

AD-A172 196

RAMBOT (RESTRUCTURING ASSOCIATIVE MEMORY BASED ON
TRAINING): A CONNECTION. (U) CALIFORNIA UNIV SAN DIEGO
LA JOLLA INST FOR COGNITIVE SCIENCE. M C MOZER AUG 86


1/1


UNCLASSIFIED

ICS-8610 N00014-85-K-0450

F/G 9/2

NL





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

AD-A172 196


RAMBOTE
A CONNECTIONIST EXPERT SYSTEM THAT
LEARNS BY EXAMPLE

Michael G. Moxer

August 1976

ICS Report 8673

COGNITIVE
SCIENCE



INSTITUTE FOR COGNITIVE SCIENCE

UNIVERSITY OF CALIFORNIA, SAN DIEGO

SAN DIEGO, CALIFORNIA 92162

12

**RAMBOT:
A CONNECTIONIST EXPERT SYSTEM THAT
LEARNS BY EXAMPLE**

Michael C. Mozer

August 1986

ICS Report 8610

*Institute for Cognitive Science
University of California, San Diego
La Jolla, California 92093*



I owe an eternal debt of gratitude to Paul Munro, who christened the system RAMBOT and even managed to make RAMBOT into an acronym on the somewhat generic phrase "Restructuring Associative Memory Based On Training." I also wish to thank Paul Smolensky and Dave Rumelhart for their thoughtful comments. The robots program available at UCSD was written by Allan R. Black of Strathclyde University, and was modified by Stephen J. Muir at Lancaster University.

This research was supported by an IBM Graduate Fellowship, a grant from the System Development Foundation, and the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research, Contract No. N00014-85-K-0450, Contract Authority Identification Number, NR 667-548. 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. Requests for reprints should be sent to Michael C. Mozer, Institute for Cognitive Science, C-015; University of California, San Diego; La Jolla, CA 92093. Copyright © 1986 by Michael C. Mozer.

DISTRIBUTION STATEMENT A

**Approved for public release;
Distribution Unlimited**

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

ADA 172 196

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited.		
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ICS 8610			7a. NAME OF MONITORING ORGANIZATION Personnel & Training Research Programs Office of Naval Research (Code 1142PT)		
6a. NAME OF PERFORMING ORGANIZATION Institute for Cognitive Science University of California, San Diego		6b. OFFICE SYMBOL (If applicable)	7b. ADDRESS (City, State, and ZIP Code) 800 North Quincy Street Arlington, VA 22217-5000		
6c. ADDRESS (City, State, and ZIP Code) C-015 La Jolla, CA 92093		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-85-K-0450			
8a. NAME OF FUNDING / SPONSORING ORGANIZATION		8b. OFFICE SYMBOL (If applicable)	10. SOURCE OF FUNDING NUMBERS		
8c. ADDRESS (City, State, and ZIP Code)		PROGRAM ELEMENT NO. 61153N	PROJECT NO. RR04206	TASK NO. RR04206-0A	WORK UNIT ACCESSION NO. NR 667-548
11. TITLE (Include Security Classification) RAMBOT: A Connectionist Expert System That Learns by Example					
12. PERSONAL AUTHOR(S) Michael C. Mozer					
13a. TYPE OF REPORT Technical		13b. TIME COVERED FROM 85 Oct TO 86 Apr	14. DATE OF REPORT (Year, Month, Day) 1986 August		15. PAGE COUNT 15
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP	Connectionism; parallel distributed processing; learning by observation; artificial intelligence; expert systems; game playing		
05	10				
19. ABSTRACT (Continue on reverse if necessary and identify by block number) <i>discussed</i>					
<p>Expert systems seem to be quite the rage in artificial intelligence, but getting expert knowledge into these systems is a difficult problem. One solution would be to endow the systems with powerful learning procedures which could discover appropriate behaviors by observing an expert in action. A promising source of such learning procedures can be found in recent work on connectionist networks, that is, massively parallel networks of simple processing elements. In this paper, I discuss a connectionist expert system that learns to play a simple video game by observing a human player. The game, Robots, is played on a two-dimensional board containing the player and a number of computer-controlled robots. The object of the game is for the player to move around the board in a manner that will force all of the robots to collide with one another before any robot is able to catch the player. The connectionist system learns to associate observed situations on the board with observed moves. It is capable not only of replicating the performance of the human player, but of learning generalizations that apply to novel situations.</p>					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL Harold Hawkins			22b. TELEPHONE (Include Area Code) (202) 696-4323		22c. OFFICE SYMBOL ONR 1142PT

Contents

INTRODUCTION	1
CONNECTIONIST (PDP) SYSTEMS	1
ROBOTS—THE GAME	3
Tricks of the Game	4
Hiding behind junk heaps.	4
Forcing robots to collide.	4
Lining up robots.	4
Teleporting.	5
Collecting a Corpus of Moves	5
RAMBOT	5
Input Representation	5
Force player into upper-left quadrant of board.	6
Draw windows onto board.	7
Lay out grid.	7
Activate units.	7
Output Representation	8
Overall Network Structure	8
Evaluation of Performance Using Corpus	8
Examples From Play	9
Getting ready to teleport.	9
Forcing robots to collide.	9
Hiding behind a junk heap.	9
Shortcomings.	10
Evaluation of Performance in Free Play	10
Level of death.	10
Average moves per completed level.	13
BEYOND RAMBOT	13
REFERENCES	14



by	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

RAMBOT: A Connectionist Expert System That Learns by Example

MICHAEL C. MOZER

INTRODUCTION

Expert systems are prominent among the successes of artificial intelligence. In fact, expert systems have become so popular that almost any program, if billed as an "expert" system, gains instant notariety. It's not always clear what is or is not an expert system, but the most interesting systems seem to operate in domains where the knowledge involved cannot be expressed in concise algorithms (Charniak & McDermott, 1985). Consequently, the most difficult task in building these systems is encoding the knowledge base (Duda & Shortliffe, 1983). Experts are often not as much help as one would like, because it is hard for experts to specify exactly what it is they're doing that makes them experts.

It would be desirable if expert systems could observe an expert in action and then discover rules of the domain based on their observations. This would allow the experts to do what they do best—perform—rather than what they do poorly—explain their own behavior. Of course, discovering the rules of any domain based on observation is a difficult task and requires powerful learning procedures. One promising source of such learning procedures can be found in the recent work on learning in multilayered connectionist networks (Ackley, Hinton, & Sejnowski, 1985; Barto & Anandan, 1985; Rumelhart, Hinton, & Williams, 1986). These networks have the ability to learn arbitrary associations from a set of known variables to a set of target variables or actions, say, from possible symptoms of a disease to possible treatments. More importantly, the networks are able to generalize from a set of examples to the broader class of situations they may be confronted with. While they generally do not discover explicit, psychologically real rules of the sort that most expert systems use, the behavior of these networks appears "rule governed" (Anderson & Hinton, 1981; Rumelhart & McClelland, 1986).

In this paper, I report on my initial efforts at constructing a connectionist expert system that learns to play a simple computer game by observing a human player. I begin by discussing some relevant properties of connectionist networks. Next, I explain the rules of the computer game, called *robots*, and present some strategy. I then describe my connectionist system, *RAMBOT*, which learns to play the robots game. Finally, I look beyond *RAMBOT* to consider the applicability of connectionist techniques to the design of expert systems in other domains.

CONNECTIONIST (PDP) SYSTEMS

Connectionist, or parallel distributed processing (PDP), systems are networks of simple processing elements that operate in parallel. The typical processing element has a large number of input lines and a single output line. The output line conveys a scalar value, called the *activation level*, and is generally

a function of the weighted sum of the input lines. The output of a unit serves as input to other units or as an output of the system. Similarly, the inputs to a unit are received from other units or may be provided as input to the system.

A typical connectionist network architecture is shown in Figure 1. The network has three layers: an input layer, an intermediate layer, and an output layer. The input units are turned on by an external source, the input units then activate the intermediate units, and the intermediate units in turn activate the output units. This type of network implements an associative memory: an input activity pattern is mapped into an output activity pattern. Another way of conceptualizing the network is to imagine that each input unit represents some feature of an external environment and each output unit represents a possible action that could be taken. In this case, the network performs a stimulus-response mapping. For instance, the input units could signify the political climate of the world, the output units the possible actions that a nation might take (e.g., launch a nuclear attack, bomb innocent children, and so forth).

Learning in this network involves adjusting the strengths of connection, or weights, between units to implement the desired mapping. Until recently, the only known weight-adjusting algorithms were for two-layered networks, that is, networks with direct input/output connections, which are unable to learn many sorts of mappings. However, Barto and Anandan (1985), Ackley, Hinton and Sejnowski (1985), and Rumelhart, Hinton, and Williams (1986) have recently developed learning algorithms for multi-layered networks like the one shown in Figure 1. I have been working with the *back propagation* algorithm of Rumelhart et al.

Using back propagation, a network can be trained to associate a set of paired input/output patterns. This training consists of two phases. In the activation phase, an input pattern is presented and is allowed to flow through the layered network to produce an output pattern. This output pattern is then compared with the target output pattern (the output that is to be associated with the given input) and a measure of discrepancy or error is computed. In the back propagation phase, the error is passed backwards through the network so that each unit has an indication of its contribution to the error. The back propagation algorithm implements gradient descent in the error measure; that is, it specifies a change in the weights that is guaranteed to decrease the error.




FIGURE 1. Typical connectionist network architecture.

Following learning, the network can perform arbitrary mappings between input and output patterns. More importantly, it is capable of generalization: If novel input patterns are presented, the network produces output patterns that appear reasonable in terms of the learned associations. To better understand the nature of such generalizations, consider the simplified case of a network that behaves linearly. Suppose this linear network has learned to associate input pattern A with output pattern A' and B with B'. Then, presenting a pattern that is halfway between A and B will result in an output that is halfway between A' and B'. Thus, generalization in this case consists of linear interpolation and extrapolation of learned patterns.

More realistically, multilayered networks require nonlinearities to achieve interesting sorts of behavior; these nonlinearities complicate the generalization issue. With the sorts of nonlinearities that are typical of back-propagation networks, it is still true that an input pattern composed of a mixture of A and B will produce an output pattern composed of a mixture of A' and B', though only if A and B are *sufficiently similar*. Unfortunately, "sufficiently similar" is difficult to define. A more sensible way of looking at generalization in a network like the one shown in Figure 1 is as follows. Think of the intermediate layer as performing a recoding of the input layer; that is, the intermediate layer constructs an *internal representation* of the inputs, one that is useful for solving the problem at hand. Generalization is then determined by the similarity among internal representations, not similarity among the actual input patterns. Thus, the response to a novel input pattern is similar to the response to known input patterns whose internal representations are similar to that of the novel pattern. Further, if the output units are linear or semilinear (see Rumelhart et al., 1986), the network performs the sort of interpolation and extrapolation described above, except it uses the internal representations of A and B, rather than A and B themselves.

The advantage of generalization should be obvious: the system needn't be trained on every point in the input space. If an unfamiliar input is presented, the system automatically determines its similarity to known inputs and produces a response based on this similarity. In contrast, many systems with explicit rules do not perform well in unfamiliar situations. Often, a missing or overspecified rule will cause the system to grind to a halt.

ROBOTS—THE GAME

I now return to the particular problem I've been working on: teaching a connectionist network to play the game robots.

The game is played on a CRT screen. The version I've worked with uses a 20×20 cell board. A sample board is shown in Figure 2A. The player is represented by an "I" and takes up one cell on the board. There are a varying number of robots, each represented by an equal sign. At the start of the game, the robots are placed on the board at random. On each turn, the player can move to an adjacent cell, remain at the current location, *teleport*, or *wait*. Teleport means that the player is lifted from the current location to a random location on the board; wait means that the player stays at the current location for the remainder of the game. The utility of these commands will be explained shortly.

After the player moves, each of the robots is allowed to move to an adjacent cell. The robots follow a simple algorithm. They march directly towards the player:

$$\begin{aligned} robot_delta_x &= sign(player_x - robot_x) \text{ and} \\ robot_delta_y &= sign(player_y - robot_y). \end{aligned}$$

If two robots land in the same cell, they collide and are replaced by a *junk heap*. Junk heaps are inert and harmless to the player. The player must walk around junk heaps, but if robots collide with a junk heap, they are destroyed and become part of the heap. Figure 2B shows the game state one move after Figure 2A, where the player has moved left. As one can see, the two robots nearest to the player have collided and formed a junk heap, represented by an "@".




FIGURE 2. A sample board before the player's move (A), and the board following one move by player and robots (B). The player moved to the left.

The player is forbidden from moving off the board, moving onto a cell with a junk heap or a robot, or moving adjacent to a robot. The player can die in one of three ways: by teleporting onto a robot, by teleporting onto a cell adjacent to a robot (in which case the robot moves over and crushes him), or by using the wait command when there is a robot directly in the line of sight (in which case the robot marches to the player and crushes him).

The object of the game is to kill off the robots by forcing them to run into each other or into junk heaps. When this happens, the game restarts at a higher level of difficulty, meaning that there are more robots on the board. The game begins at level 1 with 10 robots, and the number of robots increases to nearly 200 by level 9.

Tricks of the Game

Hiding behind junk heaps. If the player hides to one side of a junk heap and robots are approaching from the other side, the robots will march into the heap. For instance, the robot on the same row as the player in Figure 2B will eventually hit the junk heap. If in fact all robots are on the opposite side of the heap, the player can use the wait command and wipe them out in a single turn.

Forcing robots to collide. Whenever two robots are aligned in a column or row (i.e., having the same x- or y-coordinate), the player has the opportunity to force the robots to collide with one another. In the case of two horizontally aligned robots, the player must be positioned in between the two robots, and either above or below. As long as the player remains above or below the robots, the robots will attempt to converge on the player, and will collide in the process. (The same applies to two vertically aligned robots.) For example, there are two horizontally aligned robots in the lower right-hand corner of Figure 2B. If the player remains at the current location or moves down and to the left, the robots will run into one another before they reach the player. Further, the player can control exactly where the robots will collide by manipulating their speed of convergence. The robots will converge fastest if the player is between the two robots, half as fast if the player is aligned with one of the robots, or not at all if the player is off to one side of the pair.

Lining up robots. Because it is so useful to have robots aligned horizontally or vertically, a good strategy is to try forcing the robots to line themselves up. Even if lining up the robots will not help in the present situation, it may be that after teleporting, the player will be in a position where robots will

collide as they converge on him. One simple heuristic for lining up the robots is to march towards the center of mass of the robots. As the robots approach, their movement strategy will tend to place them along the horizontal and vertical axes centered on the player's location.

Teleporting. If the player becomes trapped by robots, teleporting is the only option. However, because there is no guarantee that teleporting will land the player in a "safe" position, the player should avoid teleporting unnecessarily, especially at higher levels of the game.

Collecting a Corpus of Moves

Robots is an addicting game. With practice, it also becomes fairly automatic. One can play while carrying on a conversation or eating. Over a period of several weeks, I played nearly 300 games, and recorded the games to use as examples for RAMBOT. "Recording a game" means saving an image of the board at the start of every turn, along with my move in response to that situation. In the end, 18,200 of these "situation-response" pairs were saved. This corpus did not, of course, represent optimal play; it represented my abilities and included occasional errors, which were not screened out.

RAMBOT

RAMBOT's goal was to learn associations between situations from the corpus and the corresponding responses. That is, given a board image as input, RAMBOT was to produce as output the corresponding move that I made. Following learning, RAMBOT should be capable not only of replicating my performance, but also of generalizing its learning to novel situations: when presented with a situation "similar" to ones it has observed, RAMBOT should suggest a response "similar" to the observed responses.

As in most connectionist networks, input and output representations play a critical role in determining the notion of "similarity," and hence, in determining the sort of generalizations that will be made and the overall difficulty of the learning task. In principle, input and output representations are not important, because with sufficient units in the intermediate layer, any input/output mapping can be achieved. However, practical limitations on the number of intermediate units demand careful selection of input and output representations. With appropriate representations, some of the similarity structure of the game can be built explicitly into the network.

Input Representation

The simplest input representation would be to have two units for each cell on the board. One unit would be turned on if there was a robot in the corresponding cell, the other would be turned on if there was a junk heap, and perhaps both would be turned on if the player was in that cell. However, this scheme has a serious drawback, which can be seen by considering the representations that would be generated for the situations shown in Figures 3A-D. The set of units activated in one situation does not overlap with the set activated in another. Because overlap among input patterns is necessary for the explicit representation of similarity, this encoding does not suggest that the four situations are related, when in fact they are extremely similar: the appropriate response to each situation is to move away from the robots and then stay put, allowing the robots to collide with one another. More generally, a player's response should, for the most part, depend only on the relative location of the robots with respect to the player and each other, not on the absolute location of the player, nor strictly on the absolute distance of the player to the robots, nor on the absolute orientation of the player with respect to the robots (see Figure 3E for an exception). An input representation is required that captures the location,




FIGURE 3. Four similar situations (A-D) in which the player's response should be unaffected by absolute location, scale, and orientation of the player with respect to the robots. In one situation (E), the absolute location does matter due to the edge of the board.

scale, and orientation invariances present in the game. The representation chosen for RAMBOT achieves these invariances to a degree. To explain this representation, the steps involved in constructing it from a board image are described.

Force player into upper-left quadrant of board. The board is flipped around so that the player always lies in the upper-left quadrant of the board. If the player is in the upper-right quadrant, the board (and the corresponding move made by the player) is mirrored around the central vertical axis; if the player is in the lower-left quadrant, the board is mirrored around the central horizontal axis; if the player is in the lower-right quadrant, the board is mirrored around both vertical and horizontal axes. Following this procedure, the relative position of the player is fixed with respect to the nearest corner of the board.

Draw windows onto board. Two windows are drawn on the board: a *fine-scale* window, which is centered on the player's location and covers a 13×13 region (6 cells to either side of the player), and a *coarse-scale* window, which is fixed with respect to the player's location and covers a 29×29 region (9 cells to the left of the player, 19 to the right). Because the player is located in the upper-left quadrant, the coarse-scale window is guaranteed to enclose the entire board.

Lay out grid. Within each window, a 7×7 grid of equally spaced points is laid out to span the region within the window. For the fine-scale window, this means that grid points fall on every other cell; for the coarse-scale window, grid points fall on every fourth cell.

Activate units. Each grid point represents the receptive-field center of two input units. One unit is activated by robots in its immediate vicinity, the other by junk heaps. More specifically, the receptive fields of the fine-scale and coarse-scale units are as follows:

			0.15	0.19	0.20	0.25	0.20	0.19	0.15
			0.19	0.25	0.31	0.50	0.31	0.25	0.19
			0.20	0.31	0.45	0.75	0.45	0.31	0.20
0.25	0.50	0.25	0.25	0.50	0.75	1.00	0.75	0.50	0.25
0.50	1.00	0.50	0.20	0.31	0.45	0.75	0.45	0.31	0.20
0.25	0.50	0.25	0.19	0.25	0.31	0.50	0.31	0.25	0.19
			0.15	0.19	0.20	0.25	0.20	0.19	0.15

These numbers represent the amount of activation that will be assigned to a unit given that a robot or junk heap appears in various locations with respect to the unit's receptive-field center. For example, if a robot is located at the center of a fine-scale robot-detecting unit's receptive field, 1.0 units of activation will be added to that unit's activation level. If the robot is located in the cell to the immediate lower right of the unit's receptive field center, 0.25 units of activation will be added. Because receptive fields overlap, any object on the board may produce activation in several units. However, the receptive fields are designed to guarantee that the net activity produced by any object is constant, independent of the object's location or the number of receptive fields it lies within.

The walls around the playing field are treated like junk heaps. For practical purposes, they behave the same way—they are inert and the player is not allowed to walk into them. However, because of the large number of points defining each wall, the activity of a wall point was set to only 5% of that produced by a junk heap. The aim was to prevent the presence of walls from overwhelming information about junk heaps.

This representation has many virtues. First, the player can be in various locations on the board, yet the input patterns will look similar if the local arrangement of robots and junk heaps is similar; nonetheless, activations from the walls serve to distinguish cases in which the player is trapped in a corner. Second, because the receptive fields of the units are so broadly tuned, a certain amount of scale invariance is built in. Third, the locations of objects are coarse coded (Hinton, McClelland, & Rumelhart, 1986), meaning that each object activates several nearby units. This helps to define the two-dimensional structure of the board by way of correlations in activity among neighboring units. Fourth, by coding the player's location with respect to the nearest corner, important orientation invariances are captured. Fifth, the two windows onto the board provide both a foveal and global view of the situation, with high resolution in the foveal view.¹

¹ Two ideas for improving the input representation seem promising but have not been implemented. First, the coarse-scale window wastes a large proportion of its units because they lie off the board. If the window "wrapped around" from one edge of the board to the other, the number of units could be reduced and the remaining units would be better utilized. Second, additional units could be added to the input pattern to represent a temporal context (a time-decaying trace) of previous moves (Jordan, 1985). Thus, the input pattern would specify not only the current board but also a recent history of moves; this would give the network the ability to learn plans extending over time.

Output Representation

The output representation is straightforward. There is one unit for each of the eight "directional" moves, one unit for remaining in the current location, and one unit each for teleport and wait. To provide RAMBOT with explicit information about the arrangement of the directional moves, target output patterns showed activation not only for the selected move, but also for its two "neighbors" (the directional moves to either side of the selected move). For example, when the player responded by moving directly upwards in a given situation, RAMBOT learned to associate that situation with an activity level of 1.0 for the "up" unit, and also with an activity level of 0.2 for the "up-left" and "up-right" units.

Overall Network Structure

The input layer had 196 units, the intermediate layer 74 units, and the output layer 11. The number of intermediate units selected was based on a guess of sufficiency conditions; little work has been done to estimate the necessary number of units. There was full connectivity from one layer to the next, but no direct connections from input to output layers. There were a total of 15,318 connections. The intermediate and output units were semilinear units with a logistic activation function, as described in Rumelhart, Hinton, and Williams (1986).

Evaluation of Performance Using Corpus

RAMBOT has been presented with nearly a million learning trials. This amounts to innumerable hours on our Sun-2's, but only on the order of 12 hours of Cray CPU time. Figure 4 shows performance as a function of learning trial. The bottom line indicates the percent of trials in which the most active output unit corresponded the stored response; the middle line indicates the percent of trials in which either the most active output unit or, if it was a directional move, one of its neighbors corresponded to the stored response; and the top line indicates the percent of trials in which either the most active or second most active output unit corresponded to the stored response.

Performance continues to improve, though the bottom line appears to be approaching an asymptote around 73%. This turns out to be quite impressive, for the following reason. I wrote a program that randomly selected boards from the corpus, displayed them for me, and allowed me to make a new response without knowledge of my original response. Replaying over 10% of the corpus in this fashion, I was able to match my original responses on only 66% of the trials. Thus, RAMBOT is at least as good at predicting my moves as I am.

The graph also shows RAMBOT's ability to generalize. Points labeled with *x*s immediately follow the addition of new moves to the corpus. (The corpus started with only about 4,000 moves and was gradually built up to 18,200.) Performance was barely affected when new moves were added. Thus, RAMBOT is able to respond to unfamiliar moves with almost the same degree of accuracy as to familiar moves. (The drop in performance following a rearrangement of the learning trials, the *r* points, is due to the use of momentum in the back-propagation rule; see Rumelhart et al., 1986.)

In addition to replaying old games, RAMBOT can, of course, play new games. For this purpose, RAMBOT was set up to interact with the robots game. At the start of each turn, an image of the board was encoded on the input units of the network. Activation was allowed to flow through the network to the output units, and the most active output unit was selected as RAMBOT's move. If this move was invalid (i.e., it involved walking into a wall, robot, junk heap, or adjacent to a robot), the move was discarded and the next most active move was considered. The selected move was then fed back to the robots game, the robots were allowed to move, and this cycle repeated.

Although the time required for learning was substantial, play proceeds in real time. Move selection takes about 1-2 seconds on a Sun-2 with a floating point board. Watching RAMBOT play is impressive. Most of the time it does just what I would have done—a clever program indeed.

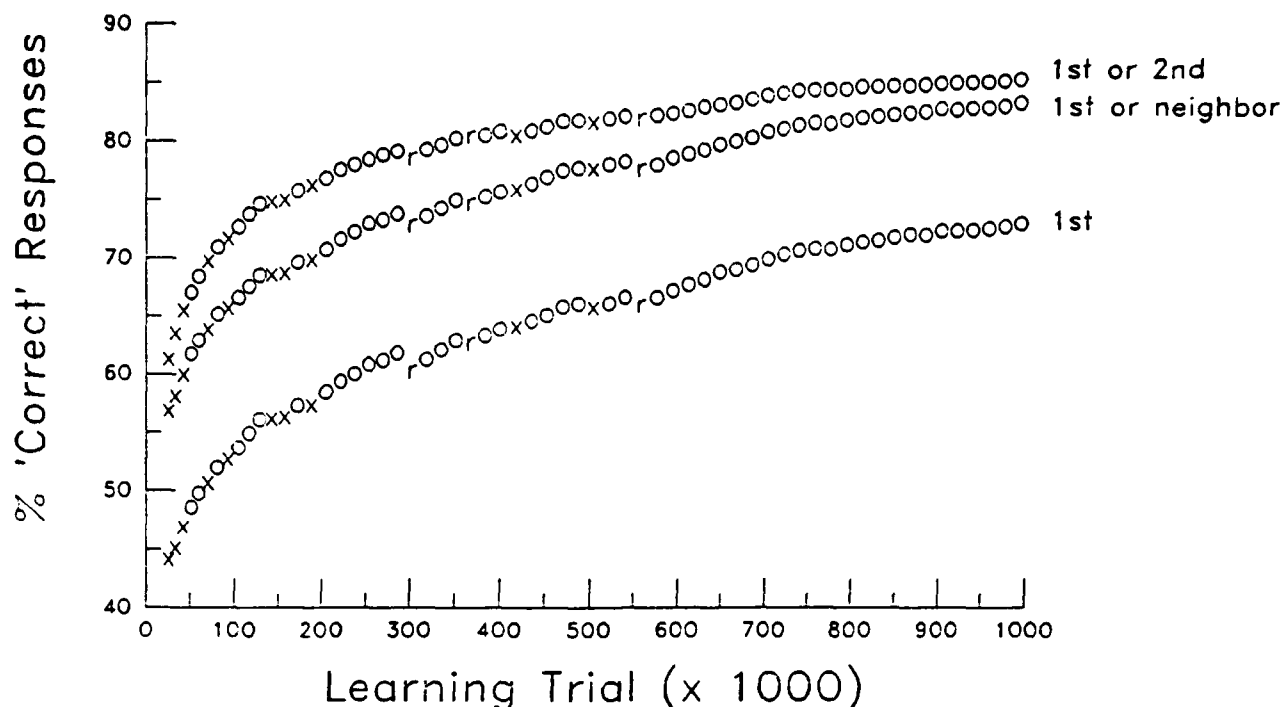


FIGURE 4. Performance as a function of learning trial. The bottom line indicates the percent of trials in which the most active output unit corresponded to the stored response; the middle line indicates the percent of trials in which either the most active output unit or, if it was a directional move, one of its neighbors corresponded to the stored response; and the top line indicates the percent of trials in which either the most active or second most active output unit corresponded to the stored response. *x* = new moves added to corpus; *r* = reordered presentation sequence.

Examples From Play

What follows are several typical examples from actual play.

Getting ready to teleport. Figure 5 shows the board at the start of a turn, as well as the activation levels of the output units in response to that board. The 3x3 array of numbers indicates the activity levels of the eight directional moves and the remain-in-current-location move, arranged by direction. The letter *t* stands for the teleport unit, *w* for the wait unit. Activation levels range from 0-1. The activation level of the selected move is flagged by an asterisk. In Figure 5, RAMBOT is trapped and its only valid move is to teleport. This is the move with the highest activation level. It is interesting to note that three other moves receive some activation: down, down-left, and left. These are the moves that one would consider if the wall were not present. Thus, this example shows that RAMBOT has learned certain facts about the game: walking into walls and robots is not an option, and when being chased by robots, move away from them.

Forcing robots to collide. Figure 6 shows a sequence of moves in which RAMBOT comes around from the right of two horizontally aligned robots and forces them to collide with one another.

Hiding behind a junk heap. Figure 7 shows a sequence of moves in which RAMBOT uses a junk heap to protect itself. RAMBOT first moves towards the junk heap, forcing the robot to its right to crash into the heap, then moves above the heap, forcing the robots below to crash into the heap.

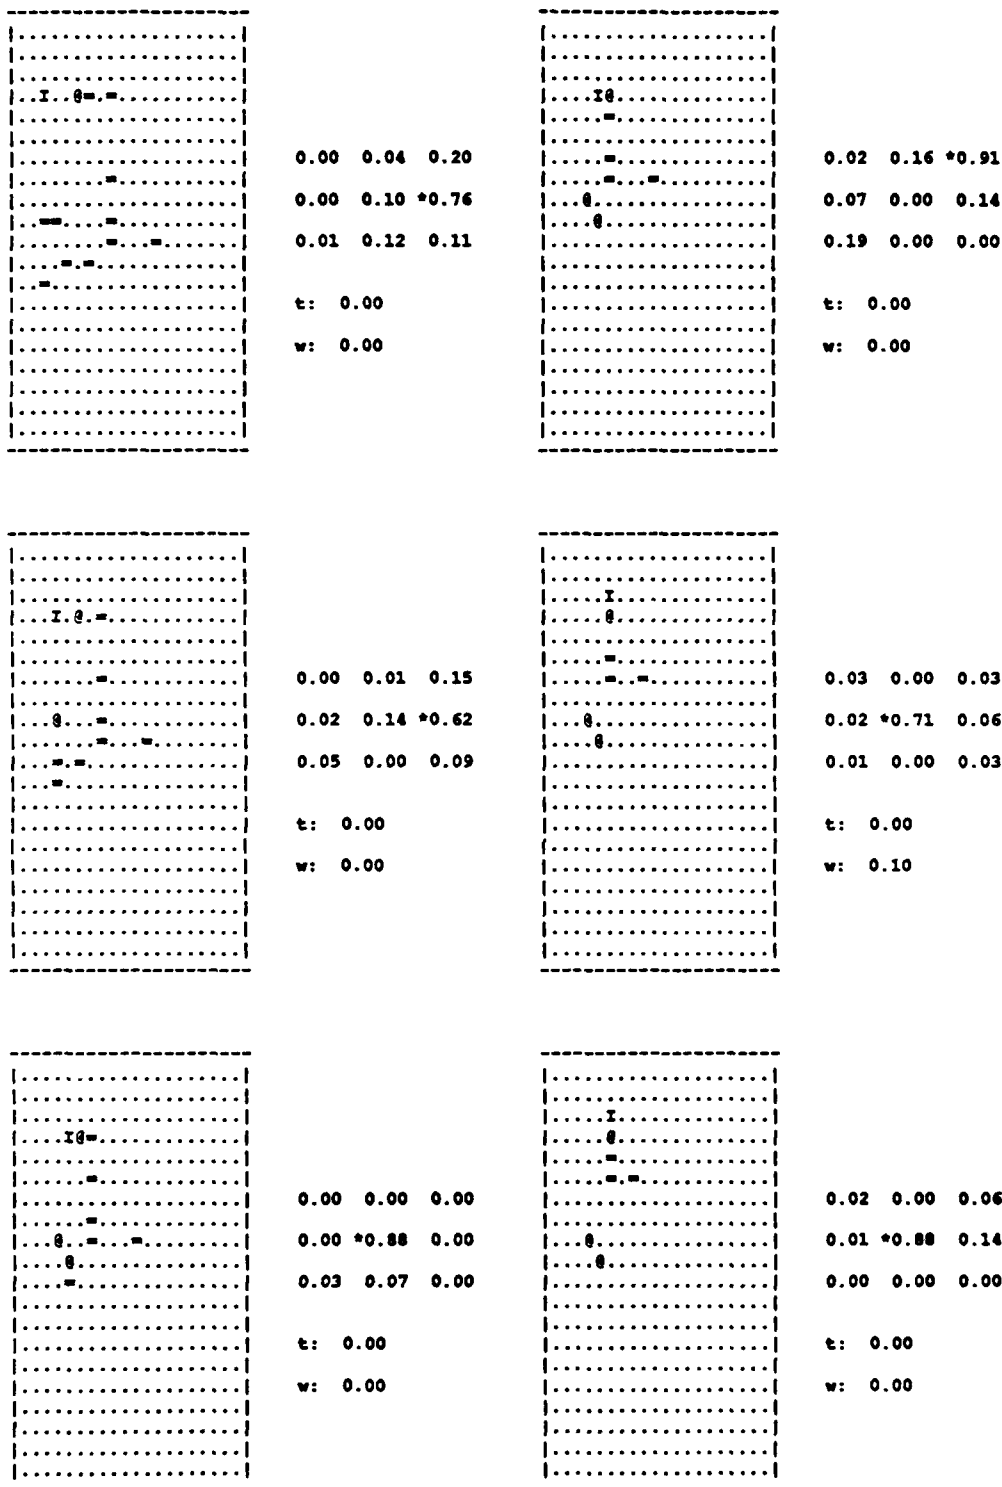


FIGURE 6. A sequence of six moves in which RAMBOT forces two robots to collide. The sequence begins in the upper-left corner, moves down the column, and then on to the right-hand column.

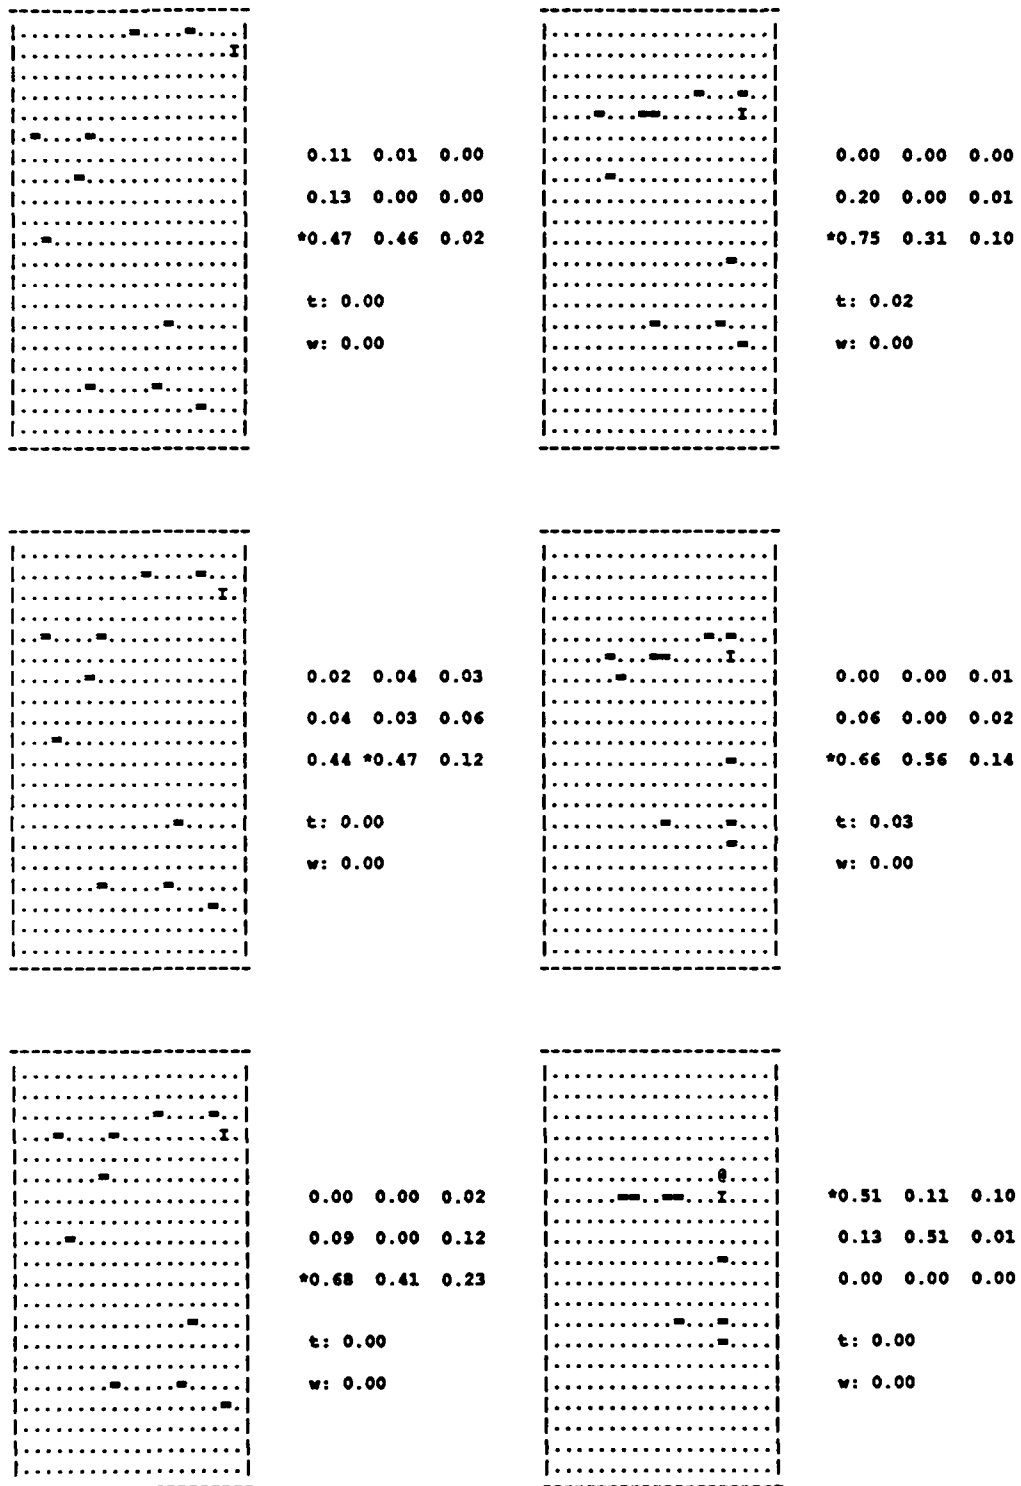


FIGURE 7. A sequence of six moves in which RAMBOT uses a junk heap to protect itself.

What are the costs and benefits of building such a system with connectionist techniques? To begin with, domain experts must specify every source of information that is potentially relevant to their decision processes (though they need not specify how the information is used). This information serves as the input to the connectionist net. Further, experts must provide a corpus of performance data, sufficiently large to sample the input space well. Without representative sampling, the system will not have solid ground on which to base generalizations.

The one serious drawback to a connectionist expert system is that the system itself has little power of explanation. It is possible to examine the outputs of the intermediate units, and to use the "internal representations" developed by the system as a justification for decisions, but generally these internal representations are so complex and highly distributed that they simply add to confusion rather than help to explain the system's behavior. A more reasonable means of increasing the explanatory power of the system is to break down its decision process. For example, in the case of medical diagnosis, the appropriate input/output mapping would not be from symptoms to diseases, but from known symptoms to possible diseases *and* further tests that could be performed to discover additional symptoms. Thus, the system could be used iteratively, performing tests suggested by the network and then feeding results of these tests back into the system. This approach at least provides a sequence of steps taken by the system to reach a decision.

RAMBOT does illustrate several important and unique properties. First and foremost, the system is able to generalize from training examples. Second, the system is able to learn behavior that is dependent on an extremely large number of variables—the robots playing board contains 400 cells—and is able to learn despite inconsistent expert behavior, as my inability to reproduce moves in the learning corpus attests to. Third, the system is able to suggest multiple hypotheses with varying degrees of certainty, as embodied by the activation levels of the output units. Fourth, the system allows for the non-linear combination of evidence, unlike many expert systems that use numerical methods (Charniak & McDermott, 1985).

Beyond these generalities, what does the success of RAMBOT have to suggest for the construction of learning connectionist expert systems in other domains? One problem with the robots game is that it can easily be thought of as a perceptual, pattern-matching task. Connectionist systems are commonly held to be good at this sort of task, but it is not as clear that connectionist techniques will prove useful in "higher-level," symbolic domains. As an argument against this point of view, consider a domain far removed from perception: using a computer operating system, say UNIX. A connectionist expert system for this domain is feasible. The idea would be to build a UNIX apprentice program (UNIXBOT?) that could learn to predict what command the user was likely to type next based on a recent history of commands and some contextual information, such as the time of day or the current working directory. If the system could make strong enough predictions, it could correct user errors, or even anticipate commands.

In principle, a system that learns to predict what command will be typed next is no different than one that learns to predict the next move of a game. It seems that much of cognitive behavior can be framed in terms of pattern recognition, even though we don't ordinarily think of that behavior as being perceptual. Experts in a domain just "see" solutions (Rumelhart, 1984). If this is indeed true, connectionist techniques may have application to a wide range of expert systems applications.

REFERENCES

- Ackley, D., Hinton, G., & Sejnowski, T. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9, 147-169.
- 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.
- Barto, A. G., & Anandan, P. (1985). Pattern recognizing stochastic learning automata. *IEEE Transactions on Systems, Man, and Cybernetics*, 15, 360-375.
- Barto, A. G., Sutton, R. S., & Brouwer, P. S. (1981). Associative search network: A reinforcement learning associative memory. *Biological Cybernetics*, 40, 201-211.
- Charniak, E., & McDermott, D. (1985). *Introduction to artificial intelligence*. Reading, MA: Addison-Wesley.
- Duda, R. O., & Shortliffe, E. H. (1983). Expert systems research. *Science*, 220, 261-268.
- Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1: Foundations*. Cambridge, MA: MIT Press/Bradford Books.
- Jordan, M. I. (1985). *The learning of representations for sequential performance*. Unpublished doctoral dissertation, University of California, San Diego.
- Rumelhart, D. E. (1984, October). *The nature of expertise*. General discussant at the meeting on The Nature of Expertise, Pittsburgh, PA.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1: Foundations*. Cambridge, MA: MIT Press/Bradford Books.
- Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models*. Cambridge, MA: MIT Press/Bradford Books.

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 and 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. Also published in A. Janda (Ed.), *Proceedings of the CHI '83 Conference on Human Factors in Computing Systems*. New York: ACM, 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. (1984), Stages and levels in human-machine interaction, *International Journal of Man-Machine Studies*, 21, 365-375; 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 and Eric M. Segal. *Reference Comprehension: A Topic-Comment Analysis of Sentence-Picture Verification*. April 1984. Also published in *Cognitive Psychology*, 16, 556-606, 1984.
8404. Paul Smolensky and Mary S. Riley. *Harmony Theory: Problem Solving, Parallel Cognitive Models, and Thermal Physics*. April 1984. The first two papers are published in *Proceedings of the Sixth Annual Meeting of the Cognitive Science Society*, Boulder, CO, 1984.
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.
8407. David E. Rumelhart and David Zipser. *Feature Discovery by Competitive Learning*. July 1984. Also published in *Cognitive Science*, 9, 75-112, 1985.
8408. David Zipser. *A Theoretical Model of Hippocampal Learning During Classical Conditioning*. December 1984.
8501. Ronald J. Williams. *Feature Discovery Through Error-Correction Learning*. May 1985.
8502. Ronald J. Williams. *Inference of Spatial Relations by Self-Organizing Networks*. May 1985.
8503. Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. *Direct Manipulation Interfaces*. May 1985. Also published in D. A. Norman & S. W. Draper (Eds.), *User Centered System Design: New Perspectives on Human-Computer Interaction*, 1986, Hillsdale, NJ: Erlbaum.
8504. Mary S. Riley. *User Understanding*. May 1985. Also published in D. A. Norman & S. W. Draper (Eds.), *User Centered System Design: New Perspectives on Human-Computer Interaction*, 1986, Hillsdale, NJ: Erlbaum.
8505. Liam J. Bannon. *Extending the Design Boundaries of Human-Computer Interaction*. May 1985.
8506. David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. *Learning Internal Representations by Error Propagation*. September 1985. Also published in D. E. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol. 1. Foundations*, 1986, Cambridge, MA: Bradford Books/MIT Press.
8507. David E. Rumelhart and James L. McClelland. *On Learning the Past Tense of English Verbs*. October 1985. Also published in J. L. McClelland, D. E. Rumelhart, & the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol. 2. Psychological and Biological Models*, 1986, Cambridge, MA: MIT Press/Bradford Books.

8601. David Navon and Jeff Miller. *The Role of Outcome Conflict in Dual-Task Interference*. January 1986.
8602. David E. Rumelhart and James L. McClelland. *PDP Models and General Issues in Cognitive Science*. April 1986. Also published in D. E. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Vol. 1: *Foundations*, 1986, Cambridge, MA: MIT Press/Bradford Books.
8603. James D. Hollan, Edwin L. Hutchins, Timothy P. McCandless, Mark Rosenstein, and Louis Weitzman. *Graphical Interfaces for Simulation*. May 1986. To be published in W. B. Rouse (Ed.), *Advances in Man-Machine Systems* (Vol. 3). Greenwich, CT: Jai Press.
8604. Michael I. Jordan. *Serial Order: A Parallel Distributed Processing Approach*. May 1986.
8605. Ronald J. Williams. *Reinforcement Learning in Connectionist Networks: A Mathematical Analysis*. June 1986.
8606. David Navon. *Visibility or Disability: Notes on Attention*. June 1986.
8607. William Appelbe, Donald Coleman, Allyn Fratkin, James Hutchison, and Walter J. Savitch. *Porting UNIX to a Network of Diskless Micros*. June 1986.
8608. David Zipser. *Programming Networks to Compute Spatial Functions*. June 1986.
8609. Louis Weitzman. *Designer: A Knowledge-Based Graphic Design Assistant*. July 1986.
8610. Michael C. Mozer. *RAMBOT: A Connectionist Expert System That Learns by Example*. August 1986.

Earlier Reports by People in the Cognitive Science Lab

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 Movement*. 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*, 20, pp. 373-419, 1984.
- 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).

Distribution List [UCSD/Rumelhart] NR 667-548

Dr. Phillip L. Ackerman
University of Minnesota
Department of Psychology
Minneapolis, MN 55455

Dr. Beth Adelson
Department of Computer Science
Tufts University
Medford, MA 02155

AFOSR,
Life Sciences Directorate
Bolling Air Force Base
Washington, DC 20332

Dr. Robert Ahlers
Code N711
Human Factors Laboratory
Naval Training Systems Center
Orlando, FL 32813

Dr. Ed Aiken
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Earl A. Alluisi
HQ, AFHRL (AFSC)
Brooks AFB, TX 78235

Dr. James Anderson
Brown University
Center for Neural Science
Providence, RI 02912

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Nancy S. Anderson
Department of Psychology
University of Maryland
College Park, MD 20742

Dr. Steve Andriole
George Mason University
School of Information
Technology & Engineering
4400 University Drive
Fairfax, VA 22030

Technical Director, ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Gary Aston-Jones
Department of Biology
New York University
1009 Main Bldg
Washington Square
New York, NY 10003

Dr. Alan Baddeley
Medical Research Council
Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EF
ENGLAND

Dr. Patricia Baggett
University of Colorado
Department of Psychology
Box 345
Boulder, CO 80309

Dr. Jackson Beatty
Department of Psychology
University of California
Los Angeles, CA 90024

Dr. Isaac Bejar
Educational Testing Service
Princeton, NJ 08450

Leo Beltracchi
United States Nuclear
Regulatory Commission
Washington DC 20555

Dr. Gautam Biswas
Department of Computer Science
University of South Carolina
Columbia, SC 29208

Dr. Alvah Bittner
Naval Biodynamics Laboratory
New Orleans, LA 70189

Dr. John Black
Teachers College
Columbia University
525 West 121st Street
New York, NY 10027

Distribution List [UCSD/Rumelhart] NR 667-548

Dr. Arthur S. Blalwes
Code N711
Naval Training Systems Center
Orlando, FL 32813

Dr. R. Darrell Bock
University of Chicago
NORC
6030 South Ellis
Chicago, IL 60637

Dr. Gordon H. Bower
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Robert Breaux
Code N-095R
Naval Training Systems Center
Orlando, FL 32813

Dr. John S. Brown
XEROX Palo Alto Research
Center
3333 Coyote Road
Palo Alto, CA 94304

Dr. Bruce Buchanan
Computer Science Department
Stanford University
Stanford, CA 94305

Joanne Capper
Center for Research into Practice
1718 Connecticut Ave., N.W.
Washington, DC 20009

Dr. Jaime Carbonell
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Gail Carpenter
Northeastern University
Department of Mathematics, 504LA
360 Huntington Avenue
Boston, MA 02115

Dr. Pat Carpenter
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

LCDR Robert Carter
Office of the Chief
of Naval Operations
OP-018
Pentagon
Washington, DC 20350-2000

Chair, Department of
Psychology
College of Arts and Sciences
Catholic University of
America
Washington, DC 20064

Dr. Alphonse Chapanis
8415 Bellona Lane
Suite 210
Buxton Towers
Baltimore, MD 21204

Dr. Paul R. Chatelier
OUSDR
Pentagon
Washington, DC 20350-2000

Dr. Michelene Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Mr. Raymond E. Christal
AFHRL/MOE
Brooks AFB, TX 78235

Dr. William Clancey
Stanford University
Knowledge Systems Laboratory
701 Welch Road, Bldg. C
Palo Alto, CA 94304

Dr. David E. Clement
Department of Psychology
University of South Carolina
Columbia, SC 29208

Chief of Naval Education
and Training
Liaison Office
Air Force Human Resource Lab
Operations Training Division
Williams AFB, AZ 85224

ONR DISTRIBUTION LIST

ONR DISTRIBUTION LIST

Distribution List [UCSD/Rumelhart] NR 667-548

Distribution List [UCSD/Rumelhart] NR 667-548

Assistant Chief of Staff
for Research, Development,
Test, and Evaluation
Naval Education and
Training Command (N-5)
NAS Pensacola, FL 32508

Dr. Michael Coles
University of Illinois
Department of Psychology
Champaign, IL 61820

Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Stanley Collyer
Office of Naval Technology
Code 222
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Leon Cooper
Brown University
Center for Neural Science
Providence, RI 02912

Dr. Lynn A. Cooper
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

CAPT P. Michael Curran
Office of Naval Research
800 N. Quincy St.
Code 125
Arlington, VA 22217-5000

Brian Dallman
3400 ITW/TTCX
Lowry AFB, CO 80230-5000

Dr. Joel Davis
Office of Naval Research
Code 1141NP
800 North Quincy Street
Arlington, VA 22217-5000

LT John Deaton
ONR Code 125
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Natalie Dehn
Department of Computer and
Information Science
University of Oregon
Eugene, OR 97403

Dr. Gerald F. DeJong
Artificial Intelligence Group
Coordinated Science Laboratory
University of Illinois
Urbana, IL 61801

Dr. R. K. Dismukes
Associate Director for Life Sciences
AFOSR
Boiling AFB
Washington, DC 20332

Dr. Emanuel Donchin
University of Illinois
Department of Psychology
Champaign, IL 61820

Defense Technical
Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC
(12 Copies)

Edward E. Eddowes
CNATRA N301
Naval Air Station
Corpus Christi, TX 78419

Dr. Jeffrey Elman
University of California,
San Diego
Department of Linguistics, C-008
La Jolla, CA 92093

Dr. Randy Engle
Department of Psychology
University of South Carolina
Columbia, SC 29208

Dr. William Epstein
University of Wisconsin
W. J. Brogden Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda, MD 20014

Dr. K. Anders Ericsson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. Marshall J. Farr
2520 North Vernon Street
Arlington, VA 22207

Dr. Pat Federico
Code 511
NPRDC
San Diego, CA 92152-6800

Dr. Jerome A. Feldman
University of Rochester
Computer Science Department
Rochester, NY 14627

Dr. Paul Felcovich
Southern Illinois University
School of Medicine
Medical Education Department
P.O. Box 3926
Springfield, IL 62708

Dr. Craig I. Fields
ARPA
1400 Wilson Blvd.
Arlington, VA 22209

J. D. Fletcher
9931 Corsica Street
Vienna VA 22180

Dr. Jane M. Flynn
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Kenneth D. Forbus
University of Illinois
Department of Computer Science
1304 West Springfield Avenue
Urbana, IL 61801

Dr. Barbara A. Fox
University of Colorado
Department of Linguistics
Boulder, CO 80309

Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

Dr. Michael Friendly
Psychology Department
York University
Toronto ONT
CANADA M3J 1P3

Julie A. Gadsden
Information Technology
Applications Division
Admiralty Research Establishment
Portsmouth, Portsmouth PO6 4AA
UNITED KINGDOM

Dr. Michaela Gallagher
University of North Carolina
Department of Psychology
Chapel Hill, NC 27514

Dr. R. Edward Geiselman
Department of Psychology
University of California
Los Angeles, CA 90024

Dr. Michael Genesereth
Stanford University
Computer Science Department
Stanford, CA 94305

Dr. Dedre Gentner
University of Illinois
Department of Psychology
603 E. Daniel St.
Champaign, IL 61820

Distribution List [UCSD/Rumelhart] NR 667-548

Distribution List [UCSD/Rumelhart] NR 667-548

ONR DISTRIBUTION LIST

Dr. Don Gentner
Center for Human
Information Processing
University of California
La Jolla, CA 92093

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Gene L. Glove
Office of Naval Research
Detachment
1030 E. Green Street
Pasadena, CA 91106-2485

Dr. Sam Glucksberg
Department of Psychology
Princeton University
Princeton, NJ 08540

Dr. Daniel Gopher
Industrial Engineering
& Management
TECHNION
Haifa 32000
ISRAEL

Dr. Sherrie Gott
AFHRL/WDJ
Brooks AFB, TX 78235

Jordan Grafman, Ph.D.
2021 Lyttonsaville Road
Silver Spring, MD 20910

Dr. Richard H. Granger
Department of Computer Science
University of California, Irvine
Irvine, CA 92717

Dr. Steven Grant
Department of Biology
New York University
1009 Main Bldg
Washington Square
New York, NY 10003

Dr. Wayne Gray
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. James G. Greene
University of California
Berkeley, CA 94720

Dr. William Greenough
University of Illinois
Department of Psychology
Champaign, IL 61820

Dr. Stephen Grossberg
Center for Adaptive Systems
Room 244
111 Cunningham Street
Boston University
Boston, MA 02215

Dr. Muhammad K. Habib
University of North Carolina
Department of Biostatistics
Chapel Hill, NC 27514

Dr. Henry M. Haliff
Haliff Resources, Inc.
4918 33rd Road, North
Arlington, VA 22207

Dr. Cheryl Hamel
NTSC
Orlando, FL 32813

Dr. Bruce W. Hamill
Johns Hopkins University
Applied Physics Laboratory
Johns Hopkins Road
Laurel, MD 20707

Dr. Ray Hannapel
Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

Stevan Harnad
Editor, The Behavioral and
Brain Sciences
20 Nassau Street, Suite 240
Princeton, NJ 08540

Dr. Steven A. Hillyard
Department of Neurosciences
University of California,
San Diego
La Jolla, CA 92093

Dr. Geoffrey Hinton
Carnegie-Mellon University
Computer Science Department
Pittsburgh, PA 15213

Dr. Jim Hollan
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. John Holland
University of Michigan
2313 East Engineering
Ann Arbor, MI 48109

Dr. Melissa Holland
Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Keith Holyoak
University of Michigan
Human Performance Center
330 Packard Road
Ann Arbor, MI 48109

Dr. James Howard
Dept. of Psychology
Human Performance Laboratory
Catholic University of
America
Washington, DC 20064

Dr. Earl Hunt
Department of Psychology
University of Washington
Seattle, WA 98105

Dr. Ed Hutchins
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. Alice Isen
Department of Psychology
University of Maryland
Catonsville, MD 21228

Dr. Robert Jannarone
Department of Psychology
University of South Carolina
Columbia, SC 29208

COL Dennis W. Jarvi
Commander
AFHRL
Brooks AFB, TX 78235-5601

Dr. Robin Jeffries
Hewlett-Packard Laboratories
P.O. Box 10490
Palo Alto, CA 94303-0971

Chair, Department of
Psychology
The Johns Hopkins University
Baltimore, MD 21218

CDR Tom Jones
ONR Code 125
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Douglas H. Jones
Thatcher Jones Associates
P.O. Box 6640
10 Trafalgar Court
Lawrenceville, NJ 08648

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Daniel Kahneman
The University of British Columbia
Department of Psychology
4154-2053 Main Mall
Vancouver, British Columbia
CANADA V6T 1Y7

Dr. Demetrios Karis
Grumman Aerospace Corporation
MS C04-14
Bethpage, NY 11714

ONR DISTRIBUTION LIST

Distribution List [UCSD/Rumelhart] NR 667-548

Distribution List [UCSD/Rumelhart] NR 667-548

Dr. Milton S. Katz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Steven W. Keele
Department of Psychology
University of Oregon
Eugene, OR 97403

Dr. Scott Kelso
Haskins Laboratories,
270 Crown Street
New Haven, CT 06510

Dr. Dennis Kibler
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. David Kieras
University of Michigan
Technical Communication
College of Engineering
1223 E. Engineering Building
Ann Arbor, MI 48109

Dr. David Klahr
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Mazie Knerr
Program Manager
Training Research Division
HUMRO
1100 S. Washington
Alexandria, VA 22314

Dr. Ronald Knoll
Bell Laboratories
Murray Hill, NJ 07974

Dr. Sylvan Kornblum
University of Michigan
Mental Health Research Institute
205 Washtenaw Place
Ann Arbor, MI 48109

Dr. Stephen Kosslyn
Harvard University
1236 William James Hall
33 Kirkland St.
Cambridge, MA 02138

Dr. Kenneth Kotovsky
Department of Psychology
Community College of
Allegheny County
800 Allegheny Avenue
Pittsburgh, PA 15233

Dr. David H. Krantz
2 Washington Square Village
Apt. # 15J
New York, NY 10012

Dr. Benjamin Kuipers
University of Texas at Austin
Department of Computer Sciences
T.S. Painter Hall 3.28
Austin, Texas 78712

Dr. David R. Lambert
Naval Ocean Systems Center
Code 441T
271 Catalina Boulevard
San Diego, CA 92152-6800

Dr. Pat Langley
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. Marcy Lansman
University of North Carolina
The L. L. Thurstone Lab.
Davis Hall 013A
Chapel Hill, NC 27514

Dr. Jill Larkin
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Robert Lawler
Information Sciences, FRL
GTE Laboratories, Inc.
40 Sylvan Road
Waltham, MA 02254

Dr. Paul E. Lehner
PAR Technology Corp.
7926 Jones Branch Drive
Suite 170
McLean, VA 22102

Dr. Alan M. Lesgold
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Alan Lechner
Deputy Division Director
Behavioral and Neural Sciences
National Science Foundation
1800 G Street
Washington, DC 20550

Dr. Jim Levin
University of California
Laboratory for Comparative
Human Cognition
D003A
La Jolla, CA 92093

Dr. Michael Levine
Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801

Dr. Clayton Lewis
University of Colorado
Department of Computer Science
Campus Box 430
Boulder, CO 80309

Library
Naval War College
Newport, RI 02940

Library
Naval Training Systems Center
Orlando, FL 32813

Dr. Bob Lloyd
Dept. of Geography
University of South Carolina
Columbia, SC 29208

Dr. Gary Lynch
University of California
Center for the Neurobiology of
Learning and Memory
Irvine, CA 92717

Dr. Don Lyon
P. O. Box 44
Higley, AZ 85236

Dr. William L. Maloy
Chief of Naval Education
and Training
Naval Air Station
Pensacola, FL 32508

Dr. Evans Mandes
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Sandra P. Marshall
Dept. of Psychology
San Diego State University
San Diego, CA 92182

Dr. Manton M. Matthews
Department of Computer Science
University of South Carolina
Columbia, SC 29208

Dr. Richard E. Mayer
Department of Psychology
University of California
Santa Barbara, CA 93106

Dr. James McBride
Psychological Corporation
c/o Harcourt, Brace,
Jovanovich Inc.
1250 West 6th Street
San Diego, CA 92101

Dr. Jay McClelland
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Distribution List [UCSD/Rumelhart] NR 667-548

Dr. James L. McLaugh
Center for the Neurobiology
of Learning and Memory
University of California, Irvine
Irvine, CA 92717

Dr. Gail McKoon
CAS/Psychology
Northwestern University
1859 Sheridan Road
Evanston, IL 60201

Dr. Joe McLachlan
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. James McMichael
Assistant for MPT Research,
Development, and Studies
OP 0187
Washington, DC 20370

Dr. Douglas L. Medin
Department of Psychology
University of Illinois
603 E. Daniel Street
Champaign, IL 61820

Dr. Arthur Meined
U. S. Department of Education
724 Brown
Washington, DC 20208

Dr. Al Meyrowitz
Office of Naval Research
Code 1133
800 N. Quincy
Arlington, VA 22217-5000

Dr. George A. Miller
Department of Psychology
Green Hall
Princeton University
Princeton, NJ 08540

Dr. Andrew R. Molnar
Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

Dr. William Montague
NPRDC Code 13
San Diego, CA 92152-6800

Dr. Tom Moran
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

Dr. Allen Munro
Behavioral Technology
Laboratories - USC
1845 S. Elena Ave., 4th Floor
Redondo Beach, CA 90277

Dr. David Navon
Institute for Cognitive Science
University of California
La Jolla, CA 92093

Dr. Allen Newell
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Mary Jo Nissen
University of Minnesota
N218 Elliott Hall
Minneapolis, MN 55455

Dr. Donald A. Norman
Institute for Cognitive Science
University of California
La Jolla, CA 92093

Director, Training Laboratory,
NPRDC (Code 05)
San Diego, CA 92152-6800

Director, Manpower and Personnel
Laboratory,
NPRDC (Code 06)
San Diego, CA 92152-6800

Director, Human Factors
& Organizational Systems Lab,
NPRDC (Code 07)
San Diego, CA 92152-6800

Fleet Support Office,
NPRDC (Code 301)
San Diego, CA 92152-6800

Distribution List [UCSD/Rumelhart] NR 667-548

Library, NPRDC
Code P201L
San Diego, CA 92152-6800

Commanding Officer,
Naval Research Laboratory
Code 2627
Washington, DC 20390

Dr. Harold F. O'Neill, Jr.
School of Education - MPH 801
Department of Educational
Psychology & Technology
University of Southern California
Los Angeles, CA 90089-0031

Dr. Michael Oberlin
Naval Training Systems Center
Code 711
Orlando, FL 32813-7100

Dr. Stellan Ohlsson
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Mathematics Group,
Office of Naval Research
Code 1111MA
800 North Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1133
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1141NP
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142
800 N. Quincy St.
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142EP
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142PT
800 N. Quincy Street
Arlington, VA 22217-5000
(6 Copies)

Psychologist
Office of Naval Research
Branch Office, London
Box 39
FPO New York, NY 09510

Special Assistant for Marine
Corps Matters,
ONR Code 00MC
800 N. Quincy St.
Arlington, VA 22217-5000

Psychologist
Office of Naval Research
Liaison Office, Far East
APO San Francisco, CA 96503

Dr. Judith Orasanu
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Jesse Orlansky
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311

Dr. Robert F. Pasnak
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Daira Paulson
Code 52 - Training Systems
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. James W. Pellegrino
University of California,
Santa Barbara
Department of Psychology
Santa Barbara, CA 93106

ONR DISTRIBUTION LIST

Distribution List [UCSD/Rumelhart] NR 667-548

Dr. Ray Perez
ARI (PERI-11)
5001 Eisenhower Avenue
Alexandria, VA 22333

Department of Computer Science,
Naval Postgraduate School
Monterey, CA 93940

Dr. Steven Pinker
Department of Psychology
E10-018
M.I.T.
Cambridge, MA 02139

Dr. Martha Polson
Department of Psychology
Campus Box 346
University of Colorado
Boulder, CO 80309

Dr. Peter Polson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. Michael I. Posner
Department of Neurology
Washington University
Medical School
St. Louis, MO 63110

Dr. Mary C. Potter
Department of Psychology
MIT (E-10-032)
Cambridge, MA 02139

Dr. Karl Pribram
Stanford University
Department of Psychology
Bldg. 4201 -- Jordan Hall
Stanford, CA 94305

Dr. Joseph Psotka
ATTN: PERI-1C
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. Lynne Reder
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. James A. Reggia
University of Maryland
School of Medicine
Department of Neurology
22 South Greene Street
Baltimore, MD 21201

Dr. Ernst Z. Rothkopf
AT&T Bell Laboratories
Room 2D-456
600 Mountain Avenue
Murray Hill, NJ 07974

Dr. William B. Rouse
Search Technology, Inc.
25-b Technology Park/Atlanta
Norcross, GA 30092

Dr. Donald Rubin
Statistics Department
Science Center, Room 608
1 Oxford Street
Harvard University
Cambridge, MA 02138

Dr. David Rumelhart
Center for Human
Information Processing
Univ. of California
La Jolla, CA 92093

Dr. E. L. Saltzman
Haskins Laboratories
270 Crown Street
New Haven, CT 06510

Dr. Fumiko Samejima
Department of Psychology
University of Tennessee
Knoxville, TN 37916

Dr. Arthur Samuel
Yale University
Department of Psychology
Box 11A, Yale Station
New Haven, CT 06520

Distribution List [UCSD/Rumelhart] NR 667-548

Dr. Walter Schneider
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Janet Schofield
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Robert J. Seidel
US Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. T. B. Sheridan
Dept. of Mechanical Engineering
MIT
Cambridge, MA 02139

Dr. Herbert A. Simon
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

ITCOL Robert Simpson
Defense Advanced Research
Projects Administration
1400 Wilson Blvd.
Arlington, VA 22209

Dr. Linda B. Smith
Department of Psychology
Indiana University
Bloomington, IN 47405

Dr. Robert F. Smith
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Richard E. Snow
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Kathryn T. Spoehr
Brown University
Department of Psychology
Providence, RI 02912

Dr. Ted Steinke
Dept. of Geography
University of South Carolina
Columbia, SC 29208

Dr. Saul Sternberg
University of Pennsylvania
Department of Psychology
3815 Walnut Street
Philadelphia, PA 19104

Dr. Albert Stevens
Bolt Beranek & Newman, Inc.
10 Moulton St.
Cambridge, MA 02238

Dr. Paul J. Sticha
Senior Staff Scientist
Training Research Division
HumRRO
1100 S. Washington
Alexandria, VA 22314

Dr. Steve Suomi
NIH Bldg. 31
Room 828-15
Bethesda, MD 20205

Dr. Patrick Suppes
Stanford University
Institute for Mathematical
Studies in the Social Sciences
Stanford, CA 94305

Dr. John Tangney
AFOSR/NL
Bolling AFB, DC 20332

Dr. Richard F. Thompson
Stanford University
Department of Psychology
Bldg. 4201 -- Jordan Hall
Stanford, CA 94305

Chair, Department of
Computer Science
Towson State University
Towson, MD 21204

Dr. Michael T. Turvey
Haskins Laboratories
270 Crown Street
New Haven, CT 06510

Distribution List [UCSD/Rumelhart] NR 667-548

Dr. Amos Tversky
Stanford University
Dept. of Psychology
Stanford, CA 94305

Dr. James Tweeddale
Technical Director
Navy Personnel R&D Center
San Diego, CA 92152-6800

Headquarters, U. S. Marine Corps
Code MPI-20
Washington, DC 20380

Dr. William Uttal
MOSC, Hawaii Lab
Box 997
Kailua, HI 96734

Dr. Kurt Van Lehn
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Norman M. Weinberger
University of California
Center for the Neurobiology
of Learning and Memory
Irvine, CA 92717

Dr. Shih-Sung Wen
Jackson State University
1325 J. R. Lynch Street
Jackson, MS 39217

Dr. Keith T. Wescourt
FMC Corporation
Central Engineering Labs
1185 Coleman Ave., Box 580
Santa Clara, CA 95052

Dr. Douglas Wetzel
Code 12
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Barry Whitset
University of North Carolina
Department of Physiology
Medical School
Chapel Hill, NC 27514

Dr. Christopher Wickens
Department of Psychology
University of Illinois
Champaign, IL 61820

Dr. Heather Wild
Naval Air Development
Center
Code 6021
Warminster, PA 18974-5000

Dr. Michael Williams
IntelliCorp
1975 El Camino Real West
Mountain View, CA 94040-2216

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

Dr. Martin F. Wiskoff
Navy Personnel R & D Center
San Diego, CA 92152-6800

Dr. Donald Woodward
Office of Naval Research
Code 1141NP
800 North Quincy Street
Arlington, VA 22217-5000

Dr. Joe Yasatuke
AFHRL/LRT
Lowry AFB, CO 80230

Dr. Masoud Yazdani
Dept. of Computer Science
University of Exeter
Exeter EX4 4QL
Devon, ENGLAND

Mr. Carl York
System Development Foundation
181 Lytton Avenue
Suite 210
Palo Alto, CA 94301

Dr. Joseph L. Young
Memory & Cognitive
Processes
National Science Foundation
Washington, DC 20550

Dr. Steven Zornetzer
Office of Naval Research
Code 1140
800 N. Quincy St.
Arlington, VA 22217-5000

Dr. Michael J. Zyda
Naval Postgraduate School
Code 52CK
Monterey, CA 93943-5100

END

DTIC

10-86