

AD-A142 712

INDUCTIVE INFORMATION RETRIEVAL USING PARALLEL
DISTRIBUTED COMPUTATION(U) CALIFORNIA UNIV SAN DIEGO LA
JOLLA INST FOR COGNITIVE SCIENCE. M C MOZER JUN 84


1/1


UNCLASSIFIED

ICS-8406 ONR-8405 N00014-79-C-0323

F/G 9/2

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

12

Inductive Information Retrieval Using Parallel Distributed Computation

Michael C. Mozer

June 1984

ICS Report No. 8406

AD-A142 712

DTIC FILE COPY

COGNITIVE
SCIENCE



UCSD

DTIC
ELECTE
JUL 05 1984
S D E

INSTITUTE FOR COGNITIVE SCIENCE

UNIVERSITY OF CALIFORNIA, SAN DIEGO

LA JOLLA, CALIFORNIA 92093

This document has been approved
for public release and sale in
distribution is unlimited.

84 07 03 009

Inductive Information Retrieval Using Parallel Distributed Computation

Michael C. Mozer

June 1984

ICS Report No. 8406

Institute for Cognitive Science, C-015
University of California, San Diego
La Jolla, CA 92093



I am grateful to David Rumelhart, Donald Norman, Paul Smolensky, and Steven Draper for their very helpful suggestions on this project and earlier drafts of the paper. Thanks also to Gary Perlman for use of his database.

This research was supported by Contract N00014-79-C-0323, NR 667-437 with the Personnel and Training Research Programs of the Office of Naval Research and by a grant from the System Development Foundation. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsoring agencies. Approved for public release; distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government.

Copyright © 1984 Michael C. Mozer

ONR Report 8405

This document has been approved
for public release and sales (S)
distribution is unlimited.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ONR #8405	2. GOVT ACCESSION NUMBER <i>A142 712</i>	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Inductive Information Retrieval Using Parallel Distributed Computation		5. TYPE OF REPORT & PERIOD COVERED Technical Report
7. AUTHOR(s) Michael C. Mozer		6. PERFORMING ORG. REPORT NUMBER ICS #8406
		8. CONTRACT OR GRANT NUMBER(s) N00014-79-C-0323
9. PERFORMING ORGANIZATION NAME AND ADDRESS Institute for Cognitive Science Center for Human Information Processing University of California San Diego La Jolla, CA 92093		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 667-437
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel & Training (Code 442PT) Office of Naval Research 800 N. Quincy St. Arlington VA 22217		12. REPORT DATE June 1984
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		13. NUMBER OF PAGES 16
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES This research was also supported by a grant from the System Development Foundation.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) parallel computation natural reasoning parallel distributed processing induction information retrieval search bibliographic retrieval		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) OVER		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102 LF 014 6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

Abstract

Massively parallel, distributed models of computation offer a new approach to the representation and manipulation of knowledge. This paper reports on an application of parallel models to the area of information retrieval. The retrieval system described makes dynamic use of the internal structure of a database to infer relationships among items in the database. Using these relationships, the system can help overcome incompleteness and imprecision in requests for information, as well as in the database itself.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

Inductive Information Retrieval Using Parallel Distributed Computation

Massively parallel models of computation have shown promising results in the areas of pattern recognition, memory, learning, and language comprehension (Anderson & Hinton, 1981; Feldman & Ballard, 1982; McClelland & Rumelhart, 1981; McClelland, Rumelhart, & Hinton, in preparation). These models, called *connectionist* or *parallel distributed processing* (PDP) models, offer a new approach to the representation and manipulation of knowledge. In this approach, computation takes place through the simultaneous interaction of many small pieces of knowledge, some pieces supporting each other and other pieces competing with each other. PDP models have demonstrated the kind of flexibility and computational power that characterize cognition, particularly those aspects of cognition that people perform so effortlessly and naturally. The PDP approach promises to be particularly useful in areas that demand flexible inferencing and reasoning with incomplete or imprecise information (which I call *inductive reasoning*). This paper reports on an application of the PDP approach to one such area, information retrieval.

I will focus on information retrieval systems used in bibliographic search, known as document retrieval systems. Although document retrieval will serve as the primary example of this paper, the discussion generalizes quite readily to information retrieval of all sorts. Indeed, the PDP conceptualization has been applied to information retrieval tasks as diverse as locating files and specifying commands on a computer system (Greenspan & Smolensky, 1984) and organizing information in an on-line manual (O'Malley, Smolensky, Bannon, Conway, Graham, Sokolov, & Monty, 1983).

I begin with a discussion of document retrieval and argue that inductive reasoning is useful in document retrieval, then present a PDP retrieval model and examine its properties.

Document Retrieval

Document retrieval systems are typically used to search for books or articles (hereafter, *documents*) on a particular topic. To aid in this process, each document is labeled by a set of *descriptor terms* that characterize the content matter of the document. For example, a book on Pascal compiler design might be identified by the descriptors PASCAL, COMPILERS, COMPUTING, PROGRAMMING, and LANGUAGES. Queries to such a system often take the form of a Boolean expression of descriptors, and the system reports all documents satisfying the expression. For example, one might request all documents about "PROGRAMMING AND (COMPILERS OR INTERPRETERS)."

There are two difficult problems with such a system. First, users have a hard time accurately specifying the information they are seeking, possibly because the document descriptors have different semantics than they realize, or because they fail to include relevant descriptors in the query, or because they include irrelevant descriptors. The end result is that the set of descriptors chosen for the query is semantically inaccurate or ill-specified. The second

problem with document retrieval systems is that the indexing of documents is itself often inconsistent and incomplete. Relevant descriptors are sometimes omitted from a document, and because documents are added to the collection over long periods of time by many individuals, indexing is often inconsistent. When either query or document descriptors are faulty, traditional retrieval systems generally perform poorly because the retrieval process is too literal minded; it simply matches query descriptors against document descriptors.

Rather than assuming that the query is a precise specification of the user's intentions and that the document descriptors precisely characterize the documents, a retrieval system might better assume that there is some uncertainty in the query and document descriptors and that the descriptors can be interpreted in a loose sense. Under these assumptions, the retrieval process must be considered inductive, not deductive.

What sort of output will an inductive retrieval system produce? The deductive Boolean retrieval system described above produced a set of documents that matched the query exactly; all other documents were rejected. However, an inductive system will, by its very nature, tend to produce a set of documents that match the query to varying degrees. Thus, an inductive system will require a procedure for assigning a quality of match, or *relevance measure*, to each document given a particular query.

There are a variety of methods in the information retrieval literature for assigning a relevance measure to a document (Bärttschi & Frei, 1982; Bookstein, 1980; Buell & Kraft, 1981; Salton, Fox, & Wu, 1983). Some use fuzzy logics to evaluate Boolean queries; others, known as *vector models*, treat documents and queries as vectors in a concept space of large dimensionality and then compute the distance between vectors. However, all of these methods base their judgments of relevance solely on the relationship between document descriptors and query descriptors.

On first thought, there doesn't seem to be any other information on which to base a judgment of relevance. However, the PDP retrieval model proposed below exploits an additional source of information: the internal structure of the database. This structure allows the model to dynamically determine the relationship between two documents or two descriptors, and to use this information in computing a relevance measure.

The PDP Retrieval Model

To understand the retrieval model, it may be useful first to outline the class of PDP models. These models consist of a large number of simple processing units operating in parallel. In most cases, each unit represents a possible hypothesis. Units have varying degrees of confidence in the truth of their hypotheses. The degree of confidence is quantified by an internal state variable of the unit, its *activation level*. Units can transmit their activation levels to one another through connecting links. There are two types of links: excitatory and inhibitory. When two units represent mutually compatible hypotheses, they will be connected by an excitatory link. Excitatory links cause the confidence in one hypothesis to increase the confidence in the other hypothesis. When two units represent mutually incompatible hypotheses, they will be connected by an inhibitory link. Inhibitory

links cause the confidence in one hypothesis to decrease the confidence in the other hypothesis. The outcome of any computation in a PDP model is thus the result of cooperation and competition among a large number of simple processes.

In the PDP retrieval model, each document and each descriptor is represented by a unit. Figure 1 shows a portion of a PDP system, with the upper row of units representing documents and the lower row descriptors. The activation level of a document unit indicates the system's belief in the relevance of the document, i.e., the relevance measure. Large activation levels indicate a high degree of relevance; low or negative activation levels indicate irrelevance. The system decides on the relevance of descriptors as well as documents. Like the documents, the relevance of a descriptor is indicated by its activation level.

Links connecting units permit the flow of activation. Mutually excitatory links connect each document and all descriptors associated with the document. Thus, evidence for a particular descriptor is evidence for all documents associated with it; and evidence for a particular document is evidence for all of its descriptors. Furthermore, there are mutually inhibitory links between every pair of documents. Thus, evidence for one document is counterevidence for all others.

The dynamics of the model are based on McClelland and Rumelhart's (1981; Rumelhart & McClelland, 1982) interactive activation model of word perception. A formal statement of the activation rules and parameter values used to implement the model are included in the Appendix. In simulating the model, time is quantized into discrete steps, and the following sequence of events occurs during each time step. If a unit has a positive activation level, it passes its activation through each of its links; otherwise, no activation is passed. Each unit computes the sum of its incoming activations, weighted by connection strengths associated with each link. This net input, modulated by the current unit activity (in order to prevent the unit's activity from exceeding a certain maximum or minimum level), is added to the current activity. Finally, a unit loses a fixed percentage of its activation during each time step, resulting in an exponential decay of activation over time. The system stabilizes when the net increase in activation to each unit equals the net decrease, that is, when the excitatory input exactly matches the combination of inhibitory input and decay. In the implementation to be described, the system approached equilibrium within about 25 time steps.

Querying the model involves activating a set of descriptor units and seeing which document units become active as a result. In contrast to the Boolean queries described earlier, the PDP-model query merely specifies a set of relevant descriptors. The set includes both positive and negative descriptors, the positive descriptors being those that should be associated with the retrieved documents and the negative descriptors those that should not. The activation levels of the positive-descriptor units are clamped to the maximum allowed level and the negative-descriptor units to the minimum allowed level. The activation levels of these units remain clamped throughout the course of processing and are not affected by decay or incoming activations.

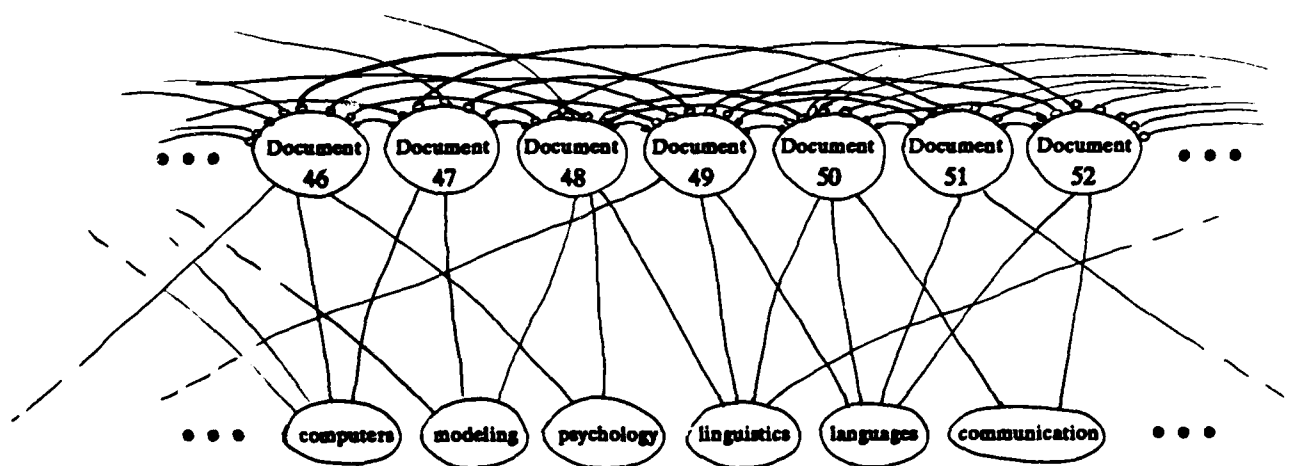


Figure 1. A portion of a PDP document retrieval system.

At an abstract level of description, the system operates as follows. The user activates a set of descriptors. These descriptors activate a set of documents. The documents in turn activate new descriptors, which will activate some new documents as well as reinforce the activation of already active documents, and so on. Activation continually flows from descriptors to documents and vice-versa. This flow of activation allows descriptors in a query to indirectly suggest other descriptors that may be useful in the document search, and it allows active documents to indirectly suggest other documents.

Implementation and Testing of the Model

A retrieval system was developed that constructs a PDP network representation of a database given to it, allows the user to activate sets of descriptor units, and displays documents ranked by activation level. The system includes a graphics interface that shows the activity of all documents and descriptors.

The document collection used to examine the system consisted of books and articles belonging to a graduate student in our lab. The collection had 407 documents, was indexed by 133 descriptors, and had an average of 3.4 descriptors per document. It had been built up over the course of several years. Indexing of the documents had been performed as the documents were added. Consequently, the indexing was somewhat inconsistent. All retrieval examples given in the next section come from this collection.

Properties of the Model

Compensation for Inaccuracy and Incompleteness in the Query

During the course of processing, an active descriptor may cause other descriptors to become active. The retrieved set of documents will be ones that match either the descriptors that were initially active (the query descriptors) or the internally activated descriptors (the *induced descriptors*). For example, consider what happens when the descriptor LINGUISTICS in Figure 1 is activated. Documents 48, 49, and 50 will become active, and each will begin activating its own descriptors; LANGUAGES will thus be activated by documents 49 and 50. Next, LANGUAGES will activate its documents, sending some activation back to documents 49 and 50, but also activating documents 51 and 52. Thus, documents 51 and 52 become active not because their descriptors were specified in the query, but because the system saw one of their descriptors as being related to descriptors in the query. In the document collection used as a testbed, the descriptors LOGIC and REASONING also became slightly active after LANGUAGES did, presumably because these two descriptors are related to LANGUAGES in the same way that LANGUAGES is related to LINGUISTICS. In essence, the system has added LANGUAGES, and to a lesser extent LOGIC and REASONING, to the list of query descriptors.

One descriptor tends to activate another to the extent that the two descriptors co-occur in the currently active subset of the document collection. The tendency for one descriptor to activate another is thus influenced by the other descriptors that are active, because the other active descriptors affect which subset of the collection is active. As an example of this context dependency, if LANGUAGES is activated along with COMPUTERS, the

descriptors PROGRAMMING, AI, and DESIGN become strongly active, whereas if LANGUAGES is activated along with PSYCHOLOGY, then COGNITIVE and LINGUISTICS become active.

Clearly, the model achieves a flexible interpretation of the query. The query is used as a guide in the retrieval process, steering the model's attention in certain directions, not as a rigid criterion that must be matched exactly. This flexibility allows the system to help compensate for inaccurate and incomplete queries.

Compensation for Inconsistency and Incompleteness in the Indexing of the Document Collection

When a document has many descriptors in common with the set of active documents, the document is likely to become active itself, even if it doesn't match the query very well. This property is quite useful if descriptors are accidentally omitted in the indexing of a document. For example, one might suppose that documents 51 and 52 in Figure 1 should have been indexed with the descriptor LINGUISTICS, because many items that concern LANGUAGES also concern LINGUISTICS. The system will make precisely this inference when LINGUISTICS is activated via the induced descriptor mechanism described above.

Thus, the system can help to overcome inconsistency and incompleteness in the indexing of the collection. Of course, the performance of the system degrades with a degradation in indexing, but because of the system's statistical nature, the degradation in performance is gradual.

Ranking of the Retrieved Documents

The activation level of a document is a measure of its relevance according to the model, with the most active documents being most relevant. Thus, ranking the retrieved documents according to their activation levels will produce a set ordered by relevance. The ordered set is quite useful when many documents are retrieved, because it suggests which documents to inspect first.

Let's examine what the ranked set of documents will look like. With the parameter values indicated in the Appendix, there will be a very strong tendency for documents matching the query exactly to be most active, followed by partial matches, followed by documents having no descriptors in common with the query but having at least one active (induced) descriptor. Conventional relevance measures cannot distinguish among the exact matches, and can make only coarse-grain distinctions among the partial matches. This is because these relevance measures are based solely on the query descriptors. However, the PDP model also uses the induced descriptors to form its relevance measure. Thus, two documents that share the same query descriptors may nonetheless be assigned different relevance measures.

Consider what happens when several documents in the collection match the query exactly, as when LANGUAGES is queried in Figure 1. Documents 49 through 52 will immediately become active, all to the same extent. Documents 49 and 50 will then support LINGUISTICS, and documents 50 and 52 COMMUNICATION. These two descriptors will in turn support

their respective documents, with document 50 receiving support from both. Thus, document 50 becomes the most active.

The most active document within the set of exact matches is the one with the greatest number of highly active induced descriptors; documents with successively less activation have fewer such descriptors. Because the induced descriptors are derived from their association with many of the active documents, one may conclude that the most active documents are those that share many descriptors with the other active documents. In this sense, the most active documents are those most *representative* or *prototypical* of the retrieved set. A representative document may serve as a useful example if a query produces a large retrieval set.

Consider now the meaning of representativeness when no document matches the query exactly but there are many partial matches. For example, this is the case in the test collection when a query is made for LEARNING and PROGRAMMING. Documents indexed by the descriptor pairs LEARNING and COMPUTERS, or PROGRAMMING and PSYCHOLOGY, end up with higher activation levels than those indexed by LEARNING but not COMPUTERS, or PROGRAMMING but not PSYCHOLOGY. The first subset is most representative of the retrieval set as a whole, and is clearly a better match to the query than the second. However, conventional relevance measures are unable to make the distinction between these two subsets.

Providing Cues for Continued Search

Information retrieval can be thought of as an iterative process (Tou, Williams, Fikes, Henderson, & Malone, 1982). The user first formulates a rough query, and in successive retrieval attempts the query is reformulated. The induced descriptors may be helpful in reformulating the query. Because these descriptors are associated with many of the retrieved documents, they are useful for partitioning the retrieval set. That is, if the user can state that a certain descriptor should or should not be present among the retrieved documents, the size of the retrieval set can be reduced. The descriptors carrying the most information are those associated with exactly half the retrieval set. Roughly, this is characteristic of the induced descriptors with highest activity.

Retrieval by Example

A query may be formulated using documents instead of descriptors, thus allowing the user to request documents that are similar to a given set of documents. Similarity is measured in terms of common descriptors. For example, if documents 47 and 50 are activated in Figure 1, document 48 is also likely to become activated via the descriptors it shares with documents 47 and 50.

Parameters of the Model

The model has over a dozen parameters (some of which have been removed from the Appendix to simplify the presentation), many more if one considers varying connection strengths for each link. Overall, it was surprising how robust the basic properties of the model were under variation in these parameters. However, several parameters are useful in controlling the behavior of the model. These parameters affect the model's behavior as follows.

Document-Document Inhibition ($w_{D_i D_i}$) and Descriptor Decay Rate (θ_D)

These parameters control how freely the model associates. If $w_{D_i D_i}$ or θ_D are large, activation cannot easily be passed through the system, and the model will tend to retrieve only documents indexed by one or more of the query descriptors; if sufficiently large, the model will behave exactly like a vector model (described in the section on Document Retrieval and in Salton, Fox, & Wu, 1983). Conversely, if these parameters are small, positive feedback reverberates throughout the model and it "hallucinates"; internal activations begin to dominate external (query) activations. To avoid hallucinatory behavior, these parameters are set so that induced descriptors never approach the activation level of the query descriptors.

Document-to-Descriptor and Descriptor-to-Document Connection Strengths ($w_{D_i d_j}$ and $w_{d_j D_i}$)

These parameters control the strength of association between a document and a descriptor. Some descriptors may be more important in characterizing a document than others, and this could be reflected in the connection strengths.

Resting Activation Levels

In the current implementation, all document units have resting activation levels of zero, that is, their activity decays back to zero. The resting activation levels can be adjusted, however, to bias retrieval in favor of certain documents. Suppose that the resting activation level of a document corresponded to its frequency of activation over long periods of time. Then documents that a user retrieved often would be retrieved more readily.

Query Descriptor Activation

Query descriptors can be assigned different levels of activation to vary their relative importance in the query. The activation levels can be set automatically (e.g., using inverse document frequency, as described by Sparck Jones, 1973) or directly by the user.

Document and Descriptor Fan-in Exponents (α_D and α_d)

These parameters help compensate for a bias built into the system. To see the bias, consider a query of MODELING in Figure 1. Documents 47 and 48 will become activated, followed by descriptors COMPUTERS, PSYCHOLOGY, and LINGUISTICS. A positive feedback loop is formed between these documents and their associated descriptors. Assuming activation of the descriptors from other sources is small, document 47 will become slightly less active than document 48, simply because it is associated with fewer descriptors. Thus, positive feedback loops bias activation in favor of documents associated with many descriptors; they do similarly for descriptors associated with many documents.

One solution to this dilemma is to replace each descriptor unit with a dipole, one unit of the dipole representing the descriptor and the other its complement. Connections would be made between each document and one of the two dipole units. Thus, all descriptors would receive an equal number of inputs, as would all dipoles, causing the bias to disappear. However, this solution requires a massive interconnection of documents and dipoles.

An inexact but practical solution is to base the connection strength of links coming into a unit on the number of incoming links. The more links a unit has, the less activity each link can provide. This solution has been implemented using α_D and α_d . These parameters were adjusted until the bias was approximately nullified.

Evaluating the Model

The best support for the model lies in its ability to produce surprising results, often retrieving documents that have no descriptors in common with the query yet are clearly relevant. These are documents that no conventional deductive retrieval system could find. It is difficult to draw firm conclusions as to the model's utility, because systematic data evaluating the model's performance has yet to be collected. However, the model has been tested on two other databases, one of operating-system commands and the other of local restaurants,¹ and the model performs well on both.

Several drawbacks of the model should be noted. First, queries lack the expressive power of a Boolean query formulation; no distinction is made between "and" and "or." This problem is inherent in all vector models (Salton, Fox, & Wu, 1983). Second, running the model on serial hardware in real time with a large database may be difficult. Third, the model requires that the documents be indexed by a highly overlapping and preferably correlated set of descriptors. Without such an indexing scheme, there is no internal structure to

1. In the restaurant database, restaurants were described by their nationality, location, and cost. Because restaurants have only one value for each of these parameters, the co-occurrence of descriptors could not be used to infer, say, that GREEK restaurants are somewhat similar to ARMENIAN restaurants, or that a restaurant in NORTH PARK is similar to one in HILLCREST. Consequently, it was necessary to specify the semantics of the descriptors explicitly. Descriptor semantics were built into the PDP system by linking related descriptors with a strength of connection corresponding to the degree of association of the descriptors. For example, the GREEK descriptor unit was linked to the ARMENIAN unit, but not to the JAPANESE unit.

the database, and the model performs exactly like a vector model (Salton, Fox, & Wu, 1983).

Conclusions

This paper has presented a new area of application for parallel distributed computation. Although evaluation has been informal, the PDP approach to information retrieval seems potentially powerful and robust. At very least, the PDP approach has pointed out that the internal structure of a database can be a useful source of knowledge in retrieval. A stronger claim, however, is that the PDP approach assists users of a retrieval system in tracking down relevant information, particularly when either the query or the database is incomplete or imprecise. The assistance provided is threefold. First, a PDP retrieval system is able to flexibly interpret descriptors in the query and the database. Second, the system ranks retrieved information precisely and in order of presumed relevance to the user. Third, the system suggests directions that users may take to further specify their query.

Appendix - Summary of Activation Rules and Parameter Values of the Test System

The activity level of document unit i at time $t+1$ is given by

$$D_i(t+1) = \begin{cases} (1-\theta_D) D_i(t) + \eta_{D_i}(t) (M - D_i(t)) & \text{if } \eta_{D_i}(t) > 0 \\ (1-\theta_D) D_i(t) + \eta_{D_i}(t) (D_i(t) - m) & \text{if } \eta_{D_i}(t) \leq 0 \end{cases}$$

where $\eta_{D_i}(t)$ is the net input to document i at time t , as given by

$$\eta_{D_i}(t) = (\bar{c}_D / c_{D_i})^{a_D} \sum_{j=1}^{n_d} w_{d_j D_i} U(d_j(t)) - \sum_{k=1}^{n_D} w_{D_i D_k} U(D_k(t)).$$

The activity level of descriptor unit j at time $t+1$ is given by

$$d_j(t+1) = \begin{cases} (1-\theta_d) d_j(t) + \eta_{d_j}(t) (M - d_j(t)) & \text{if } \eta_{d_j}(t) > 0 \\ (1-\theta_d) d_j(t) + \eta_{d_j}(t) (d_j(t) - m) & \text{if } \eta_{d_j}(t) \leq 0 \end{cases}$$

where $\eta_{d_j}(t)$ is the net input to descriptor j at time t , as given by

$$\eta_{d_j}(t) = (\bar{c}_d / c_{d_j})^{a_d} \sum_{i=1}^{n_D} w_{D_i d_j} U(D_i(t)).$$

$U(x)$ is the zero-threshold identity function:

$$U(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ x & \text{if } x > 0. \end{cases}$$

All other quantities are parameters of the document collection or are constants

(the values in parentheses are those used in the test system):

n_D	number of documents in the collection
n_d	number of descriptors in the collection
$w_{d,j}$	strength of connection from descriptor j to document i (0 if no connection; .10 otherwise)
$w_{D,i}$	strength of connection from document i to descriptor j (0 if no connection; .03 otherwise)
$w_{D_i D_i}$	strength of (inhibitory) connection from document k to document i (0 if $k=i$; .0075 otherwise)
θ_d	descriptor decay rate (.25)
θ_D	document decay rate (.10)
M	maximum activity level (1.0)
m	minimum activity level (-0.2)
c_{D_i}	number of descriptors in document i
c_{d_j}	number of documents indexed under descriptor j
\bar{c}_D	average number of descriptors per document in collection
\bar{c}_d	average number of documents per descriptor in collection
α_D	document fan-in exponent (.10)
α_d	descriptor fan-in exponent (.30)

References

- Anderson, J. A., & Hinton, G. E. (1981). Models of information processing in the brain. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bärtschi, M., & Frei, H. P. (1982). Adapting a data organization to the structure of stored information. In G. Salton & H. J. Schneider (Eds.), *Research and development in information retrieval*. Berlin: Springer-Verlag.
- Bookstein, A. (1980). Fuzzy requests: An approach to weighted Boolean searches. *Journal of the American Society for Information Science*, 31, 240-247.
- Buell, D. A., & Kraft D. H. (1981). A model for a weighted retrieval system. *Journal of the American Society for Information Science*, 32, 211-216.
- Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. *Cognitive Science*, 6, 205-254.
- Greenspan, S., & Smolensky, P. (1984). DESCRIBE: Environments for specifying commands and retrieving information by elaboration. *User Centered System Design, Part II: Collected Papers from the UCSD HMI Project* (Tech. Rep. No. 8402). San Diego: University of California, Institute for Cognitive Science.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception, Part 1: An account of the basic findings. *Psychological Review*, 88, 375-407.
- McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (in preparation). Interactive activation: A framework for information processing. To appear in D. E. Rumelhart & J. L. McClelland (Eds.), *Explorations in natural computation*.
- O'Malley, C., Smolensky, P., Bannon, L., Conway, E., Graham, J., Sokolov, J., & Monty, M. L. (1983). A proposal for user centered system documentation. In A. Janda (Ed.), *Proceedings of the CHI '83 Conference on Human Factors in Computing Systems* (pp. 282-285). New York: ACM.
- Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception, Part 2: The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, 89, 60-94.

Salton, G., Fox, E. A., & Wu, H. (1983). Extended Boolean information retrieval. *Communications of the ACM*, 26, 1022-1036.

Sparck Jones, K. (1973). Index term weighting. *Information Storage and Retrieval*, 9, 619-633.

Tou, F. N., Williams, M. D., Fikes, R., Henderson, A., & Malone, T. (1982). RABBIT: An intelligent database assistant. *Proceedings of the National Conference on Artificial Intelligence*, Pittsburgh, PA.

Cognitive Science ONR Technical Report List

The following is a list of publications by people in the Cognitive Science Lab and the Institute for Cognitive Science. For reprints, write or call:

Institute for Cognitive Science, C-015
University of California, San Diego
La Jolla, CA 92093
(619) 452-6771

- ONR-8001. Donald R. Gentner, Jonathan Grudin, and Eileen Conway. *Finger Movements in Transcription Typing*. May 1980.
- ONR-8002. James L. McClelland and David E. Rumelhart. *An Interactive Activation Model of the Effect of Context in Perception: Part I*. May 1980. Also published in *Psychological Review*, 88.5, pp. 375-401, 1981.
- ONR-8003. David E. Rumelhart and James L. McClelland. *An Interactive Activation Model of the Effect of Context in Perception: Part II*. July 1980. Also published in *Psychological Review*, 89, 1, pp. 60-94, 1982.
- ONR-8004. Donald A. Norman. *Errors in Human Performance*. August 1980.
- ONR-8005. David E. Rumelhart and Donald A. Norman. *Analogical Processes in Learning*. September 1980. Also published in J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum, 1981.
- ONR-8006. Donald A. Norman and Tim Shallice. *Attention to Action: Willed and Automatic Control of Behavior*. December 1980.
- ONR-8101. David E. Rumelhart. *Understanding Understanding*. January 1981.
- ONR-8102. David E. Rumelhart and Donald A. Norman. *Simulating a Skilled Typist: A Study of Skilled Cognitive-Motor Performance*. May 1981. Also published in *Cognitive Science*, 6, pp. 1-36, 1982.
- ONR-8103. Donald R. Gentner. *Skilled Finger Movements in Typing*. July 1981.
- ONR-8104. Michael I. Jordan. *The Timing of Endpoints in Movements*. November 1981.

- ONR-8105. Gary Perlman. *Two Papers in Cognitive Engineering: The Design of an Interface to a Programming System and MENUNIX: A Menu-Based Interface to UNIX (User Manual)*. November 1981. Also published in *Proceedings of the 1982 USENIX Conference*, San Diego, CA, 1982.
- ONR-8106. Donald A. Norman and Diane Fisher. *Why Alphabetic Keyboards Are Not Easy to Use: Keyboard Layout Doesn't Much Matter*. November 1981. Also published in *Human Factors*, 24, pp. 509-515, 1982.
- ONR-8107. Donald R. Gentner. *Evidence Against a Central Control Model of Timing in Typing*. December 1981. Also published in *Journal of Experimental Psychology: Human Perception and Performance*, 8, pp. 793-810, 1982.
- ONR-8201. Jonathan T. Grudin and Serge Larochelle. *Digraph Frequency Effects in Skilled Typing*. February 1982.
- ONR-8202. Jonathan T. Grudin. *Central Control of Timing in Skilled Typing*. February 1982.
- ONR-8203. Amy Geoffroy and Donald A. Norman. *Ease of Tapping the Fingers in a Sequence Depends on the Mental Encoding*. March 1982.
- ONR-8204. LNR Research Group. *Studies of Typing from the LNR Research Group: The role of context, differences in skill level, errors, hand movements, and a computer simulation*. May 1982. Also published in W. E. Cooper (Ed.), *Cognitive aspects of skilled typewriting*. New York: Springer-Verlag, 1983.
- ONR-8205. Donald A. Norman. *Five Papers on Human-Machine Interaction*. May 1982. Also published individually as follows: Some observations on mental models, in D. Gentner and A. Stevens (Eds.), *Mental models*, Hillsdale, NJ: Erlbaum, 1983; A psychologist views human processing: Human errors and other phenomena suggest processing mechanisms, in *Proceedings of the International Joint Conference on Artificial Intelligence*, Vancouver, 1981; Steps toward a cognitive engineering: Design rules based on analyses of human error, in *Proceedings of the Conference on Human Factors in Computer Systems*, Gaithersburg, MD, 1982; The trouble with UNIX, in *Datamation*, 27/12, November 1981, pp. 139-150; The trouble with networks, in *Datamation*, January 1982, pp. 188-192.
- ONR-8206. Naomi Miyake. *Constructive Interaction*. June 1982.
- ONR-8207. Donald R. Gentner. *The Development of Typewriting Skill*. September 1982. Also published as Acquisition of typewriting skill, in *Acta Psychologica*, 54, pp. 233-248, 1983.
- ONR-8208. Gary Perlman. *Natural Artificial Languages: Low-Level Processes*. December 1982. Also published in *The International Journal of Man-Machine Studies* (in press).
- ONR-8301. Michael C. Mozer. *Letter Migration in Word Perception*. April 1983. Also published in *Journal of Experimental Psychology: Human Perception and Performance*, 9, 4, pp. 531-546, 1983.

- ONR-8302. David E. Rumelhart and Donald A. Norman. *Representation in Memory*. June 1983. To appear in R. C. Atkinson, G. Lindzey, & R. D. Luce (Eds.), *Handbook of experimental psychology*. New York: Wiley (in press).
- ONR-8303. The HMI Project at University of California, San Diego. *User Centered System Design: Part I, Papers for the CHI 1983 Conference on Human Factors in Computer Systems*. November 1983. Also published in A. Janda (Ed.), *Proceedings of the CHI '83 Conference on Human Factors in Computing Systems*. New York: ACM, 1983.
- ONR-8304. Paul Smolensky. *Harmony Theory: A Mathematical Framework for Stochastic Parallel Processing*. December 1983. Also published in *Proceedings of the National Conference on Artificial Intelligence, AAAI-83*, Washington DC, 1983.
- ONR-8401. Stephen W. Draper and Donald A. Norman. *Software Engineering for User Interfaces*. January 1984. Also published in *Proceedings of the Seventh International Conference on Software Engineering*, Orlando, FL, 1984.
- ONR-8402. The UCSD HMI Project. *User Centered System Design: Part II, Collected Papers*. March 1984. Also published individually as follows: Norman, D.A. (in press), Stages and levels in human-machine interaction, *International Journal of Man-Machine Studies*; Draper, S.W., The nature of expertise in UNIX; Owen, D., Users in the real world; O'Malley, C., Draper, S.W., & Riley, M., Constructive interaction: A method for studying user-computer-user interaction; Smolensky, P., Monty, M.L., & Conway, E., Formalizing task descriptions for command specification and documentation; Bannon, L.J., & O'Malley, C., Problems in evaluation of human-computer interfaces: A case study; Riley, M., & O'Malley, C., Planning nets: A framework for analyzing user-computer interactions; all published in B. Shackel (Ed.), *INTERACT '84, First Conference on Human-Computer Interaction*, Amsterdam: North-Holland, 1984; Norman, D.A., & Draper, S.W., Software engineering for user interfaces, *Proceedings of the Seventh International Conference on Software Engineering*, Orlando, FL, 1984.
- ONR-8403. Paul Smolensky and Mary S. Riley. *Harmony Theory: Problem Solving, Parallel Cognitive Models, and Thermal Physics*. April 1984. The first two papers will appear in *Proceedings of the Sixth Annual Meeting of the Cognitive Science Society*, Boulder, CO, 1984.
- ONR-8404. Donald R. Gentner. *Expertise in Typewriting*. April 1984. To appear in M. T. Chi (Ed.), *The nature of expertise* (in press).
- ONR-8405. Michael C. Mozer. *Inductive Information Retrieval Using Parallel Distributed Computation*. May 1984.

ICS Technical Report List

The following is a list of publications by people in the Institute for Cognitive Science. For reprints, write or call:

Institute for Cognitive Science, C-015
University of California, San Diego
La Jolla, CA 92093
(619) 452-6771

- 8301. David Zipser. *The Representation of Location*. May 1983.
- 8302. Jeffrey Elman & Jay McClelland. *Speech Perception as a Cognitive Process: The Interactive Activation Model*. April 1983. Also published in N. Lass (Ed.), *Speech and language: Volume 10*, New York: Academic Press, 1983.
- 8303. Ron Williams. *Unit Activation Rules for Cognitive Networks*. November 1983.
- 8304. David Zipser. *The Representation of Maps*. November 1983.
- 8305. The HMI Project. *User Centered System Design: Part I, Papers for the CHI '83 Conference on Human Factors in Computer Systems*. November 1983.
- 8306. Paul Smolensky. *Harmony Theory: A Mathematical Framework for Stochastic Parallel Processing*. December 1983. Also published in *Proceedings of the National Conference on Artificial Intelligence, AAAI-83*, Washington DC, 1983.
- 8401. Stephen W. Draper and Donald A. Norman. *Software Engineering for User Interfaces*. January 1984. Also published in *Proceedings of the Seventh International Conference on Software Engineering*, Orlando, FL, 1984.
- 8402. The UCSD HMI Project. *User Centered System Design: Part II, Collected Papers*. March 1984. Also published individually as follows: Norman, D.A. (in press), Stages and levels in human-machine interaction, *International Journal of Man-Machine Studies*; Draper, S.W., The nature of expertise in UNIX; Owen, D., Users in the real world; O'Malley, C., Draper, S.W., & Riley, M., Constructive interaction: A method for studying user-computer-user interaction; Smolensky, P., Monty, M.L., & Conway, E., Formalizing task descriptions for command specification and documentation; Bannon, L.J., & O'Malley, C., Problems in evaluation of human-computer interfaces: A case study; Riley, M., & O'Malley, C., Planning nets: A framework for analyzing user-computer interactions; all published in B. Shackel (Ed.), *INTERACT '84, First Conference on*

Human-Computer Interaction, Amsterdam: North-Holland, 1984; Norman, D.A., & Draper, S.W., Software engineering for user interfaces, *Proceedings of the Seventh International Conference on Software Engineering*, Orlando, FL, 1984.

- 8403. Steven L. Greenspan. *Reference Comprehension: A Topic-Comment Analysis of Sentence-Picture Verification*. April 1984.
- 8404. Paul Smolensky and Mary S. Riley. *Harmony Theory: Problem Solving, Parallel Cognitive Models, and Thermal Physics*. April 1984. The first two papers will appear in *Proceedings of the Sixth Annual Meeting of the Cognitive Science Society*, Boulder, CO.
- 8405. David Zipser. *A Computational Model of Hippocampus Place-fields*. April 1984.
- 8406. Michael C. Mozer. *Inductive Information Retrieval Using Parallel Distributed Computation*. May 1984.

Navy

1 Robert Mohr
Code 8711
Naval Factors Laboratory
NAVTRAC/DPEN
Orlando, FL 32813

1 Dr. Nick Bond
Office of Naval Research
Liaison Office, Far East
APO San Francisco, CA 96303

1 Dr. Fred Chang
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Stanley Gallyer
Office of Naval Technology
800 N. Quincy Street
Arlington, VA 22217

1 CDR Mike Curran
Office of Naval Research
800 N. Quincy St.
Code 270
Arlington, VA 22217

1 Dr. Jude Franklin
Code 7510
Navy Research Laboratory
Washington, DC 20375

1 Dr. Jio Hollar
Code 16
Navy Personnel B & B Center
San Diego, CA 92152

1 Dr. Ed Hitchens
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Moran J. Kerr
Chief of Naval Technical Training
Naval Air Station Memphis (75)
Hillingdon, TN 38054

1 Dr. William L. Haley (82)
Chief of Naval Education and Training
Naval Air Station
Pensacola, FL 32508

1 Dr. James McMichael
Navy Personnel R&D Center
San Diego, CA 92152

Navy

1 Dr. William Montague
WPBC Code 13
San Diego, CA 92152

1 Library, Code P201L
Navy Personnel R&D Center
San Diego, CA 92152

1 Technical Director
Navy Personnel R&D Center
San Diego, CA 92152

1 Commanding Officer
Naval Research Laboratory
Code 2427
Washington, DC 20390

1 Office of Naval Research
Code 433
800 N. Quincy Street
Arlington, VA 22217

1 Office of Naval Research
Code 4410P
800 N. Quincy Street
Arlington, VA 22217

1 Personnel & Training Research Group
Code 41271
Office of Naval Research
Arlington, VA 22217

1 Dr. Bernard Rindland (OIC)
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Robert B. Smith
Office of Chief of Naval Operations
OP-997N
Washington, DC 20350

1 Dr. Alfred F. Soeda, Director
Department N-7
Naval Training Equipment Center
Orlando, FL 32813

1 Dr. Richard Snow
Liaison Scientist
Office of Naval Research
Branch Office, London
Box 39
FPO New York, NY 09510

Marine Corps

1 H. William Greenup
Education Advisor (EC-1)
Education Center, MCEC
Quantico, VA 22134

1 Special Assistant for Marine
Corps Matters
Code 100H
Office of Naval Research
800 N. Quincy St.
Arlington, VA 22217

1 Dr. A.L. Slafobsky
SCIENTIFIC ADVISOR (CODE 69-1)
HQ. U.S. MARINE CORPS
WASHINGTON, DC 20380

Air Force

1 U.S. Air Force Office of Scientific
Research
Life Sciences Directorate, WL
Bolling Air Force Base
Washington, DC 20332

1 Bryan Gallan
AFMEL/LRT
Lowry AFB, CO 80230

1 Dr. Genevieve Haddad
Program Manager
Life Sciences Directorate
AFOSR
Bolling AFB, DC 20332

1 Dr. John Tangney
AFOSR-WL
Bolling AFB, DC 20332

1 Dr. Joseph Yasutake
AFMEL/LRT
Lowry AFB, CO 80230

Air Force

1 Technical Director
U. S. Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Commander, U.S. Army Research Institute
for the Behavioral & Social Sciences
ATTN: PERI-9R (Dr. Judith Grassano)
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Robert Sasner
U. S. Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Army

1 Dr. Richard Snow
Liaison Scientist
Office of Naval Research
Branch Office, London
Box 39
FPO New York, NY 09510

ONR Distribution List

Page 5

Page 4 UCSD/Horran & Rueschhart (NH 667-437) 21-May-84

UCSD/Horran & Rueschhart (NH 667-437) 21-May-84

Department of Defense	Civilian Agencies	Private Sector	Private Sector
12 Defense Technical Information Center Cameron Station, Bldg 5 Alexandria, VA 22314 Attn: IC	1 Dr. Patricia A. Butler NIE-DM Bldg., Mail Stop 87 National Institute of Education 1200 19th Street NW Washington, DC 20036	1 Dr. John R. Anderson Department of Psychology Carnegie-Mellon University Pittsburgh, PA 15213	1 Dr. Jaime Carbonell Carnegie-Mellon University Department of Psychology Pittsburgh, PA 15213
1 Military Assistant for Training and Personnel Technology Office of the Under Secretary of Defense for Research & Engineering Room 3B129, The Pentagon Washington, DC 20301	1 Dr. Paul G. Chapin Linguistics Program National Science Foundation Washington, DC 20550	1 Dr. Alan Baddley Medical Research Council Applied Psychology Unit 15 Chaucer Road Cambridge CB2 2EF ENGLAND	1 Dr. Pat Carpenter Department of Psychology Carnegie-Mellon University Pittsburgh, PA 15213
1 Major Jack Thorpe JAGPA 1400 Wilson Blvd. Arlington, VA 22209	1 Dr. Edward C. Weiss National Science Foundation 1800 K Street, NW Washington, DC 20550	1 Patricia Baggett Department of Psychology University of Colorado Boulder, CO 80309	1 Dr. Michelle Ch. Learning & D Center University of Pittsburgh 3939 O'Hara Street Pittsburgh, PA 15261
1 Dr. Robert A. Nisner OUEHME (ELIS) The Pentagon, Room 3B129 Washington, DC 20301	1 Dr. Joseph L. Young, Director Memory & Cognitive Processes National Science Foundation Washington, DC 20550	1 Dr. Avron Barr Department of Computer Science Stanford University Stanford, CA 94305	1 Dr. William Clancy Department of Computer Science Stanford University Stanford, CA 94305
		1 Dr. Henrich Birenbaum School of Education Tel Aviv University Tel Aviv, Ramat Aviv 69978 Israel	1 Dr. Michael Cole University of California at San Diego Laboratory of Comparative Human Cognition - B005A La Jolla, CA 92093
		1 Dr. John Black Yale University Box 11A, Yale Station New Haven, CT 06520	1 Dr. Allan M. Collins Bolt Beranet & Newman, Inc. 50 Moulton Street Cambridge, MA 02138
		1 Dr. John S. Brown SEROS Palo Alto Research Center 3333 Coyote Road Palo Alto, CA 94304	1 Dr. Lynn A. Cooper LNSC University of Pittsburgh 3939 O'Hara Street Pittsburgh, PA 15213
		1 Dr. Glenn Bryan 6208 Poe Road Bethesda, MD 20817	1 Dr. Emanuel Berchin Department of Psychology University of Illinois Champaign, IL 61820
		1 Dr. Bruce Buchanan Department of Computer Science Stanford University Stanford, CA 94305	1 ERIC Facility-Acquisitions 4833 Pugh Avenue Bethesda, MD 20814
		1 Dr. Alan Baddley Medical Research Council Applied Psychology Unit 15 Chaucer Road Cambridge CB2 2EF ENGLAND	1 Dr. Anders Ericsson Department of Psychology University of Colorado Boulder, CO 80309

Private Sector	Private Sector	Private Sector	Private Sector
1 Dr. Wallace Fourzeng Department of Educational Technology Bolt Beranek & Newman 10 Moulton St. Cambridge, MA 02138	1 Dr. Frederick Hayes-Roth Technology 525 University Ave. Palo Alto, CA 94301	1 Dr. Walter Kintsch Department of Psychology University of Colorado Boulder, CO 80302	1 Dr. Mark Miller Computers/Thought Corporation 1731 West Platte Parkway Plano, TX 75075
1 Dr. John R. Frederiksen Bolt Beranek & Newman 50 Moulton Street Cambridge, MA 02138	1 Dr. Joan L. Heller Graduate Group in Science and Mathematics Education c/o School of Education University of California Berkeley, CA 94720	1 Dr. Stephen Kosslyn 1236 William James Hall 33 Kirtland St. Cambridge, MA 02138	1 Dr. Tom Moran Verco PAC 7333 Coyote Hill Road Palo Alto, CA 94304
1 Dr. Michael Gossensath Department of Computer Science Stanford University Stanford, CA 94305	1 Melissa Holland American Institutes for Research 1055 Thomas Jefferson St., N.W. Washington, DC 20007	1 Dr. Pak Langley The Robotics Institute Carnegie-Mellon University Pittsburgh, PA 15213	1 Dr. Allen Munro Behavioral Technology Laboratories 1815 Elena Ave., Fourth Floor Redondo Beach, CA 90277
1 Dr. Bob Sentner Center for Human Information Processing University of California, San Diego La Jolla, CA 92093	1 Glenda Greenwald, Ed. Human Intelligence Newsletter P. O. Box 1163 Birmingham, AL 35202	1 Dr. Jill Lartin Department of Psychology Carnegie Mellon University Pittsburgh, PA 15213	1 Dr. Jesse Orlansky Institute for Defense Analyses 1801 N. Beauregard St. Alexandria, VA 22311
1 Dr. Robert Slaser Learning Research & Development Center University of Pittsburgh 3939 O'Hara Street Pittsburgh, PA 15260	1 Dr. Earl Hunt Dept. of Psychology University of Washington Seattle, WA 98105	1 Dr. Alan Leopold Learning R&D Center University of Pittsburgh 3939 O'Hara Street Pittsburgh, PA 15260	1 Dr. Nancy Fennington University of Chicago Graduate School of Business 1101 E. 58th St. Chicago, IL 60637
1 Dr. Joseph Sequen SRI International 333 Ravenswood Avenue Menlo Park, CA 94025	1 Robin Jeffries Computer Research Center Hewlett-Packard Laboratories 1501 Page Mill Road Palo Alto, CA 94304	1 Dr. Jia Levin University of California at San Diego Laboratory for Cooperative Human Cognition - 8003A La Jolla, CA 92093	1 Dr. Peter Polson DEPT. OF PSYCHOLOGY UNIVERSITY OF COLORADO BOULDER, CO 80309
1 Dr. Daniel Sopher Faculty of Industrial Engineering & Management TECHNION Haifa 32000 ISRAEL	1 Dr. Marcel Just Department of Psychology Carnegie-Mellon University Pittsburgh, PA 15213	1 Dr. Michael Levine Department of Educational Psychology 210 Education Bldg. University of Illinois Champaign, IL 61801	1 Dr. Mike Posner Department of Psychology University of Oregon Eugene, OR 97403
1 DR. JAMES G. BREWED LRDC UNIVERSITY OF PITTSBURGH 3939 O'HARA STREET PITTSBURGH, PA 15213	1 Dr. Steven M. Keale Dept. of Psychology University of Oregon Eugene, OR 97403	1 Dr. Don Lyon P. O. Box 44 Higley, AZ 85326	1 Dr. Lynn Rader Department of Psychology Carnegie-Mellon University Schensley Park Pittsburgh, PA 15213
1 Dr. Barbara Hayes-Roth Department of Computer Science Stanford University Stanford, CA 94305	1 Dr. Scott Kelso Maximus Laboratories, Inc 270 Crown Street New Haven, CT 06510	1 Dr. Jay McClelland Department of Psychology MIT Cambridge, MA 02139	1 Dr. Fred Reif Physics Department University of California Berkeley, CA 94720

ONR Distribution List

Page 1

Page 8 UCSD/Norase & Ruelhart (NR 467-437) 21-May-84

UCSD/Norase & Ruelhart (NR 467-437) 21-May-84

Private Sector	Private Sector	Private Sector
1 Dr. S. Riley Program in Cognitive Science Center for Human Information Processing University of California, San Diego La Jolla, CA 92093	1 Dr. N. Wallace Sinsko Program Director Hampover Research and Advisory Services Southwestern Institution 801 North Pitt Street Alexandria, VA 22314	1 Dr. Thomas Wictens Department of Psychology Franz Hall University of California 405 Wilgarden Avenue Los Angeles, CA 90024
1 Dr. Andrew R. Rose American Institutes for Research 1955 Thomas Jefferson St. NW Washington, DC 20007	1 Dr. Edward E. Smith Bolt Beranet & Newman, Inc. 50 Houlton Street Cambridge, MA 02138	1 Dr. Mike Williams Intelligence 124 University Avenue Palo Alto, CA 94301
1 Dr. Ernest Z. Rothkopf Bell Laboratories Murray Hill, NJ 07974	1 Dr. Elliott Solomay Yale University Department of Computer Science P.O. Box 2158 New Haven, CT 06520	1 Dr. Joseph Wohl Alphatech, Inc. 2 Burlington Executive Center 111 Middlesex Turnpike Burlington, MA 01803
1 Dr. William B. Rouse Georgia Institute of Technology School of Industrial & Systems Engineering Atlanta, GA 30332	1 Dr. Kathryn T. Spahr Psychology Department Brown University Providence, RI 02912	
1 Dr. David Ruelhart Center for Human Information Processing Univ. of California, San Diego La Jolla, CA 92093	1 Dr. Robert Sternberg Dept. of Psychology Yale University Box 114, Yale Station New Haven, CT 06520	
1 Dr. Roger Schank Yale University Department of Computer Science P.O. Box 2158 New Haven, CT 06520	1 Dr. Albert Stevens Bolt Beranet & Newman, Inc. 10 Houlton St. Cambridge, MA 02238	
1 Dr. Alan Schoenfeld Mathematics and Education The University of Rochester Rochester, NY 14627	1 Dr. Perry N. Thorndyke Perceptronics, Inc. 545 Middlefield Road, Suite 140 Menlo Park, CA 94025	
1 Dr. Alan Schoenfeld Mathematics and Education The University of Rochester Rochester, NY 14627	1 Dr. Kurt Van Lehn Xerox PARC 3333 Coyote Hill Road Palo Alto, CA 94304	
1 Dr. Emanuel Boush Department of Psychology University of Illinois Champaign, IL 61820	1 Dr. Keith T. Wescourt Perceptronics, Inc. 545 Middlefield Road, Suite 140 Menlo Park, CA 94025	
1 Dr. Colin Sheppard Applied Psychology Unit Naval Air Warfare Technology Est. Teddington, Middlesex United Kingdom		

END

FILMED

8

24

100