSA can also measure the overlap between course content and the full range of tasks performed in many different occupations. Information from such analyses will suggest where the training needed for different occupations overlaps and might be combined, where training is lacking, point to components that may not actually be needed at all, and, in some instances, suggest ways in which occupations might be restructured to increase training efficiency. LSA methods will not solve these problems completely, but we believe they can offer highly useful information for planners that is currently unobtainable or prohibitively expensive.

*Just-in-time training materials***.** In brief, the way in which we envision that LSA would be employed in helping to rapidly create new targeted training programs might be as follows. The component knowledge needed and tasks to be performed for a new device, system, or procedure would be carefully described by designers and relevant subject matter experts. LSA would determine the degree of match of each component to a wide range of tasks performed in the Air Force and to every paragraph in every possibly relevant training or operations manual. Tasks and paragraphs would be ranked by estimated relevance to the new system, and the LSA similarity of each paragraph to each task determined. In the quickest and dirtiest version, a custom retraining document for each candidate could initially be compiled from paragraphs highly relevant to the new system that are not highly similar to tasks the candidate has previously performed. In the case of urgent need for a small number of trainees, a subject matter or training expert could then edit each version. In case of need for large numbers and more available time, the collection of paragraphs could be crafted into a simple computer-based training program with branching to permit trainees to skip parts they already know.

MISSION COMMUNICATIONS ANALYSIS

Introduction

In the second ongoing AFRL/HEA application, LSA is being used to assess communication in Distributed Mission Training. The goal here is to develop and implement an LSA-based "Automated Communications Analysis" pipeline for performance assessment of mission communications applicable to both simulated and live Distributed Mission Training. The analysis of communications will be used to inform instructors and students for feedback both during mission performance and in related After Action Briefings. (Figure 6).

Team Mission Communications
Speech Recognition speech-to-text
LSA analyses and performance scores
After action briefings & Performance feedback

Figure 6. Automated Communications Analysis Pipeline

As a proof of concept, LSA was successfully able to predict team performance in a simulated UAV task environment (Kiekel, Cooke, Foltz, Gorman, and Martin, 2001) based only on communications transcripts. Using human transcriptions of 67 team missions in the UAV environment, LSA predicts objective team performance scores at a very high level of reliability (LSA alone, $r = 0.74$; LSA combined with additional text analysis measures, $r = 0.85$) The Team Performance Score used as the criterion measure is a composite of objective measures including the amount of fuel and film used, the number and type of photographic errors, route deviations, time spent in warning and alarm states, unvisited waypoints and violations in route rules. In this analysis, LSA compares the content of a mission transcript of unknown performance quality to those of known performance quality to generate the LSA Performance scores. A weighted average of the objective scores of the most semantically similar transcripts is calculated as the LSA score. The strong performance of this automated technique, also validated by KAT in its Intelligent Essay Assessor software, suggests that it could be a very valuable tool for both summative assessment of performance and for feedback—similarity of a new transcript to known performance deficits could be used to provide the most applicable feedback to individual team members.

The Speech Recognition problem

For use in the proposed Analysis Pipeline, either in near-real time or in an After Action Briefing, human typed transcription of the speech to text is not possible, therefore the speech-to-text transcription must be produced automatically. Output produced by commercial Speech Recognition (SR) systems is known to contain errors, even under the best of conditions. The question we want to answer is how robust is LSA in the presence of such noise? In particular how well does LSA correlate with human assessment of performance as errors are introduced into mission communications transcripts?

Synthesizing Noisy Data

Because transcripts produced by a SR system were not yet available to us, we evaluated the robustness of LSA using synthetic SR output. We developed a program to add noise to human-created transcripts of the UAV mission communications. Noise is defined by three types of errors:

- **Insertion**. Insert a word from an LSA space. Inserted words are limited to no more than m characters. In this study we used m=8.
- **Deletion**. Delete a word in the original transcript.
- **Substitution**. Substitute an original word with a word from an LSA space. Substituted words have two constraints. The first p characters must match the original word and the length must be within l characters of the original word. In this study we used p=2 and l=4.

In this study the LSA space from which insertion and substitution words were selected was created from the corpus of mission communications transcripts, which contained 6103 unique terms. Words are chosen randomly, subject to the constraints described above. The constrained space of terms mimics those vocabularies seen in military applications of SR systems.

Original	Sample 1	Sample 2
this is	this is	this is198
Intelligence to AVO	Intelligence to AVO	Intelligence to AVO
Intelligence this is the AVO how many targets have you taken so far?	Intelligence this is thank AVO houses many targets have you tactful edges so?	Intelligence this the AVO how many tazsar have you taken so far?
we've taken two pictures we are on the third one.	maneuver we've taken twice pictures we are the third one.	wearing taken two we on thirty one.
thank you.	thank you.	bunched thank yoda.
I haven't taken the picture yet hold on.	I happen taken the pieces yet hold on.	I haven't taken the picture yeas hold on.
go ahead AVO	go ahead AVO	go ahead aviator
okay DEMPC my question is what is my effective time for change over to MSTE over	okay DEMPC my question is what is my effective time for change over to MSTE over	okay DEMPC myself question issues what is my effective time for checkpoint over
as soon as she take the picture you can switch over to MSTE	asks soon as she tasks the piece you can switch over to mst	kicks to MSTE over asap soon asked she take the picture you can over to MSTE
there's no effective on them.	thick effective systems on thirty.	there's no effective on them.
this is an effective radius of 5.0	is an effective racks of 55	this is an radius of 5.0
picture taken let's go	pitch taken	picture taken let's go
let's change over we are a little off course but we'll get back on track.	let's change over we are secrets a little off course but we'll generic back radio on.	let's change over we are lit wph off course but we'll generic back on 140

Figure 7: Sample SR Degraded Transcripts

Each noise or degradation level is defined in terms of an overall per-word error rate and component insertion, deletion and substitution rates. In this study we used twelve different degradation levels. The first four levels represent "best" and "typical" error rates for two speech recognition algorithms, Linear Predictive Coding (LPC) and Mixed Excitation Linear Prediction (MELP). The remaining eight levels were created using the insertion, deletion and substitution rates of "typical LPC" and "typical MELP" and 57%, 71%, 85% and 99% for the overall error rates. All twelve degradation levels are shown in Table 1.

Table 1. Degradation levels as defined by error rates (percent).

Degradation Level	Overall Error Rate	Insertion Rate	Deletion Rate	Substitution Rate
Best LPC	29	10	28	62
Best MELP	29	14	31	55
Typical LPC	44	11	18	71
Typical MELP	42	14	17	69
57% LPC	57	11	18	71
57% MELP	57	14	17	69
71% LPC	71	11	18	71
71% MELP	71	14	17	69
85% LPC	85	11	18	71
85% MELP	85	14	17	69
99% LPC	99	11	18	71
99% MELP	99	14	17	69

Performance Assessment of Synthetic SR missions

The evaluation corpus consists of 67 simulated mission communication transcripts, produced by human listeners. This evaluation corpus is termed the *verbatim* corpus and is assumed to have an error rate of 0%. The verbatim transcripts were evaluated by LSA to produce a set of *text* and *comparison* measures. Text measures are based on properties of each transcript. Comparison measures are obtained by comparing a transcript to its *k-nearest* neighbors in the LSA space. From these measures, two LSA scores were produced for each transcript. The LSA+ score was produced using stepwise linear regression to build a model from the measures, which predicts the human scores for each transcript. The LSA score is the single LSA *k-near* measure that has the highest correlation with human scores. The reliability of Verbatim LSA+ with human scores is 0.85, while the reliability of Verbatim LSA is 0.74 (see Figure 8).

For each of the twelve degradation levels, five samples of the corpus were generated using the program described earlier. Each sample was then evaluated by LSA to produce a set of text and comparison measures.

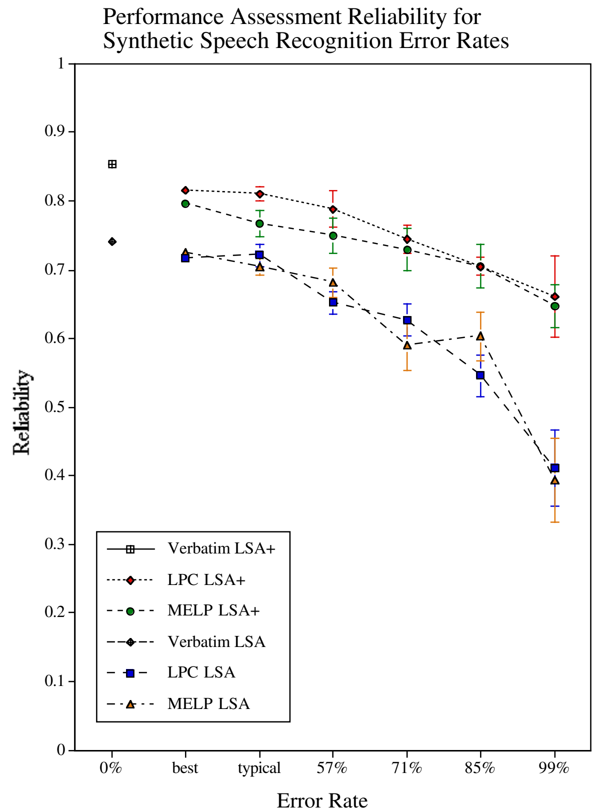

Figure 8. Performance of LSA using SR degraded transcripts

Stepwise regressions and correlations were performed to obtain LSA+ and LSA scores for each sample and to compute reliability with human scores. The reliability measures were averaged over the five samples for each of the twelve degradation levels. Reliability measures, along with standard error bars, are presented in Figure 8. The points connected by the two lines show reliability for LSA+ scores on LPC and MELP samples. The points connected by the bottom two lines show for LSA scores on LPC and MELP samples. Table 2 provides the average reliability scores for the models.

Table 2. Reliability for models

		Best		Typ		57 %	
	Y	L	M	L	M	L	M
LSA	**.74**	.72	.73	.72	.71	.68	.68
LSA+	**.85**	.82	.80	.81	.77	.79	.75

		71 %		85 %		99 %	
	Y	L	M	L	M	L	M
LSA	**.74**	.64	.60	.55	.60	.41	.40
LSA+	**.85**	.74	.73	.71	.71	.66	.65

Y – Verbatim text (human transcriptions
L – Average LPC SR text
M – Average MELP SR text

A second analysis was conducted which compared the original transcripts to their SR counterparts in LSA space. A Cosine Similarity judgment was made between each Original and its SR transcripts. A Cosine Similarity of 1.0 indicates perfect agreement. The average agreements decreased steadily as more and more noise was introduced, but as can be seen in Figure 9, both the best and typical error rates from commercial SR systems are judged very similar by LSA with scores above 0.90.

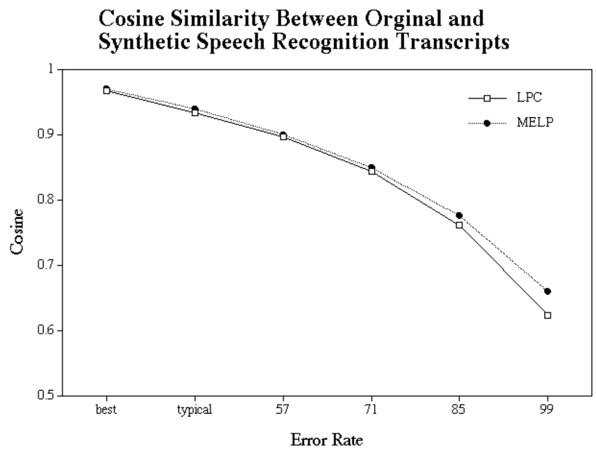

Figure 9. Average similarity judgments

CONCLUSIONS

CareerMap

The CareerMap software represents an initial demonstration of a usable World Wide Web based Intelligent Search Agent based on the LSA technology. Even with its limited knowledge base, it demonstrates the necessary capabilities to match mission and job requirement statements with military personnel and training data. By measuring semantic similarity of training materials and tests, it facilitates combining occupations based on core competencies and similar work activities. It also helps to identify individuals qualified for work activities for which no current occupation exists.

An organization that has acquired a new or revised system can develop detailed descriptions of the activities required to operate or to maintain the system, based on system requirements documents, operations manuals or provided by subject matter experts. Given such descriptions, and assuming an increase in its knowledge, CareerMap could automatically identify current jobs on fielded systems that are similar in component work activities and in their requirements for training. It could also identify similar paragraphs in existing course materials and rank them by probable relevance to work with the new system.

New occupations could be structured around these activities and new sets of training materials assembled, at least in major part, from subsets of existing material. In addition, individuals who work in jobs that use subsets of the competencies and experience required can be identified. This may permit the immediate employment of appropriate personnel or their more rapid and effective re-training for work in support of new systems. In this way, CareerMap could help the military to exploit Internet resources to achieve information superiority.

In the occupational domain, this effort may ultimately produce a cost-effective capability to systematically mine occupational personnel and training databases to develop new job and training structures to support a variety of requirements. This capability will help employers identify critical characteristics and competencies associated with work activities and then to identify individuals who have the requisite experience and competencies to perform the identified work activities.

Automated Mission Communications Analysis

The initial research suggests that LSA will be an effective analysis tool even in conditions where the text to be analyzed has been significantly degraded. The noise introduced by SR systems is essentially random—enough of the original signal survives to be effectively analyzed—even at today's less than optimal SR error rates.

The capabilities suggested by these studies—to automatically and in real-time predict levels of team performance based on their communications and to identify and diagnose common error patterns should provide near future DMT systems with an enormous instructional advantage over current systems. These early success of Latent Semantic Analysis based tools are indicators of continuing improvement in simulator systems which will ultimately lead to better and more cost effective training for our allied warfighters.

REFERENCES

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.

Kiekel, P. A., Cooke, N. J., Foltz, P. W., Gorman, J. & Martin, M. (In press). Some Promising Results of Communication-Based Automatic Measures of Team Cognition. Proceedings of Human Factors 2002.

Laham, D., Bennett, W., and Landauer, T. K. (2000). An LSA-based software tool for matching jobs, people, and instruction. Interactive Learning Environments, 9, 1-15.

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

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of the acquisition, induction, and representation of knowledge. Psychological Review, 104, 211-240.

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). Introduction to Latent Semantic Analysis. Discourse Processes, 25, 259-284.