# Naval Research Reviews

Office of Naval Research
Four / 1986
Vol XXXVIII

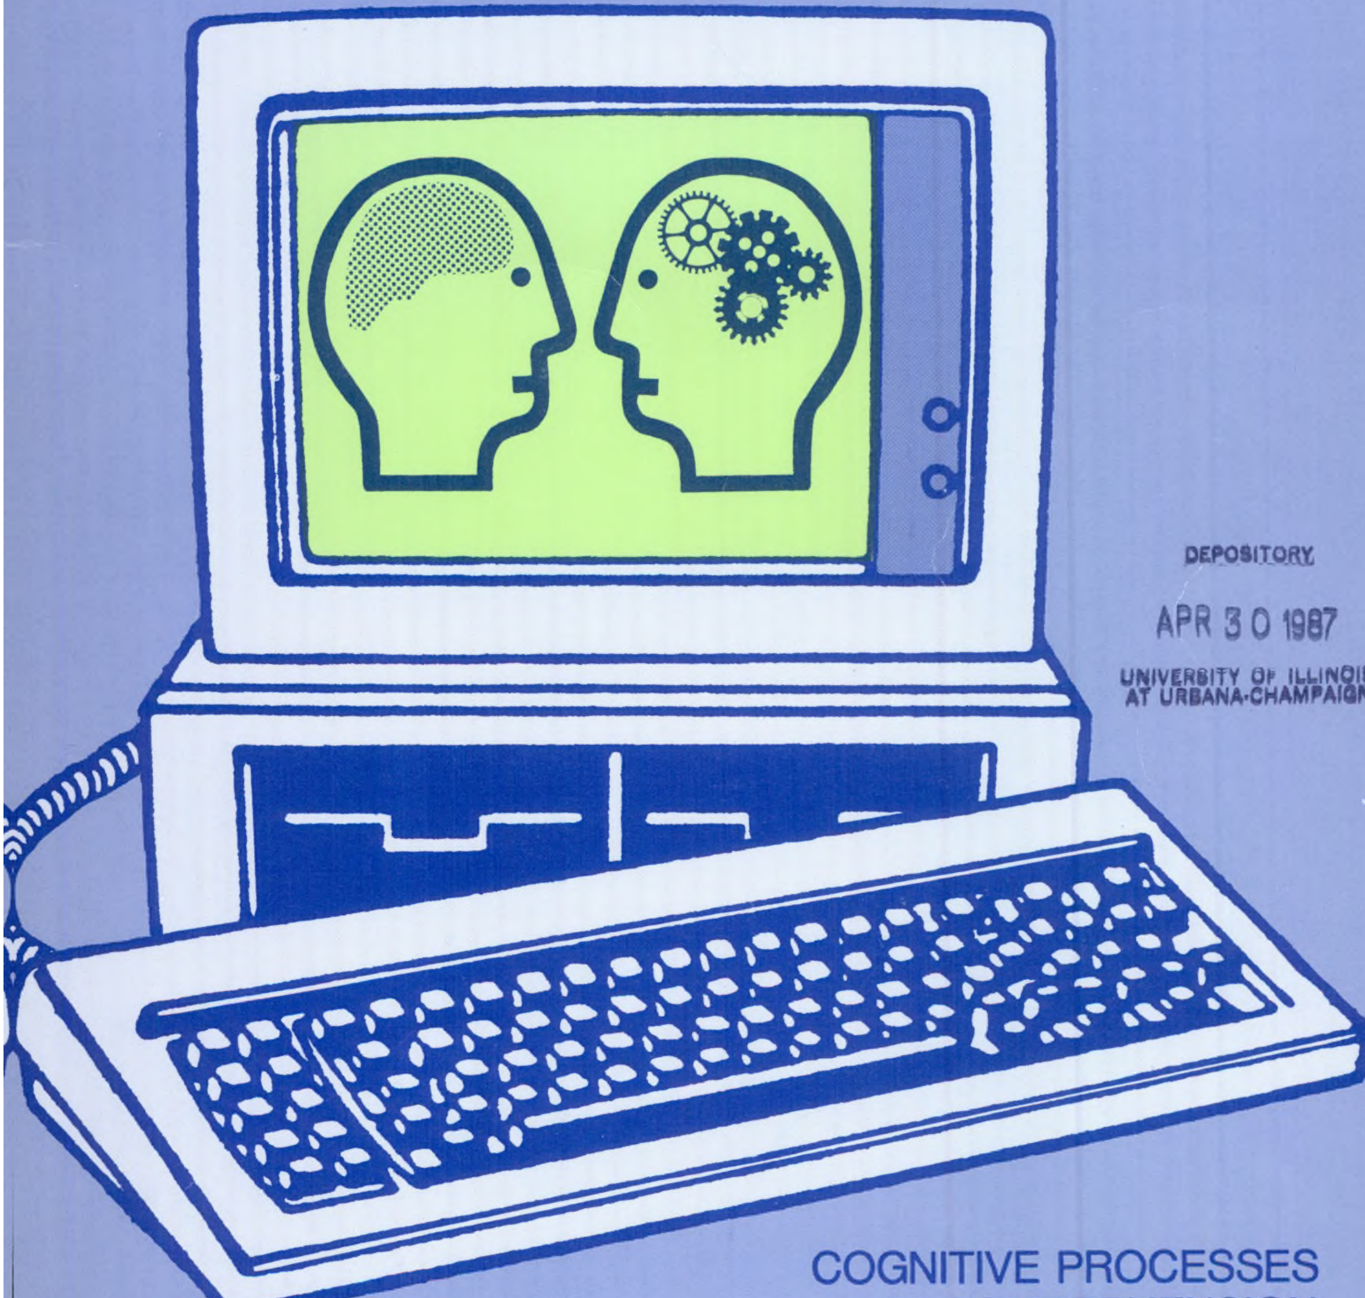



COGNITIVE PROCESSES
IN COMPREHENSION

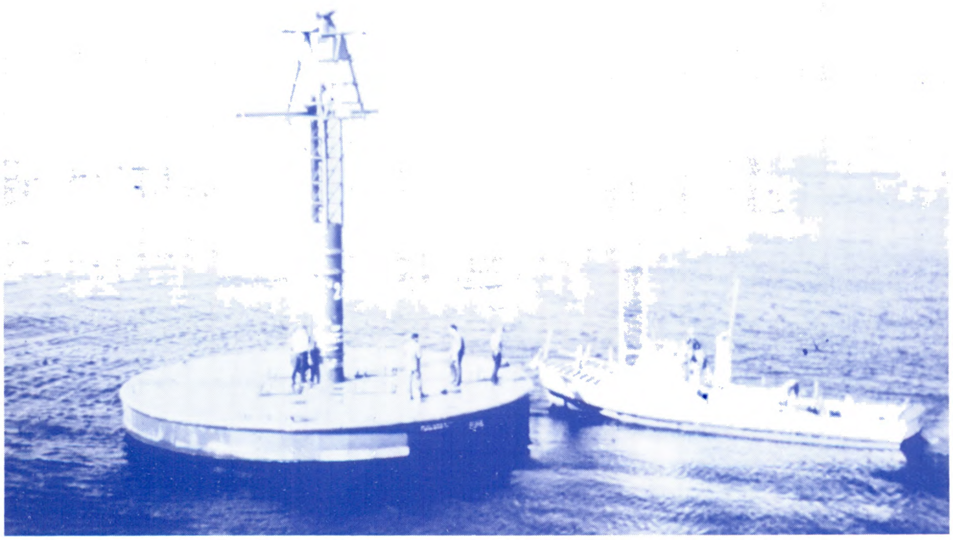

# The Monster Buoy

The first of the "monster buoys" which was designed by a committee of scientists from the Office of Naval Research and fabricated at the Convair Division of General Dynamics, San Diego, in 1960. This buoy, was successfully tested in 1965 in the Gulf Stream and withstood the fury of Hurricane Betsy. Gathering synoptic data over broad geographical areas for long periods of time was the fundamental mission of the buoy, which has proven an invaluable aid for oceanography and weather prediction. Measuring 40 feet in diameter and weighing 100 tons, such ocean stations were capable of operating unattended for more than a year.

These buoys were equipped with a variety of sensors to measure wind direction and speed, precipitation, radiation, barometric pressure and compass heading. Instruments inside the hull record water temperatures at the surface, wave height, period and direction. Clamped along the mooring line, other instruments measured water temperatures, pressure and salinity and determined the structure of subsurface layers. Data stored aboard the buoy in a magnetic core memory is transmitted every six hours to a shore-based station. "Monster buoys" were built for the U.S.—most of them for the National Oceanographic and Atmospheric Administration; others were built for international users.

Over the years, ONR has continued to develop techniques and systems for collecting oceanographic data such as remote sensing and tethered and untethered profilers, and long-term oceanographic moorings. Today, the "Monster Buoy" remains a reliable workhorse for meteorologists and oceanographers. These buoys are still moored in the oceans of the world gathering and transmitting data for daily weather prediction and general information about the ocean.

# Naval Research Reviews

**Office of Naval Research**
**Four / 1986**
**Vol XXXVIII**


## Articles



## Departments







**About the Cover**

A graphic design depicting cognitive science which is the focus of the article about The Personnel and Research Program of the Office of Naval Research on page 2.


*Naval Research Reviews* publishes articles about research conducted by the laboratories and contractors of the Office of Naval Research and describes important naval experimental activities. Manuscripts submitted for publication, correspondence concerning prospective articles, and changes of address, should be directed to Code OPARI, Office of Naval Research, Arlington, Va. 22217-5000. Requests for subscriptions should be directed to the Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402. *Naval Research Reviews* is published from appropriated funds by authority of the Office of Naval Research in accordance with Navy Publications and Printing Regulations. NAVSO P-35.

# Personnel and Training Research Program: Cognitive Science at ONR

Susan F. Chipman,
Charles Davis,
Michael G. Shafto

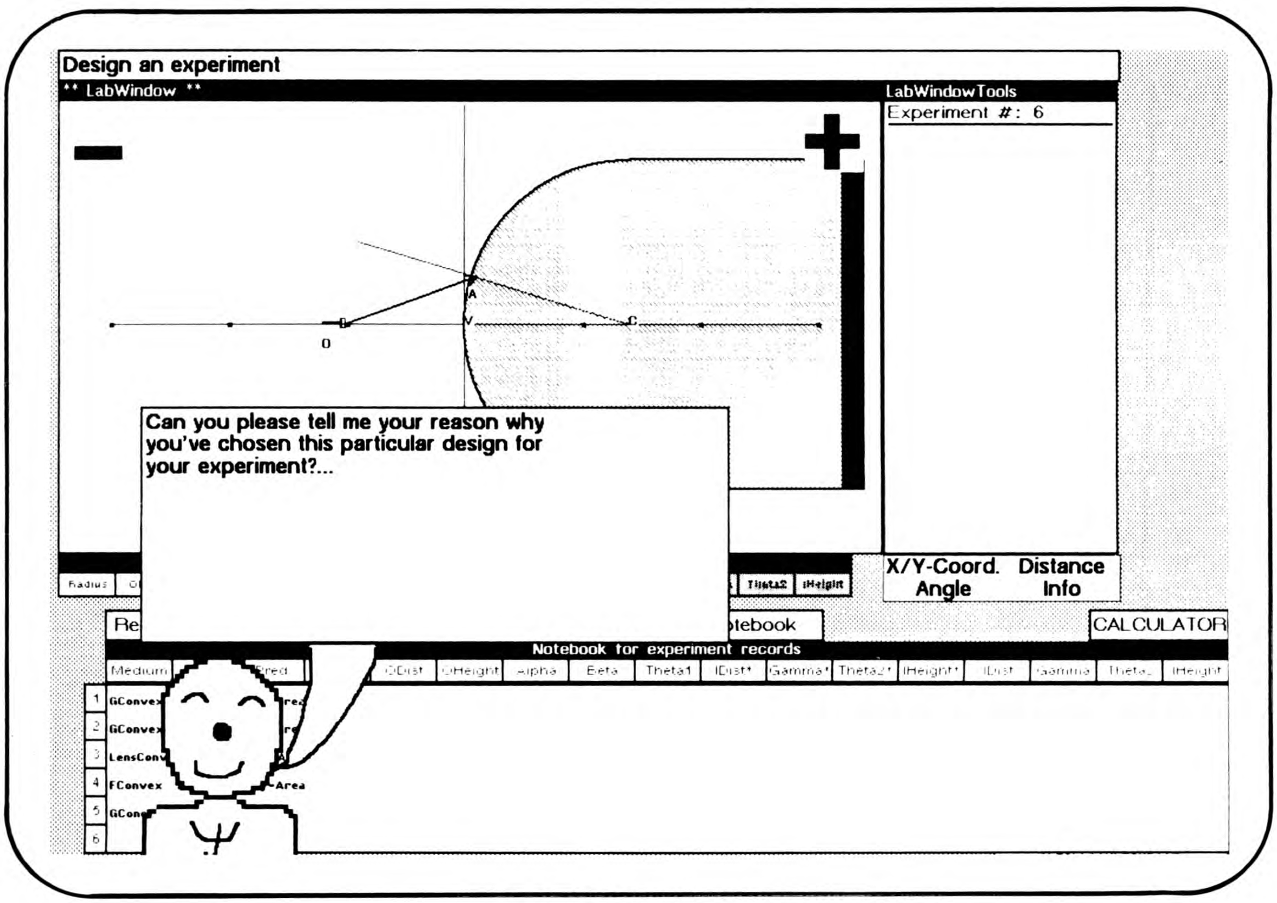

# Introduction

The Personnel and Training Research Program at ONR focuses, as the long-term goal, on increasing the effectiveness of the people who run the Navy and the Marine Corps. In 1985, Secretary Weinberger testified in his annual report to Congress, "The single most important indicator of the strength of any organization is its people—their competence, commitment, and morale. It is these professionals . . . who constitute the core of our military strength." Most of ONR's research programs are concerned with improving Naval hardware; this program aims to ensure that the personnel of the Navy and Marine Corps will be able to keep that hardware functioning and to use it effectively to the limits of its capabilities.

Training in the Navy is itself an enterprise with a budget of more than $5 Billion annually. The Personnel and Training Research program represents an investment of about one tenth of 1% in basic research to improve that enterprise. There is every reason to believe that investment can be repaid many times over, through improvements in efficiency but especially through improvements in effectiveness. Improved training techniques can cut training time and therefore training costs. For example, it has been shown that conventional computer-assisted instruction can reduce training time by 30% while maintaining equal effectiveness. Below, ONR's effort in the newer technology of artificially intelligent computerized tutors will be described. The aim of this research is to approximate the striking instructional effectiveness of individual tutoring by human tutors. It has been shown that individual tutoring typically increases student and trainee performance levels by two standard deviations, as compared to conventional group classroom instruction. That is the equivalent of increasing an SAT score from 500 to 700.

There is ample room for improvement. Dropout rates of 30% are not uncommon in Navy courses, a costly waste of resources and a discouraging personal experience for many individuals. In part, this problem can be addressed through improved selection—one of the concerns of this program—but the declining size and changing composition of the youth population dictate that large improvements in training effectiveness are essential if the manpower needs of the high technology, 600 ship Navy are to be met. At the same time, the Navy's jobs are characterized by increasingly high technology demanding more capable and more highly trained personnel.

As Figure 1 indicates, there is ample room for improvement in the performance of successful trainees as well. This graph shows the rate of success of technicians with varying degrees of training and experience in locating which of about 20 major modules in the Mark 45 Gun was defective and in need of repair. If the knowledge of expert technicians could be more effectively analyzed and transmitted through training—a general strategy advocated by this research program—substantial savings could be realized. Maintenance failures generate unnecessary replacement costs and reduce the effective size of the force. Comparable shortcomings in operational training result in the underutilization of the technical capabilities of equipment and may disable or destroy equipment and endanger lives.

**Figure 1**

*Locating Defective Module of Mark 45 Gun. This graph shows the percentage of the time that gunner's mates with varying degrees of training and experience were able to correctly identify which major module of the Mark 45 gun was defective. These data come from a study conducted by* NPRDC.

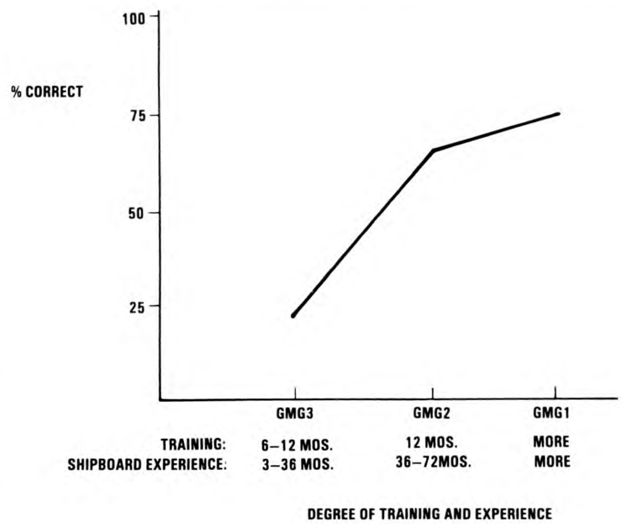


# Scientific Issues

In order to achieve these longterm goals, the Personnel and Training Research Program focuses on scientific issues in four primary areas:

- ☐ Fundamental Characteristics of the Human Learner and Performer
- ☐ The Nature of Learned Skills and Knowledge
- ☐ Theory for Effective Instruction
- ☐ Mathematical Models for Measuring Individual Characteristics, Skill and Knowledge: the theoretical basis for testing.

If we understand what we have to build with—the fundamental capabilities of the human learner—and what has to be built—the detailed nature of the knowledge and skills that are the goals of instruction, then we should be able to figure out how to build it—a theory for instruction.

## Distinctive Character of the ONR Program

The issues outlined above are of broad and general interest in psychological research. But the ONR research program has a very distinctive character among basic research programs in psychology, not coincidentally related to its purpose in contributing to Navy goals.

The program is characterized by an emphasis on understanding complex performance, as contrasted to the very simplified experimental situations that are typical in psychological research. In his 1985 keynote address to the inaugural meeting of a European Society for Cognitive Psychology, Daniel Broadbent described a fundamental dichotomy in the theoretical positions taken by basic researchers and theorists in psychology. Some believe that it will be possible to analyze human psychological functioning into "invariant atoms" that can be studied in isolation in very simplified experimental situations. Others believe that human psychological functioning is always fundamentally complex, displaying rich interconnections among different functions and strong dependency upon the particular context of performance. They doubt that the "atoms" will suffice to account for complex performances. Scientific episodes like the discovery that people employ complex strategies involving visual images and verbal games in order to cope with apparently pure and simple learning experiments that ask them to remember associations between pairs of nonsense syllables have tended to cast doubt upon the atomistic approach. Clearly, the ONR program has placed its bets on the challenging effort to deal with the full complexity of human performance. This fact is intimately related to its second distinctive characteristic.

The ONR program is characterized by a cognitive science approach. This represents the tools that make it possible to attempt scientific explanation of complex human performances. Computer simulations are used extensively to express and test complex psychological theories. This includes the extensive use of artificial intelligence programs as psychological models. Artificial intelligence and psychological theory have in common the goal of having models that can produce—and therefore account for—intelligent behavior. The critical difference is that psychologists are concerned with understanding the particular way—among all possibilities—that humans actually do perform intelligently (or sometimes not so intelligently). ONR has had a strong leadership role in the emergence of the exciting new field of cognitive science, in which these two concerns come together. ONR contractors such as David Kieras at the University of Michigan and Kurt VanLehn at Carnegie Mellon University are also taking a leading role in thinking about how computer simulations should be formulated in order to rigorously express and test psychological theories.

The ONR Personnel and Training Research Program is also characterized by a strong interest in providing the research foundations for advanced instructional technology, particularly the use of artificial intelligence techniques in instruction. ONR has been much more willing than other sources of support to invest in research with very long-term payoff, to explore the theoretical limits of what can be done, without concern that computer equipment involved might be too expensive for practical training use. Of course, this strategy is being vindicated by the rapidly declining cost of computer equipment. The Navy does have a particularly strong interest in the potential of instructional technology because so many of its personnel are at sea, where the provision of traditional classroom instruction for upgrading and maintaining skills is often impractical. From a scientific point of view, computerized instructional systems are laboratories which provide a degree of control over instructional actions and a data collection capability that make it possible to study the processes of learning and instruction scientifically. The results of this research may be as applicable to conventional instruction as to automated systems, but only automated systems make the conduct of the research practical.

Together, these characteristics define an unusual program of basic research in psychology that can be realistically expected to have significant practical impact on the manpower needs and operational effectiveness of the high technology Navy we will have in the year 2000.

# The Human Learner

This facet of the program aims to understand the human system that we are trying to train. What resources does it have available that we can work with? There are three primary questions about the characteristics of this system:

- ☐ How is the information that is presented encoded and stored for future use?
- ☐ What processes are available within the system to operate on that information?
- ☐ What are the mechanisms of control that govern the selection and operation of those processes?

We shall discuss three lines of work that correspond to each of these major questions:

- ☐ The dynamics of memory in skilled performance.
- ☐ Visuo-spatial cognitive processes
- ☐ Automatic and controlled processing.

This facet of the program shares the primary questions and concerns that have characterized traditional experimental psychology, but there is a particular emphasis on developing those aspects of psychological theory that will help us to understand how the system acquires and performs complex skills. A central challenge is to factor out invariant structural characteristics from the modifications that learning builds upon those characteristics, at the same time achieving an understanding of the way desirable modifications of system characteristics can be brought about. Studies of short-term memory are a prime example of this challenge.

## The Dynamics of Memory in Skilled Performance

Short-term memory is the sort of memory that we use to remember a telephone number between looking it up in the directory and dialing the telephone. As we all know, the capacity of human short-term memory is limited, and information in short-term memory is quickly lost if our attention is distracted. One of the most cited papers in all of psychology is George Miller's paper, "The magical number seven, plus or minus two," (work supported by ONR many years ago). In this paper, he argued that short term memory has a constant capacity in terms of items of information, which he labeled *chunks,* as long as one recognizes that the nature of those items is variable: they might be letters or digits, but they might be words (containing several letters) or well-known dates, or they might even be short sentences. Obviously, the very limited capacity of short-term memory has important implications for the design of tasks, for person-machine interface design, and for instructional design. We have difficulty, for example, when telephone numbers or zip codes have too many digits.

Yet, persons who are expert in complex tasks seem to be able to rapidly take in large amounts of information, and they appear to have rapid access to that information in support of their problem solving activity. The late William Chase of Carnegie-Mellon University, Anders Ericsson, now of the University of Colorado, and Chase's student James Staszewski have been engaged in a program of research aiming to explain how experts overcome the limits of short-term memory. Figure 2, from Ericsson's work is just one of many possible illustrations of the way experts use memory to store large amounts of information in support of performance. It shows the number of seconds an expert chess player required to make a chess move with a chess board physically present or on the basis of memory only when playing blind-fold chess (mental moves). Performance on the basis of information in memory is at least as fast. Chess playing is a particularly well-defined (and thus easier to study) representative example of many tasks that require holding a lot of information in a state of ready access for problem solving.

**Figure 2**

*Time Required to Make Chess Moves. This graph, from a study by Anders Ericsson, shows the number of seconds that an expert chess player required to make each successive move when playing with the board present (physical moves) or playing blindfold chess (mental moves).*

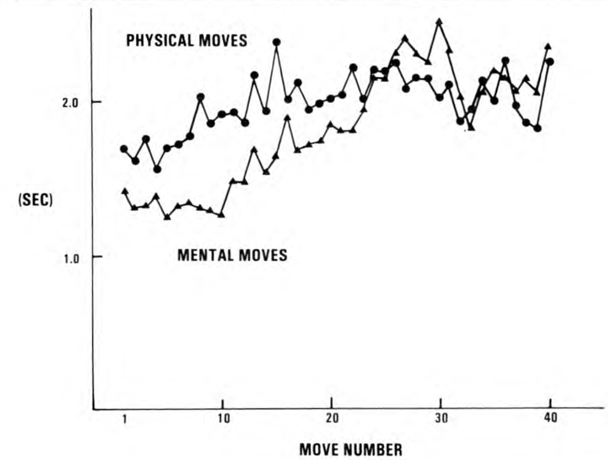


Examples of such tasks include designing complex computer programs, doing tactical planning, or diagnosing illness based upon the information provided by a standard patient interview.

The strategy of this research involves several steps: 1) studying a simplified case of artificial expertise, 2) analyzing the mechanisms that support that performance, 3) testing the applicability of that analysis to some additional simple cases that involve problem solving: mental calculation and blindfold chess, 4) extending the analysis to more realistic tasks, and 5) throughout the research, demonstrating the trainability of the mechanisms behind performance—as contrasted to the hypothesis that apparently exceptional memory performance might reflect special innate abilities.

Partially for dramatic effect, Chase and his associates chose—as their special case of artificial expertise—to create an expert in the task that is ordinarily used to assess a person's short-term memory capacity (which does vary somewhat from person to person). This is the digit span task in which random digits are read to a person at a rate of 1 per second in order to determine how many can be recalled. Figure 3 shows the extraordinary levels of performance that their two "experts", initially college students with typical levels of short-term memory performance, were able to achieve. As shown on the graph, they are very far above previous world-record levels. SF was the most successful of several students who were given a lot of practice and urged to try to remember more digits. DD is a student who was trained to use SF's methods after they had been analyzed, thus demonstrating that it is not a case of special individual ability.

**Figure 3**

*Mean Digit Span For* SF *and* DD *as a Function of Practice. This graph shows the length of the sequence of random digits that two subjects trained by Chase, Ericsson and Staszewski were able to recall as training progressed. The lines across the graph show average human performance and previous world record levels.*

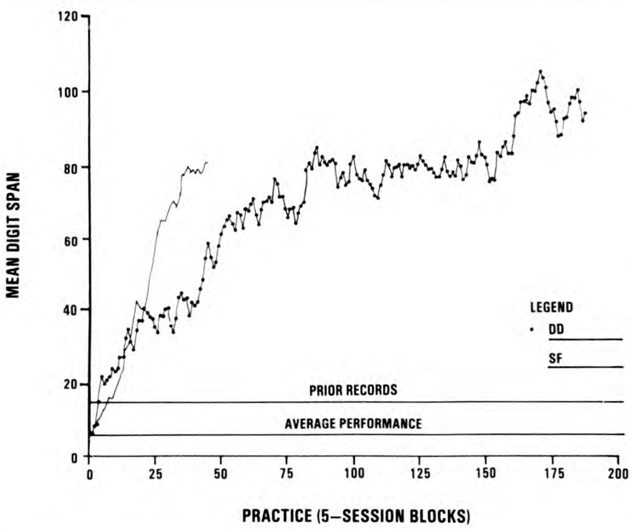


How did these students overcome the normal limits on short-term memory? Short-term memory is coordinated with long-term memory in order to effectively expand capacity. Established knowledge in long term memory is used to code incoming information into larger chunks. For example, many of us would find it easy to remember this particular long sequence of digits: 617212202619213415. In addition, these students created special structures in long term memory that hold information from the task, "retrieval structures". A third factor in their performance is that the speed of the processing done to code the information increases with practice.

Both of these students happened to be long-distance runners; they used their extensive knowledge of past, present and personal record running times to code sequences of digits, as we might use our knowledge of the area codes of major cities to code the sequence above. This tactic alone can enable one to remember 18 to 20 digits. Not only did students report that this was how they were doing it, the use of the coding scheme was confirmed by giving them sequences of digits that were particularly easy or difficult to code in this scheme. Their performance improved or declined as would be expected. The key point about this is that an established structure of knowledge in long term memory serves to finely differentiate and organize the items of information that need to be remembered. Experts have this type of organized knowledge in their field, whatever it may be.

The second factor in the extraordinary performance of these two "digit span experts" was their development of "retrieval structures", a set of abstract locations or bins established in long-term memory in order to support performance in the task. Figure 4 shows the retrieval structure that DD used to recall a sequence of 100 digits. Each of the ticks at the bottom of the tree represents a digit, as exemplified in the blow-up beneath. At the first node above, the type of code such as "a two-mile time" or "an age" is stored. We know that the retrieval structure is in long-term memory because large amounts of the information stored in it can be remembered after a considerable period of time, like an hour, well beyond the limits of short-term memory. The subjects are aware of the retrieval structure, but its existence and organization are also confirmed by data such as the amount of time required to go from one digit to the next, which is greater if the link is at a high level of the tree. There are traces of the limits of short-term memory in the characteristics of the retrieval structure and process. Only 3 to 5 items are grouped at each level of the retrieval structure, and if one stops the expert in the process of recall and asks what he has in mind at that very moment, 3 to 5 items will be reported, which may be either digits or code types.

**Figure 4**

*Subject DD's Retrieval Structure. This is a graphic illustration of the internal mental organization that subject DD used to recall sequences of 100 random digits. As the blow-up illustrates, each tick at the base of the tree represents a digit, while information about the type of code being used for a 3 or 4 digit sequence is recorded at the next higher node in the tree.*

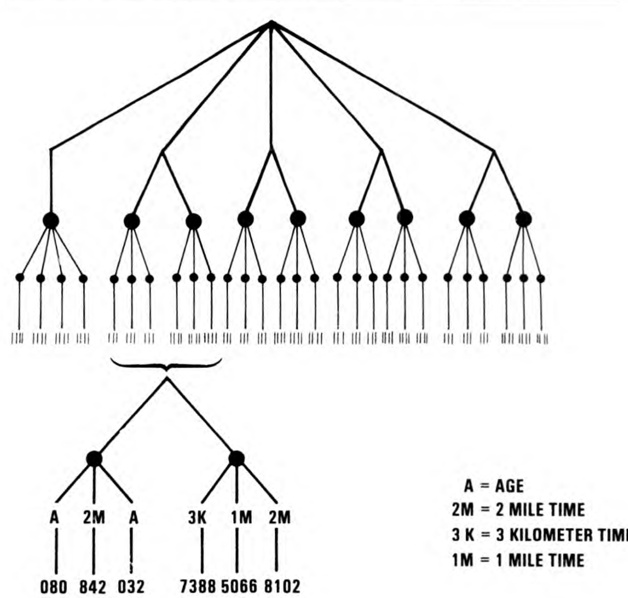

The potential payoff of this work lies in the implication that we can create new structural resources in the human system to enhance the ability to deal with complex and information intensive tasks. By identifying and deliberately training the knowledge that an expert in an area like tactical planning uses to organize and handle in-coming information, we should be able to greatly enhance the speed and efficiency of training new experts. We may even be able to design retrieval structures particularly apt for handling the information germane to a particular task, raising performance to levels never observed before.

Scientifically, this work is at the stage of determining the general applicability of the analysis to a wide range of tasks. In addition, the researchers are moving towards the use of computer simulations to specify the memory processes more finely and more formally, and to link their work with theoretical work like that of Anderson's discussed below.

## Visuo-spatial Cognitive Processes

The second line of research that we are discussing focuses on understanding some of the processes that are available to manipulate information within the human system. The study of the human processing of information that has a visual character – rather than a symbolic character – has been emphasized in ONR's program. This is a topic of high current interest in cognitive science, partially as a counterbalance to past heavy emphasis on verbal and other symbolic modes of thinking. It is also closely related to current work on computational theories of visual perception that interests ONR's Perceptual Science program. Visual thinking has high relevance to many Navy tasks, such as radar interpretation and maintenance tasks.

The work of Patricia Carpenter and Marcel Just of Carnegie-Mellon University builds upon the past accomplishments of psychometric research that attempted to isolate and measure important human abilities. They are looking beyond and beneath these measures of ability to determine the essence of the individual differences the tests detect. The research has involved detailed analyses of the cognitive processes used in performing the tasks presented in tests with a visual character: *Cube Comparisons, Raven's Progressive Matrices,* and *Mechanical Reasoning.* Figure 5 shows examples of these test items. The research method involves collecting several kinds of data as people do the test items. A high technology computerized laboratory is used to record eye movements, showing precisely where the person is looking at each moment during the test. The total time required to answer each test item is recorded, providing the reaction times that are the most commonly used type of data in experimental psychology. Verbal reports of what the subjects are aware of doing as they perform the test items are collected also. A second step in the research is formulating computer simulation models, production system models, of cognitive processing that can

**Figure 5**

*Ability Test Items. These are illustrative items for the* Cube Comparison Test *(a), the* Raven's Progressive Matrices *(b), and the* Mechanical Reasoning Test *(c).*

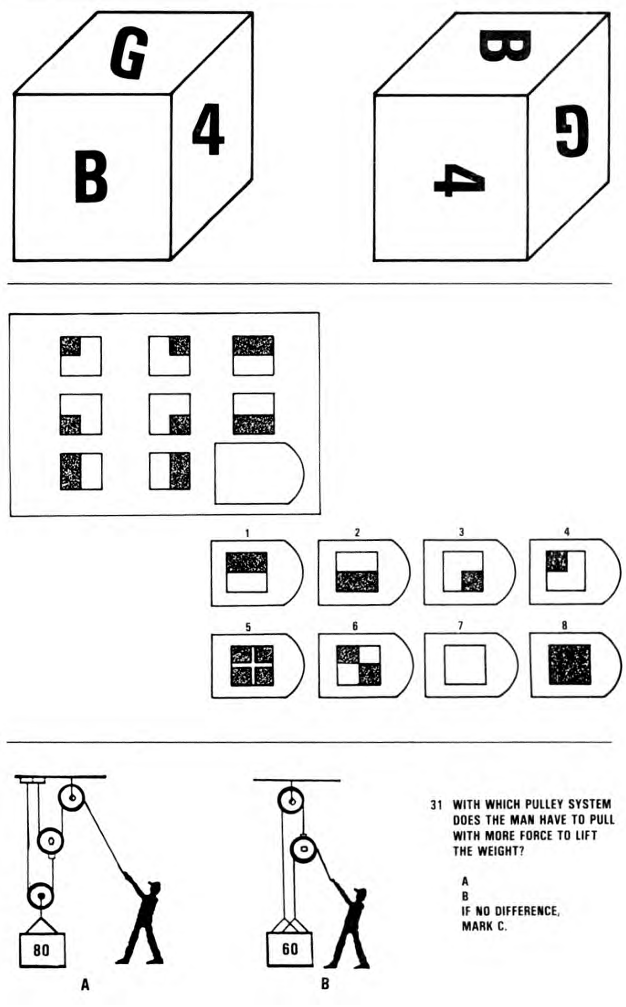


account for the data. These experiments use subjects selected for high and low ability performance on a previously administered battery of spatial tests. Thus, separate models can be developed to simulate the performance of high and low ability subjects, and the models can be compared to identify the source of the differences in performance.

A production system model of a psychological process is a computer program with two main elements – a working memory in which problem information is stored and a set of rules. The rules consist of a condition that is tested against the information in working memory and of an action that is taken if that condition is met. The action may change the contents of working memory, or it may be an overtly observable action.

In the *Cube Comparison* test, the subject's task is to determine whether or not the two cubes pictured could be the same (subject to the condition that each letter or number appears only once). It was possible to model the performance of high and low ability subjects with very similar production systems. Most people perform this task by mentally rotating the cubes, although this research did detect other successful strategies. High ability individuals can mentally rotate the cubes around any arbitrary axis to make the comparison, whereas low ability individuals only rotate around the canonical x, y and z axes so that they require several steps to complete many problems. In addition, the rates of mental rotation are faster for high ability individuals. Figure 6 shows the fits of these models to reaction time data for problems of varying difficulty. The production system model involves an arbitrary cycle time so the amount of time required to respond to a display of two identical drawings was used to calibrate the cycle time. The fit of the model to the data is very good, accounting for 94.2% of the variance.

Comparable research on the *Raven's Progressive Matrices,* which has a spatial character but is also the most generally accepted pure measure of general intelligence, had very interesting results. The primary difference between the simulation of the performance of average college students and the simulation of the performance of outstanding students was the addition of processes that systematically manage other processes. In order to perform well in this test, you must very systematically compare and keep track of changes in the figures in each row or column. Unlike the example we have shown, many of these items are very complex. In contrast, the research on the *Mechanical Reasoning* test seems to be leading to the conclusion that the important factor is knowledge about which variables in the situation are relevant, knowledge that may be acquired either from practical experience or from physics courses.

Stephen Kosslyn of Harvard University is another ONR supported researcher interested in the cognitive processing of visual information. While Carpenter and Just have focused on a few processes like mental rotation, Kosslyn has proposed a much more complete "cognitive architecture" for visual processing. It includes the proposal of a special short-term memory storage for visual information that is functionally like a map of the perceived visual field, long term memory for information about the literal appearance of objects as well as the more usually discussed long term memory for facts about objects, and very specific hypotheses about the processes that exist to manipulate information in these memory stores. In a contract that is now completed, Kosslyn sought to demonstrate that there are separate human subabilities for performing tasks that require mental imagery, subabilities which correspond to his hypothesized processes. Currently, Kosslyn is studying patients with brain injuries to obtain stronger evidence for the separateness of these processes and to locate the portions of the brain that carry out these processes. This work was described in *Research Reviews,* Four/1985. Martha Farah of Carnegie-Mellon University, who was recently designated as a recipient of one of ONR's special Young Investigator awards, is doing related work to identify the brain localization of processes of mental rotation.

**Figure 6**

*Time Required to Respond Correctly to Cube Comparison Items. This graph illustrates the time required to respond correctly to Cube Comparison test items that vary in difficulty because of the number of degrees of rotation, the number of matching letters and the type of axis rotation. The open circles represent predictions made by the theoretical model developed by Carpenter and Just.*

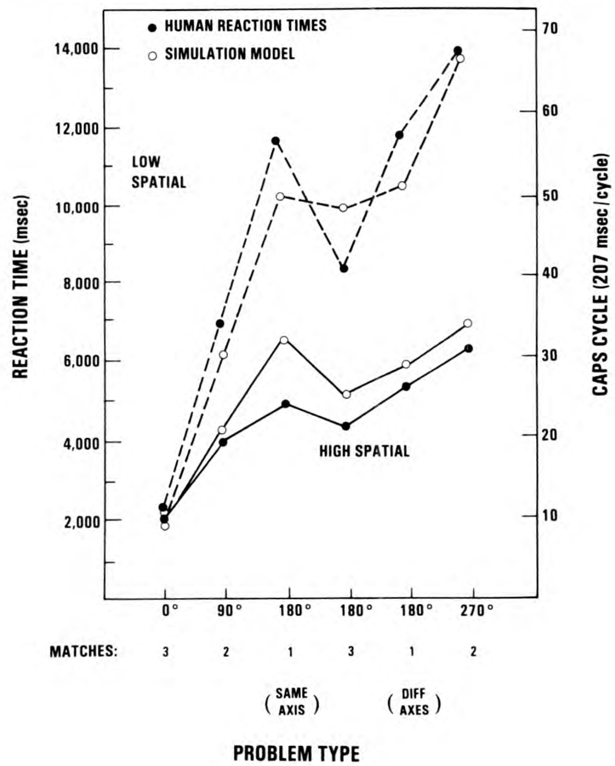


Like Carpenter and Just, James Pellegrino of the University of California, Santa Barbara, has been interested in identifying the sources of individual differences in performance on spatial ability test items. But he has also had a special interest in determining whether these processes and abilities are modifiable by training. Figure 7 shows the effects of the training Pellegrino gave on a spatial ability test; similar results were obtained on 5 different tests. There are two groups of subjects—initially very high ability and initially very low ability. The figure shows (1) initial scores (in percentiles), (2) scores after training, and (3) scores after a delay of several weeks. The training provided prac-

**Figure 7**

*Training Effects on the Paper Form Board Test. This graph shows that Pellegrino's training of component spatial skills had a large effect on the performance of subjects who initially had very low scores on the* Minnesota Paper Form Board *test of spatial abilities, which requires one to identify which of several whole figures could be assembled from a group of puzzle pieces. Percentile scores for high and low ability groups of subjects are shown: (1) before training, (2) immediately after training, and (3) after several week's delay.*

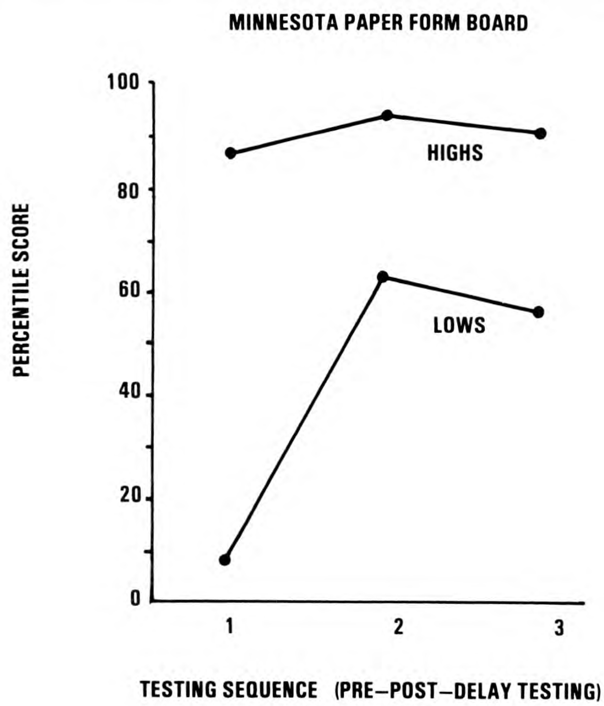


tice on component processes of the spatial task, like mental rotation of letters. It was not practice on the actual test items. Figure 8 shows how the speed of mental rotation increased with practice, more so for "low ability" subjects and shows that increased speed was retained after a delay of several weeks.

Reflecting upon these studies of visuo-spatial processes as a group, a number of important implications can be identified. Visualization ability decomposes into a number of relatively independent processes that can provide a sounder basis for measurement. Currently used psychometric tests mask strategy differences that can lead to prediction errors. More importantly, the research indicates that some of these abilities may be modifiable with appropriate training. Thus, a greater proportion of the population may be able to master the skills they predict. On the other hand, there are some individuals who do not profit from training and who may possibly lack the neural apparatus to perform these mental operations, just as some people are colorblind. Furthermore, such results are not confined to the visuospatial domain. Another ONR contractor, Robert Sternberg of Yale University, has conducted similar research analyzing verbal abilities and successfully training their components.

**Figure 8**

*Improving Mental Rotation Speed. Underlying the improved spatial ability scores shown in Figure 7, was improvement in component skills that were practiced, like the mental rotation of forms. Here, the amount of time required to verify the identity of two alphanumeric characters or unfamiliar letter-like forms is shown in milliseconds per degree of rotation at the time of initial testing (1), after several practice sessions (2), and after several weeks of delay without additional practice (3).*

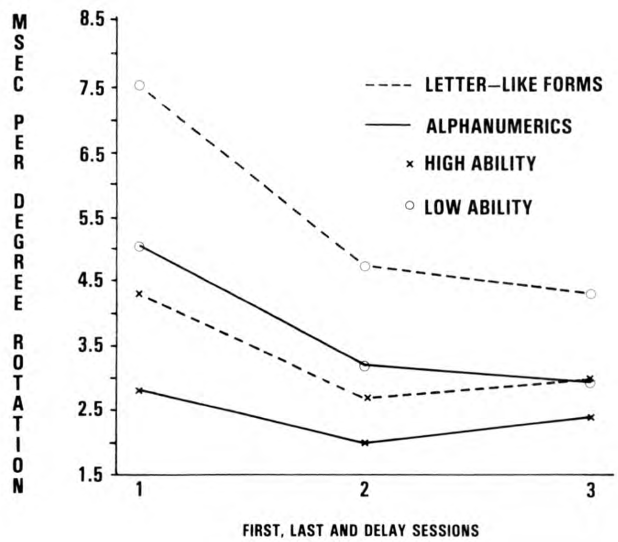


From a scientific point of view, many important questions remain in this research area. For example, Kosslyn and Carpenter and Just disagree about such details as the nature of the storage of visual information and the degree to which mental rotation approximates a continuous process. There is a need to account for individual differences in the rate of rotation or the efficiency of processes. There is a need to account for the changes in speed and efficiency that occur with practice. It is likely that we will need to turn to different modeling approaches, at a finer level of explanation, in order to account for these phenomena.

## Automaticity and Control

Understanding the operation of the human information processing system involves having to understand how that system is internally controlled. How is the information that will be processed selected? How are the processes that will be executed selected? Notice that the production system models discussed above incorporate particular notions of control. Selection of processing actions occurs as a result of matches between information in working memory and the conditions in production rules. There are additional issues in these models about priority ordering or (conceptual) parallelism of productions. The outcome of the research on individual differences indicates that the issue of the control and management of cognitive processes deserves more research attention in the future: it seems to be a key to understanding high levels of intellectual performance.

The research of Michael Posner (University of Oregon/ Washington University) is directed at understanding the operation of visual attention. This is a mechanism which provides favored processing to information that is present in the attended location. Posner is delineating the component processes of visual attention, processes that disengage, move and engage visual attention in a new location. Currently, he is using the performance of individuals with different kinds of brain damage and PET scan data from normal individuals in order to provide strong evidence that these processes are distinct and separable and to locate the regions of the brain that carry out those processes. This work is of considerable interest to neuroscientists working with animals, who are using Posner's ideas to guide the design their experiments on the neural control of visual attention.

The work of Walter Schneider (University of Illinois, now University of Pittsburgh) may be the most significant research effort focused directly on issues of control. Two of Schneider's papers have rapidly reached the status of Science Citation Classics. Schneider has attempted to understand how complex tasks such as flying and driving can become apparently "automatic", undemanding of attention and easy to combine with other tasks. He distinguishes "automatic" from "controlled" processes by these characteristics:

"Automatic" processes are—

- fast (up to 100 times faster)
- "parallel"
- undemanding of cognitive effort
- not limited by short-term memory
- resistant to disruption by fatigue, alcohol, etc.
- rigid, and difficult to consciously modify (as when we find it difficult to deviate from our normal path to work)
- require practice under consistent conditions in order to develop.

"Controlled" processes are—

- slow, serial and effortful
- limited by short-term memory
- responsive to novelty, flexible
- easily disrupted
- interactive with automatic processes

Of course, what is under discussion here is two different mechanisms or types of control.

Schneider's work is now moving in two disparate directions. One is toward practical application, the other toward deeper theoretical development. He has published general training guidelines for the development of automated processing, and is himself applying these ideas to the training of air intercept skills for the Naval Training Systems Center and to automated components of electronic trouble-shooting skill for the Air Force. These applications are scientifically important, too, because they will determine whether Schneider's results can be generalized successfully from the very simplified experimental situations he has used.

Schneider is also working on a computational theory of automatization, trying to provide an explanatory account based on changing strengths of connections among neuron-like units. Such a theory would have very wide significance. For example, we have just discussed the speed-up of memory encoding (shown in the work of Chase and his associates) and the similar speed-up of mental rotation. In general, automaticity of low level components of skill (like phonetic decoding in reading) is thought to be important in making higher and more complex forms of skill possible. In moving to neural network models, Schneider is linking up with another important theme in the ONR PT research program.

Several researchers, among whom David Rumelhart of the University of California, San Diego, is a central figure, are engaged in a theoretical enterprise which they call the *pararallel distributed porocessing* approach. In a newly published book, Rumelhart describes the effort to develop computational models within the constraints of "brain-style" processing:

- neurons are slow
- there are a large but not unlimited number ($10^{10}$ to $10^{11}$)
- large fan-into and fan-out from each unit
- learning involves modifying connections
- excitatory and inhibitory messages—not symbol passing
- continuously informative states, not stages of processing
- graceful degradation of performance with damage
- no central processing unit.

Rumelhart and others of this group, such as McClelland and Hinton of Carnegie-Mellon, have outlined an agenda of major classes of psychological phenomena (well-described in the results of decades of previous psychological research) that they are trying to account for in this theoretical framework. Phenomena of motor control, language grammar, and pattern recognition are among them. In part, this work is a challenge to the more dominant production rule approach to psychological modeling that we have described, an approach that is yielding increasing accomplishments. But really it is an effort to develop explanations at a more microscopic level of description that may link more readily to the operations of the brain. In Rumelhart's approach, rules (productions) would be viewed as expressing regularities that emerge from interactions among simple processing units. Thus, they would remain valid theories at a grosser level of description.

As we conclude this description of research on the human learner, it is appropriate to point out the rich interconnections that exist among the different research efforts. It is difficult to talk about processes without also talking about the encoding and storage of information. And it turns out that some rather complex processing may be involved in that encoding and storage. Furthermore, it seems that what had been regarded as basic abilities, basic characteristics of the system that learns, may be modifiable skills of the sort that we now turn to discussing. And large, complex bodies of knowledge may be used in performing even the most apparently simple task – the digit span test of memory. Thus, we turn to research on the nature of knowledge and skill.

## Knowledge and Skill: The Clearly Flexible Human Repertoire

This is the second major theme of our research program. The first was concerned with understanding the parts that the human system makes available for the building of skilled performance. Here, we are concerned with understanding the specifications of what we are trying to build in training. The objective of this research is to give us a better definition – a definition at the same refined level of description as the work on cognitive processes – of the knowledge and skill we seek to produce as a result of training. General goals such as achieving "learning with understanding" or of teaching general scientific reasoning or problem solving skill have been around for a long time. But this research aims to reveal exactly what those goals mean in terms of detailed cognitive structures and processes that the student or trainee must come to possess. Secondly, this research aims to characterize the cognitive change that occurs as the individual progresses from ignorance to expertise: to produce a cognitive learning theory. If we understand, in the same detail, how change can occur, that should be a big step towards our goal of making change occur in the way we want when we educate or train.

The research emphasizes problem solving in relatively well understood math and physical science domains, like basic electricity, that are also directly relevant to naval training problems. The general approach is to contrast performance at varying levels of expertise. A wide variety of psychological research techniques are used to generate evidence concerning the way people think about and solve problems. People may be asked to classify problems in order to reveal what characteristics they consider important, may be interviewed with a carefully designed set of questions, or may be asked to think aloud as they endeavor to solve problems. Frequently, a computer simulation of performance is used as a way to articulate and fully specify one's theory of the performance.

Once such specific and formal theories of performance have been achieved for two or more levels of expertise, it should be possible to advance very specific hypotheses for mechanisms that could transform one state of knowledge to the next – to propose the desired theory of learning. This area of research is now facing that challenge. Another present challenge is to begin to seek generalities, common features, in the specific cases of expertise development that are being studied: that is, to find general characteristics of human understanding of the world and of the mechanical devices in it. Metaphorically stated, we are trying to specify the computational language of the mind in which these learned skills are written.

### Some Interim Conclusions

Experts and novices differ in their major organizing conceptual categories; they see problems differently. For example, Michelene Chi and Robert Glaser of the University of Pittsburgh found that experts classify mechanics problems according to the involvement of principles of conservation of energy or conservation of momentum. In contrast novices (this time, students who have already done well in first year college physics) classify problems according to surface features like the presence of an inclined plane or a pulley. Similarly, James Greeno (Pittsburgh, now Berkeley) described the underlying semantic structures behind arithmetic and algebraic word problems that "expert" performers understand. Qualitative reasoning about situations has been found to be pervasive in experts. Such reasoning explores possible outcomes and guides the selection of formulas and other problem solving approaches that will be used to solve the problem. For this reason, qualitative reasoning, the so-called "mental models" that people use to think about phenomena and devices, have become a major focus of research. Such qualitative reasoning processes are not easy to express in present computer languages – that is part of the challenge and possible value of our attempt to understand the language of the mind. This work already seems to have important implications for education and training. It would seem desirable to attempt to establish these conceptual categories and qualitative thinking processes in students. In the next section of this article, we discuss a few examples of current research that is revealing more about those thinking processes.

## Computer Simulation of Expertise Variation

Mary Riley of the University of California, San Diego, (now at SDSU) has been studying different levels of expertise in solving electrical circuit problems of the sort that appear in the Navy's Basic Electricity and Electronics course. This work also aims to elucidate what defines "learning with understanding" in this domain. The research approach involved having individuals "think aloud" as they solved problems in order to provide information about a) the sequence and type of solution steps used, b) reasoning with equations vs. reasoning by analogy, and c) use of quantitative or qualitative relations. Computer simulations were developed as a theoretical exercise in order to test the sufficiency of the account for performance and, as it turns out, to generate new hypotheses for requisite problem solving knowledge. That is, the effort to produce a computer program that can solve the problems as one believes students are doing it reveals that some of the necessary components have not been specified. Riley's computer simulations are production system models, complex computer programs that include a database of problem situation information, of formulas for Ohm's Law and Kirchoff's Law, etc. In addition, they include productions (condition-action rules) that 1) parse (read) circuit schematics, (2) provide strategies to coordinate this information with representations internal to the model, (3) express knowledge of relations among entities in circuits, and (4) express knowledge of additive and multiplicative whole-parts relations, including procedures for designating quantities as wholes and parts and procedures for operating on wholes and parts to infer quantities or changes in value. This last is a type of general knowledge that can be applied to the specific subject matter of electric circuits.

At present, Riley has developed three variants of her computer model that are based on different patterns observed in the human problem solving protocols (records of thinking aloud). Two are fully adequate to solve all of the problems used to test mastery of the BEE course. One is a set of disconnected rules that does not include part-whole relations. That is, this program does not "know" that the currents I1 and I2 flowing through parallel resistors like those shown in Figure 9 are parts of the total current I and that the parts must sum to the whole. A second model does include part-whole relations, and it does qualitative reasoning directly with circuit quantities. It is more efficient, more tightly structured, contains redundant constraints, and is more likely to be robust in obtaining correct solutions even if some information in the model were "forgotten". Riley also has a third model that simulates defective performance. It fails to distinguish between voltage drops and resistances. It mimics observed student errors but still succeeds in getting 84% correct on BEE problems. No wonder some graduating students seem to have inadequate understanding later on!

**Figure 9**

*Part-Whole Relations Among Currents.*

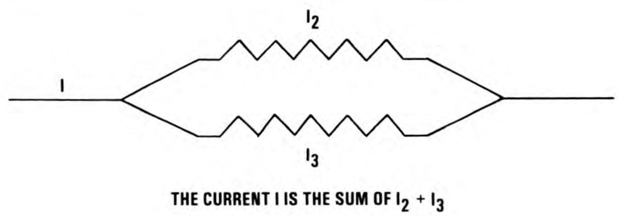


Riley's work is now moving into a more empirical phase, explicitly testing students who are studying BEE materials for the knowledge specified in the various computer models. Very high failure rates are being found on qualitative questions (e.g. If the value of this resistor is increased, what will happen to the current?) about relations among circuit quantities. Thus, despite the very theoretical orientation of Riley's computer modeling, there is immediate transition potential for her work. Coordinated work is being done at the Navy Personnel Research and Development Center in order to develop more sensitive testing of what trainees are learning from instruction about electricity.

Another interesting project is using the cognitive modeling approach to specify the knowledge required to transfer general scientific reasoning principles from one domain to another. *FERMI,* the work of Jill Larkin and Jaime Carbonell of Carnegie-Mellon University, is the simulation of an idealized expert in general scientific reasoning. It uses general principles like invariance and conservation of quantities, but also general strategies for exploiting those principles—strategies like decomposition, analogy to similar problems, and subgoaling. The notion is that this project will be used to define explicit objectives for an instructional experiment in teaching general scientific reasoning. This can be regarded as a first step toward a future goal: Once you have decoded the computational language of human thinking, it can be used to design new skills and more expert experts.

## Mental Models and Conceptual Change

Allan Collins, Dedre Gentner and Al Stevens of Bolt, Beranek and Newman, Inc., have been engaged in a research project with a rather different style. It began as a rather informal approach in order to reveal information about the nature of qualitative reasoning, aiming to minimize bias in what is observed by avoiding highly structured questions and problems. In order to permit the observation of general characteristics of thinking, people's understanding of three topic areas was explored: evaporation, home heating, and electricity. Questions were asked to provoke thought and explanations of phenomena. Does evaporation affect water temperature? How do you adjust your thermostat? Why?

A number of different features of thinking about such phenomena have been revealed in this way. There are three levels of thinking about these phenomena. People have macroscopic causal models linking variables like temperature and pressure. They also have molecular models that describe individual particles and their interaction. In addition, they have aggregate models that link the macroscopic variables to the behavior of aggregates of particles. A complex phenomenon or problem breaks down into subproblems. Thus, one finds a variety of different models that people have for components of the phenomenon. For example, mental models of the molecules in water may treat them like inert grains of sand, may include the notion of either uniform or random speeds, and may or may not include the idea of molecular attraction (see Figure 10). Various correct and incorrect models may exist for the different components of the total phenomenon. For example, some people think that when water becomes a vapor, it is disassembled into hydrogen and oxygen gas.

Last year, Collins conducted an experiment designed to test various extant theories about cognitive change, a point of considerable theoretical controversy now. Do novices adhere to consistent but wrong theories? Is there a progression through a sequence of increasingly adequate mental models? Does change involve differentiation and/or generalization of concepts? The answer to these questions is important because it has implications for the way one should go about instruction in order to bring about conceptual advance.

The prior work on people's correct and incorrect component models was used to generate test questions to be used in the experiment. Intensive individual tutoring by an expert was used to induce as much conceptual change as possible. The results indicated that novices accept statements consistent with many different models; they show no signs of adhering to a single consistent model. The conceptual change that occurs involves the acquisition of new basic mechanisms—like molecular attraction—that have broad effects. Now, we might ask whether there is a limited vocabulary of such basic mechanisms that people use to build mental models of phenomena.

**Figure 10**

*Mental Models of Component Processes in Evaporation. These are graphic illustrations of some of the differing mental models that Collins, Gentner, and Stevens of BBN found people using to think about the phenomena of evaporation. Various models of molecules in water are shown, along with various models of the process by which molecules return to a body of water from the air.*

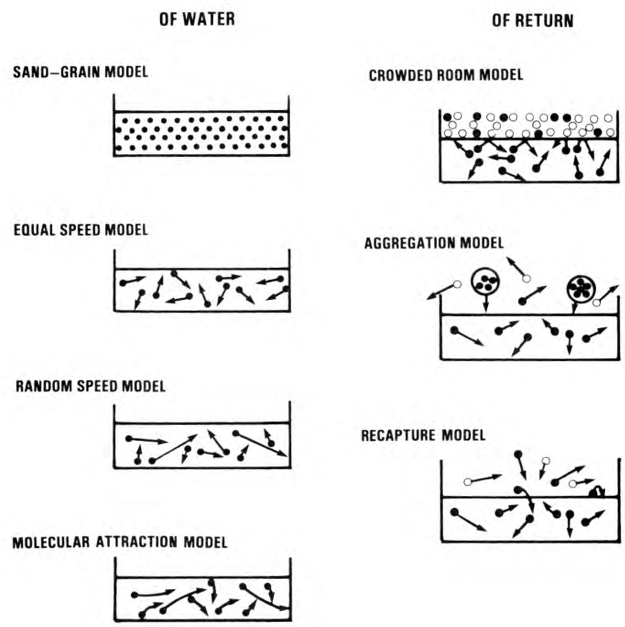


## New Directions

The most recent contracts awarded in this general area examine the sources of the mental models that people use to think about phenomena and mechanical devices. This work should provide a relatively direct basis for instructional theory. For some time, Dedre Gentner, now at the University of Illinois, has studied the use of analogies to known systems, such as the use of hydraulic analogies to think about electricity. She is now studying the factors that govern the formation of an analogy and its quality. Several researchers are simulating the process by which direct instruction is translated into internal production systems to govern action, an important element of John Anderson's theory of human learning. These include David Kieras of the University of Michigan and Carpenter and Just of Carnegie-Mellon. In a new contract, they will integrate their earlier work on reading comprehension with their work on comprehension of diagrams in mechanical reasoning (mentioned above) in order to simulate the understanding of device operation that results from reading an instructional text. In a very interesting new contract, Clayton Lewis of the University of Colorado is seeking to identify a set of simple inference rules that people use to connect the events that they experience when interacting with a device such as a computer system. Based on principles such as simple contiguity in time, the appearance of common elements, and the presumption that everything must be connected somehow, people draw conclusions about the causal connections between their actions and the behavior of the device. This work bears an interesting relationship to the work of Kahneman & Tversky (sponsored by Engineering Psychology), who tend to lament the way that people misuse evidence and leap to conclusions. For better or worse, we are learning something about how the human system makes use of the evidence it observes. Work like Lewis' can be used to design systems that will lead people to draw the right conclusions about their operation, thus ease training and reduce operational errors.

Metaphorically, the mental models that people use to think about phenomena and devices can be thought of as writing a simulation program – in the "language of the mind"– that is characterized by human information processing limitations and strengths. Several of the contracts just mentioned are doing exactly that. But we do not yet have good computer languages to express the qualitative reasoning that people seem to do in their mental models. Ken Forbus of the University of Illinois is working on a computer science formalization called Qualitative Process Theory. He is a contractor of the ONR computer science program, who has been inspired in his work by his association with the Collins group at BBN, and is now working with Gentner to formalize her theories about analogy formation and qualitative reasoning. It is interesting that the approximations of human qualitative reasoning that are expressed in Qualitative Process Theory are very, very demanding of computer resources. There are some ways in which the human mind/brain seems to be a very impressive computational device.

## Other Themes

There are other major themes in the PT research program that there has not been space to discuss on this occasion. One is research on technical and procedural instructional texts. This work is now showing an interesting change of focus. In the past, most research on text design has been concerned with the effects of surface characteristics – like grammatical structure or rarity of vocabulary – upon text comprehensibility. Now these researchers are saying that the more important issue is whether or not the text provides the information that enables the reader to build an effective mental model to guide correct procedural actions.

Another very major theme is research on problem solving, including work on problem representation, work on the way that insight restructures problem representations, and work on the transfer of problem solving skill to new situations. Of course, this work has strong linkages to what we have discussed. The assembly of complex mental models from component mechanisms – as Collins has described – is an example of problem representation.

## Some Major Accomplishments

Before moving on to the topic of research on instruction, let us recapitulate some of the major recent accomplishments in this area:

- ☐ Riley's work resulted in a theory-driven discovery of serious short-comings in the understanding that results from BEE instruction.
- ☐ Collins provided new insight into the way conceptual change comes about, rejecting previous theories on this issue.
- ☐ Gentner obtained experimental results showing that surface features of situations determine the formation of analogies.
- ☐ A new conception of what is critical in instructional texts emerged.

# Intelligent Computer-assisted Instruction

A third major theme of the Personnel and Training Research Program has been work aimed at developing a theory of instruction sufficiently precise to support the use of artificial intelligence technology to simulate the effectiveness of individual human tutoring. In 1985, an Accelerated Research Initiative, Cognitive and Technological Approaches to Navy Training, began to provide the resources to substantially expand this effort, with emphasis on relationships to Naval needs in the area of remedial basic skills instruction. Thus, there are many new contracts in this area.

As we mentioned at the outset, individualized human tutoring has a very large effect upon instructional outcomes, typically increasing achievement by two standard deviations, as compared to conventional classroom instruction. Other instructional variables that have been shown to be particularly effective are closely associated with the tutorial situation: performance feedback, student time on task, and the provision of cues and explanations. The benefits of one-on-one tutoring have been obtained in technical subjects such as cartography, probability, and computer programming as well as in traditional school subjects. For example, John Anderson of Carnegie-Mellon University found that only 11 hours of individual tutoring produced the same level of mastery of the LISP programming language as 43 hours of traditional classroom instruction with supplementary student homework on programming problems. However, one-on-one tutoring is a prohibitively expensive instructional approach. The technological opportunity here is the possibility that artificial intelligence techniques might be used to simulate human tutoring.

The problem of artificially intelligent computer-assisted instruction is also one that integrates many of the basic scientific concerns of the PT program. Furthermore, intelligent tutoring systems are uniquely appropriate as laboratories for the analysis of complex human learning. Because of their automated capabilities, the individual's learning and problem-solving can be followed and recorded step-by-step, providing objective protocols of performance from large numbers of students learning over long periods of time. New methods are being developed, and old methods adapted, for the analysis of these data, producing a major step forward in theory testing and refinement. At the same time, we are seeing new breakthoughs in the study of individual differences in learning, under conditions of controlled instruction that have never before been possible. The precise tuning and control of explanations and feedback in intelligent tutoring systems provides an novel opportunity for investigating questions of instructional theory and motivation in realistically complex learning situations. Many of the findings will be applicable to traditional classroom instruction as well as to automated instruction, but the questions could never be effectively investigated in traditional classrooms.

Intelligent tutoring systems, like human tutors, have a general knowledge of the domain that they teach. An intelligent tutoring system must be able to solve for itself each problem that it presents to a student. The tutoring system follows the student's progress through a problem step-by-step, inferring the student's plans and checking for possible errors and poor strategies. Students working with an intelligent tutoring system can try to solve complex, multiple-step problems in a natural fashion, receiving guidance along the way. This is very different from the rigid question-answer-feedback style of older forms of computer-assisted instruction.

An intelligent tutoring system is a complex system, involving several major components. In this way, it integrates many of the research concerns discussed in earlier sections of this article. Most individual contracts contribute knowledge that helps to realize one of the following necessary components of a tutoring system:

I. *An Ideal Student Model.* This is a type of *expert system* that mimics the performance one would like to see in an ideal student. Work on expert and novice problem solving performance, like that of Riley, Chi, Larkin, and Greeno contributes to our understanding here, defining instructional goals more precisely.

II. *Diagnostic Capability.* The tutor must be able to identify misconceptions and ways that the tutee is deviating from the ideal student model. ONR sponsored work that contributes to this capability includes that of John Seeley Brown (then Bolt, Beranek, and Newman) on diagnosing mental "bugs" in arithmetic skills, that of Patrick Langley and Stellan Ohlsson (then Carnegie-Mellon University) on automated inference of the (possibly faulty) rules governing arithmetic performance, and that of Derek Sleeman (Stanford) on diagnosing errors in algebra. A particularly complex example is Elliot Soloway's (Yale) PROUST system for identifying types of errors in student computer programs.

III. *Tutoring Knowledge.* This includes strategies used to teach the domain knowledge. The tutor determines when to interrupt the student, what instruction to give the student, and what problems to present next. Researchers supported by ONR have studied the strategies of effective human tutors. Collins and Stevens of BBN characterized the rules that govern the generation of questions by experienced Socratic tutors. Clancy of Stanford has shown how the knowledge that is required to provide adequate explanations of medical diagnostic reasoning differs from and goes beyond what is sufficient as an expert system to do the reasoning. Both Anderson of Carnegie-Mellon and Soloway of Yale have conducted studies of human tutors of computer programming in order to define strategies specific to that domain. Micheline Chi and Robert Glaser of the University of Pittsburgh are studying how and what students learn from worked example problems in a physics text, research that should provide significant guidance in the selection and use of example problems, an important instructional strategy. Similarly, Collins of BBN is following up the work on naive mental models that we described earlier with a computer-controlled instructional experiment that will compare the effectiveness of various instructional strategies for overcoming misconceptions. A new 6.2 contract with Thorndyke and Macmillan of FMC Corp. will work out the AI programming problems associated with the selection of particular instructional acts.

IV. *Simulation Capabilities.* Many training situations of interest to the Navy require that the trainee be able to interact with a dynamic simulation of an active system, such as a ship power plant or electronic circuitry. A research effort conducted by the Navy Personnel Research and Development Center and BBN, STEAMER, was a pioneering effort of this kind. Generalization of these capabilities (which involve the use of advanced AI programming technology) is now be carried out in an ONR sponsored contract, Direct Manipulation Programming, that involves collaboration between NPRDC and researchers at UCSD. This work allows creation of a working simulation by simply constructing a graphic diagram in which system components are appropriately connected. In a 6.2 contract jointly managed by NPRDC and ONR, similar capabilities are being developed for an inexpensive AI work station as part of an effort to produce a generalized intelligent tutor for maintenance tasks (Douglas Towne, USC). Similarly, Govindaraj at Georgia Tech has developed a capability to approximate complex computations in order to produce qualitatively correct simulations of ship power plant operation, for example, with high speed on a relatively inexpensive computer.

V. *The Student-Tutor Interface.* The effectiveness of this interface determines how easily the student can access the knowledge contained in the tutor's ideal student model. Research is being done on graphic means to communicate the goal and plan structure of problem solutions, as well as other forms of knowledge that cognitive research has shown to be both important and neglected in most current instruction. Research is also being done on the problem of adapting instruction to limitations on the tutee's attentional and information processing capabilities. Because many of the problems concerning natural language comprehension and production by computer are still unsolved, Barbara Fox, a linguist at the University of Colorado, is studying the ability of students to adapt to communicating with a computer that has limited language capabilities.

Although many problems remain unsolved, it is becoming possible to create relatively complete intelligent computer tutors with at least limited capabilities in all of these areas. John Anderson of Carnegie-Mellon University has produced an intelligent tutor of LISP (the primary AI programming language) that has actually been used in teaching a programming course at Carnegie-Mellon. This tutor also functions as a laboratory in which Anderson can collect data to test his cognitive theory of human skill learning. The National Science Foundation Information Science Program has supported Anderson in developing an intelligent tutor of geometry proof-making (based on Greeno's earlier ONR funded research) that is now being tested in the Pittsburgh public schools, in a project supported by the Carnegie Corporation. (Responsive to concern within the Navy about what will happen when this advanced technology is introduced into Navy training, ONR is sponsoring a study by Janet Schofield of the University of Pittsburgh of the impact of the intelligent tutoring system upon the roles of students and teachers, upon competitive motivation for student performance, etc.) NSF Science Education is now supporting both Anderson's development of an intelligent tutor for high school algebra and Soloway's development of a tutor for Pascal programming, as well as several other projects directly related to earlier ONR research. In this way, ONR's long-term investment in the research foundations of intelligent computer tutoring shows promise of improving the nation's capability to prepare the technically competent recruits the Navy needs. The products of this NSF development work may also be directly useful to the Navy and Marine Corps in their own remedial education programs. More specifically directed at Navy problems, is the Special Technology 6.2 contract (jointly managed by NPRDC and ONR) under which Towne at USC is producing a generalized tutor for maintenance training. This system will make it possible for an experienced maintenance technician to enter information about specific system components and failures and thereby produce both a system-specific simulation and an intelligent tutoring system. This system is built upon a previously developed generalized expert system for trouble-shooting that was sponsored by ONR's Engineering Psychology Program.

Towne, Anderson, and many of the other researchers mentioned here have been provided with common AI computer equipment through a major Defense University Research Instrumentation Program grant to the University of Pittsburgh. This facilitates the sharing of costly software developments in this research. A centerpiece to this effort is a contract with Alan Lesgold at the University of Pittsburgh, who is exploring whether one can find a generally applicable way to organize both subject-matter and tutoring knowledge so that the production of intelligent tutoring systems could become feasible for those who are not trained cognitive scientists. This basic inquiry is being supplemented and enhanced by more applied contracts from the Air Force for the production of specific tutors in such areas as basic electricity and airplane engine maintenance, projects with potential application for the Navy, too.

Although it is gratifying to see the beginnings of a practical payoff for ONR's investment, we think it is no accident that the individual, John Anderson of Carnegie-Mellon, who is leading this activity is also one of the leading theorists and basic researchers in cognitive psychology, author of some of the most influential books and articles in modern psychology. Anderson's work in intelligent tutoring is founded on his ACT theory of human learning and cognition and is largely motivated by his desire to test this theory in situations where people are learning real skills of real significance to them over prolonged periods of time. The ACT theory is the most complete and adequate cognitive theory yet proposed. It was created under ONR support through 1980. Additional ONR supported work on problem solving and learning in mathematics and computer programming, along with related work supported by NSF and ARI, established the ACT theory as the focal point of research on complex human learning. The theory and related experimental work are described in Anderson's 1983 book, *The Architecture of Cognition*.

Figure 11 diagrams the main elements of the theory. It postulates two long-term memory stores, a production memory that holds pattern-action associations such as those we have discussed earlier in this article. These associations are not literally verbal rules, but rather tendencies for certain states of working memory to elicit other states or to produce behavioral responses. Declarative memory holds factual or propositional information. This type of information is directly available from a textbook or from a classroom lecture. Unfortunately, it is not directly useful in solving problems. The ACT theory describes mechanisms for converting declarative textbook knowledge into the pattern-action associations which are directly used in problem solving. This conversion process is slow, effortful, and error-prone, and made more difficult by the well-documented limits of working memory capacity (cf. the work of Schneider discussed earlier). The theory also describes a mechanism by which pattern-action associations become stronger through repeated practice, leading

**Figure 11**

*John Anderson's ACT Model.*

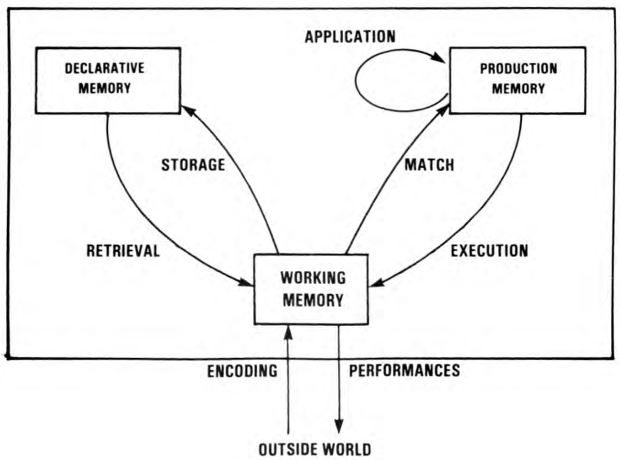


eventually to the high levels of automaticity characteristic of expert performance.

The ACT theory is proving to have impressive predictive power for learning situations. For example, David Kieras of the University of Michigan used it to make predictions about human learning of a sequence of operating procedures for an artificial device. The theory predicts that the difficulty of learning a procedure will depend upon the number of novel productions involved, upon the number of productions that are exactly identical with those learned as a part of previous procedures, and upon the number of productions that are easily generalized because they differ only slightly in the conditions to which they are applicable. Figure 12 illustrates the impressive fits that these theoretical models can give to the zig-zag curves of learning times required for procedures in sequence. Obviously, this theory provides a basis for optimum training sequences, a notion that Kieras is now pursuing with support from IBM for application to training in computer and word processing systems.

Similarly, Anderson himself recently carried out experiments concerning transfer of learning among three computer file editors. Two of the editors were line editors that are quite similar in their deep goal structure but very different in the superficial appearance of their commands. The third editor was a screen editor that has a very different goal structure. A production system analysis within the ACT framework predicted strong transfer between the two line editors, despite the difference in their commands. That is indeed what happened. Switching from one line editor to another in the course of learning had almost no impact on the learning curve of improving performance speed. Specific difficulties that did occur as a result of the switch were also predicted.

**Figure 12**

*Predicted and Observed Training Times. The average amount of time which a subject required to learn each of a set of procedures (identified by numbers on the x-axis) when they were presented in a particular order is shown. Open circles are the times predicted by Kieras' production system analysis, taking into account transfer of training among the successive procedures.*

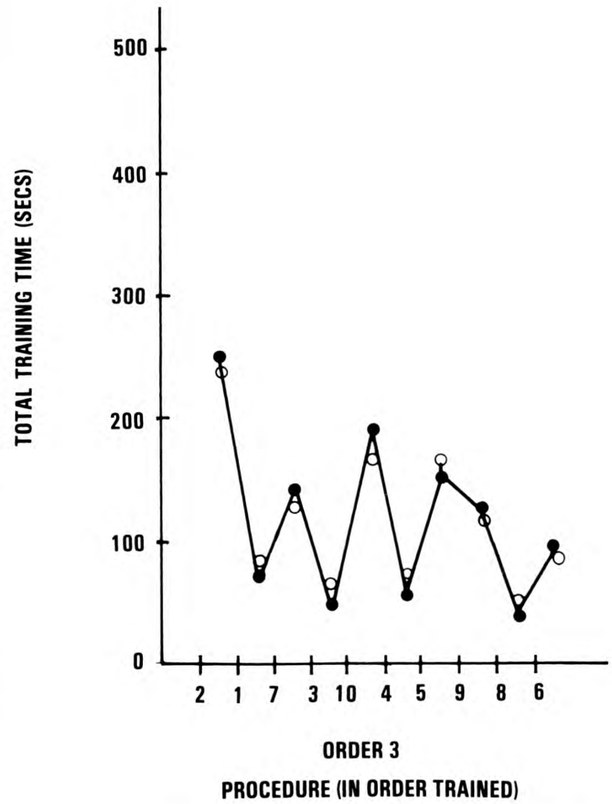


Thus, evidence is accumulating that production system analysis can predict the fine structure of transfer of training. This is an exciting accomplishment in one of the central areas of psychological inquiry. Since the Navy spends $2 to $2.5 billion per year on simulator-based training, transfer of that training to real situations is also an issue of major practical importance. We believe that the controlled laboratories being made available by artifically intelligent tutoring systems will permit unprecedented advances in the theoretical understanding of human learning and teaching in the next few years.

# Personnel Measurement

All of the research described above is associated with one of the two core programs in Personnel and Training, Learning and Training, which has been supplemented by funds from three Accelerated Research Initiatives. The second core program is Personnel Measurement, which addresses the basic theoretical foundations for the testing and selection activities of the Navy and Marine Corps. Indeed, despite its relatively modest size, this core program provides most of the support for basic research on test theory in the U.S. In recent years, it has made a particularly important contribution by providing the theoretical basis for new, efficient computerized adaptive testing techniques that are now moving toward operational use with the ASVAB (Armed Services Vocational Aptitude Battery) and elsewhere.

The theoretical basis for testing and selection involves hypothetical constructs in psychology such as ability, achievement, and capacity. These theoretical constructs are not directly observed, but must be inferred on the basis of the individual's responses to concrete observable stimuli, the test items. This research program focuses on model-based psychological measurement. It is concerned with the development of a theoretical framework and associated modeling techniques that will allow us to draw inferences about those theoretical constructs from the individual's responses to test items.

An important tool in model-based measurement is the item response function. Figure 13 shows the item response functions for the various alternative answers to a test item. (In order to keep the item confidential, we have substituted other words for the actual ones in the test item.) The graph plots the probability of selecting the particular response alternative against the individual's ability level, measured in standard deviation units with a mean of zero. Notice that the response function for the correct answer is a monotonically increasing function, but there is another alternative which is a very attractive wrong answer even among moderately capable examinees.

The theory involves two sorts of variables: item performance, which is observed, and abilities, which are not observed, but are inferred. Item response functions like those plotted in the graph play a central role in the theory because they tie the unobservable abilities to the observable item performance. Notice that these item response functions are themselves not observable. First of all, these functions describe relationships in the population, and we see only a sample. But in addition, we do not even see the abilities of the examinees in our sample. All we see is a collection of individual examinees' patterns of responses to questions like the one illustrated. So, to infer the shapes of these item response functions from data, we need to infer something about the abilities in our population. Finally, although many psychologists have been willing to stipulate constraints on the response function shape for the *correct* response, there is little agreement on what the functions should look like for *incorrect* alternatives. Consequently, there is usually insufficient basis for estimating those functions, and the information that might be provided by the selection of particular incorrect answers typically is not used.

A number of different technical issues are of central concern in the program. The most fundamental and far-reaching issue concerns whether two measurement paradigms are measuring the same thing. The research scientist is frequently faced with a situation in which multiple tasks can be constructed that theoretically should measure the same psychological process. In the practical testing world, one must have multiple forms of the same test in use simultaneously. The need for new theory is particularly acute in modern computerized adaptive testing, where each examinee is given a different set of items selected to zero-in on his ability level, just as an eye-doctor zeroes-in on the right prescription without wasting time on completely inappropriate lenses.

The problem of determining the shape of item response functions is equally important and equally difficult. The item response functions tell us how ability is governing performance. In the multi-dimensional case, they describe how two or more abilities interact to determine performance. For the testing practitioner, item response functions guide the design of efficient and effective tests. Items with flat functions are not useful in measuring ability because persons of high and low ability are equally likely to get them right.

**Figure 13**

*Item Response Curves. This graph shows the probability of selecting particular response alternatives as a function of the individual's underlying ability, for a particular verbal test item.*

*For example, the item might have been (though it was not, for reasons of test security) the following.*

*To draw conclusions from information is to: (a) imply, (b) think, (c) infer, (d) equivocate.*

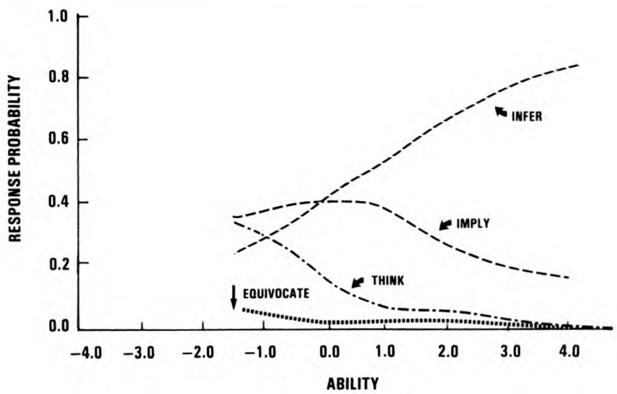

In order to have models of test performance that are identifiable, it is necessary to stipulate various constraints on the solution—conditional independence of items, continuity of the item response functions, and monotonicity are examples of such constraints. It is one of the major theoretical goals of this program to determine which constraints are necessary and how they interact with each other.

Another major goal is to develop techniques for on-line calibration of new items from data obtained when they are introduced unobtrusively into an examinee's operational test. This provides data from a realistic testing situation—as opposed to an experimental situation in which the test does not "count"—and reduces the cost of test item development. It also provides a way to monitor possible changes in item properties after initial calibration. For example, if a word becomes popular, knowledge of it may no longer be the indicator of high ability that it once was.

Appropriateness measurement is another focus for research. This refers to techniques for determining whether the pattern of responses given by the examinee is consistent with the model. Practically, this can be used to identify individuals who are trying to fail or who may be hampered by language difficulties, individuals for whom the test has not been an appropriate description of capabilities. For the researcher, it can help to identify groups of subjects with different response patterns, who may merit further study. For example, this approach could be used to identify groups of subjects who are responding consistently or inconsistently with one of the models of expert problem solving behavior described in earlier sections of this article. Dr. Kikumi Tatsuoka of the University of Illinois is pursing such possibilities in her research.

The technical approach in this program is quite comprehensive and includes theoretical development of the measurement foundations, mathematical and numerical development of modeling techniques, and Monte Carlo and empirical studies to assess the adequacy of those techniques.

Typical questions in measurement foundations include:

- What do we mean, in a rigorous mathematical sense, by the dimensionality of a test?
- What does it mean to say that two tests have been equated?
- When are two ability distributions equivalent?

Typical questions of a psychological nature include:

- What is the nature of dependencies among items in various domains and how might they be modeled?
- What is the nature of item omitting behavior, and how might it be modeled?
- What is the nature of "slips", temporary lapses of performance, and how might they be modeled?

Typical mathematical/numerical questions include:

- How can we improve the fidelity, robustness, and numerical properties of our modeling techniques, especially with small samples of people?
- What sorts of statistical tests will inform goodness-of-fit and model validity?

We have space to describe only a few of the recent accomplishments of the program. Mike Levine at the University of Illinois has obtained a rigorous mathematical representation of the unobservable ability distribution on a test. This representation decomposes the abstract ability distribution into that part which is *identifiable,* in the sense that it has implications for the probability distributions of some of the observables on the test, and that part which is *not identifiable,* that is, has no implications for any of the observables on the test. Levine has also reduced the problem of estimating the ability distribution to the mathematically tractable problem of maximizing a convex function.

Fumiko Samejima at the University of Tennessee has been developing techniques to discover the shapes of item response functions without making strong *a priori* assumptions about constraints on shapes. Recently, she has developed efficient and robust techniques. For the first time, we have theoretically sound, computationally feasible, and logistically feasible approaches for studying the shapes of the item response functions for the incorrect alternatives on a multiple-choice item and for studying omitting behavior. The item response functions given at the outset of this section were obtained with Samejima's techniques. They allow us to get more information about examinees' abilities from a single test item.

Bill Stout of the University of Illinois has developed a powerful test of the hypothesis that the underlying space of ability determining test performance is of dimension 1, as against the alternative that its dimensionality is greater than 1. This result is important to the computerized adaptive testing programs of the Navy and Marine Corps because they rely upon the well-developed unidimensional item response theory. Therefore, it is important to know whether one's set of items is consistent with the unidimensional assumption of the theory.

## Future Directions

We expect to continue investigations of the challenging theoretical problems described above. In addition, one of the major goals and management challenges of the Personnel and Training Research Program as a whole is to strengthen the connections between the research program in model-based measurement and the cognitive research elsewhere in the program. The theoretical constructs postulated in modern cognitive psychology—semantic networks of concepts, sets of productions, etc.—are quite different in character from the concepts of ability dimensions that underlie the theory of model-based measurement. In most applications, ability dimensions are a very global level of description which has considerable validity and predictive power in personnel selection but which is not useful in providing detailed guidance to instructional activities. Typically, the process by which test items themselves are generated and developed cannot be described as theory-driven or even theory-guided, although items are selected and evaluated by the test theory. As Riley's research indicates, conventionally developed test items may not effectively measure what you had hoped. On the other hand, the processes of inference by which cognitive psychologists test the validity of their theories—like most inference in scientific research—are often considerably less rigorous than the techniques of model-based measurement. Using the tools of item response theory, microabilities might be defined that correspond to particular aspects of the skill or knowledge being learned. Alternatively, similar but novel approaches might be used to obtain statistical estimates classifying trainee's mastery in the terms provided by a cognitive theory of the skill. We believe that both research communities—now quite separate and distinct—could profit from the achievements of the other. In addition to encouraging specific contracts that bridge these research approaches, the program is sponsoring activities like an ONR-Educational Testing Service conference on cognitive diagnosis, discussing the problem of diagnosing cognitive states of knowledge and skill acquisition with the rigor characteristic of model-based measurement.

# Conclusions

We believe that we are participating in a very exciting period in the history of psychology. A mature science of cognition is emerging. A substantial body of researchers is beginning to work within the common theoretical framework of cognitive architectures. The ONR supported researchers Anderson and Kosslyn are leading exponents of this general framework. It generates decideable theoretical questions. The disagreements between Kosslyn and Carpenter and Just about the form in which information with a visual character is internally represented and processed is an excellent example of such a theoretical question. This theoretical framework has the power to reveal new information about the nature of cognitive understanding, as we saw in the work of Riley. It can make impressive predictions, as shown in Kieras' and Anderson's results on the transfer of learning. Perhaps most importantly, it is integrating a wide range of phenomena, linking up what had been separate islands of psychological research investigation. These links reflect the actual interconnections among these phenomena in the functioning human system.

**Figure 14**

*Future Directions for Personnel and Training Research.*

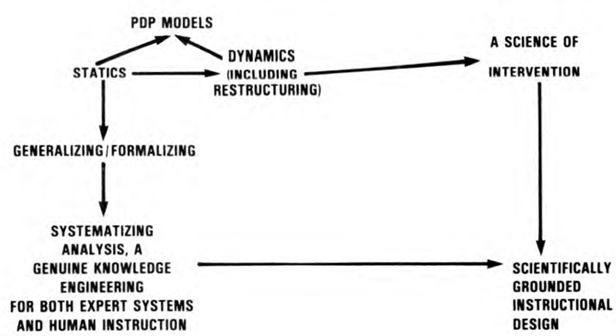


Of course, the exploration of the full scope and power of this theoretical framework is just beginning. The results we have mentioned are no more than a tantalizing sample of what we expect to see achieved. Figure 14 is a schematic representation of the future directions that we see for the program. At present, the primary focus is on description of static states of cognitive functioning, of expertise levels, of mental models. Now, researchers are beginning to address the problem of dynamics, to characterize cognitive change and to propose genuinely cognitive theories of learning. Situations in which knowledge seems to be fundamentally restructured will pose a particular challenge. More explicit cognitive theories of learning should lead naturally into a scientifically grounded instructional theory that will enable us to intervene in learning in a principled way. Another direction of future development of this research is toward greater formality of expression and toward greater generality. Indeed, only by developing abstract, formal expressions of the cognitive structures being described in much of this work will we be able to determine whether the same theoretical descriptions are valid across a wide range of specific domains of knowledge. If those formal expressions are achieved and prove reasonably general, they will provide a basis for systematizing the analysis of knowledge and skill. Then, there could be a *knowledge engineering* that—unlike the situation today—deserves the label *engineering*. This would be of great practical importance, because knowledge engineering, the analysis of expertise, is a serious and costly bottleneck in both the production of expert systems and the design of instruction. By the year 2000, these two directions of research progress could reasonably come together to provide a scientifically grounded technology of instructional design that would enable us to effectively and efficiently meet the severe training demands of a highly technological Navy.

At the same time, we expect to continue pursuing the alternative theoretical framework represented by parallel distributed processing. It promises a finer level of explanation with closer links to the neurological substrate of cognition. And it may offer insights into the way that artificial systems can be built to mimic some of the accomplishments of human thought that now seem intractable.

The ONR Personnel and Training Research Program has played a critical leadership role in the rise of cognitive science in the U.S. And cognitive science has been identified as one of the critical areas of modern science in which the U.S. now occupies a leading position. In order to have the resources to maintain that important role, the program staff have been actively participating in ONR's Research Option process.

## Authors

Dr. Susan F. Chipman was Program Manager for Personnel and Training Research and is now, under a new organization, Program Manager for Cognitive Science. Dr. Charles Davis manages the core program in Personnel Measurement. Dr. Michael Shafto participates in the management of the cognitive research described, especially that involving artificial intelligence. This article is largely based on the program review that was presented to ONR management in December, 1985. The authors would like to thank another PT staff member, Dr. Harold Hawkins, for his assistance in preparing the review. Dr. Willard Vaughan, acting Division Head of Psychological Sciences, and Dr. Steven Zornetzer, Associate Director for Life Sciences, provided comments helpful in shaping the contents of that review. The authors would also like to acknowledge the vital contributions of two former members of the ONR staff, Dr. Marshall Farr and Dr. Henry Halff, who managed the Personnel and Training Research Programs and are largely responsible for the research accomplishments which we are able to report here. ONR's psychological research programs have been reorganized and renamed; most of the work described here will be continuing under the new Cognitive Science Program within the Cognitive and Neural Sciences Division.

---

Title Illustration on page 2: A screen display from an artificially intelligent tutor of optics that is under development at the Learning Research and Development Center at the University of Pittsburgh. This particular tutor is primarily supported by the Office of Educational Research and Improvement of the Department of Education. Professor Robert Glaser is the Principal Investigator, the designer is Peter Reimann, and Professor Hans Spada of the University of Freiburg in Breisgau also contributed. The tutor relies upon software tools for intelligent tutoring produced under an ONR contract with Alan Lesgold and Jeffrey Bonar of LRDC and contributes to the general understanding of issues in intelligent tutoring sought under that contract.

# Profiles in Science

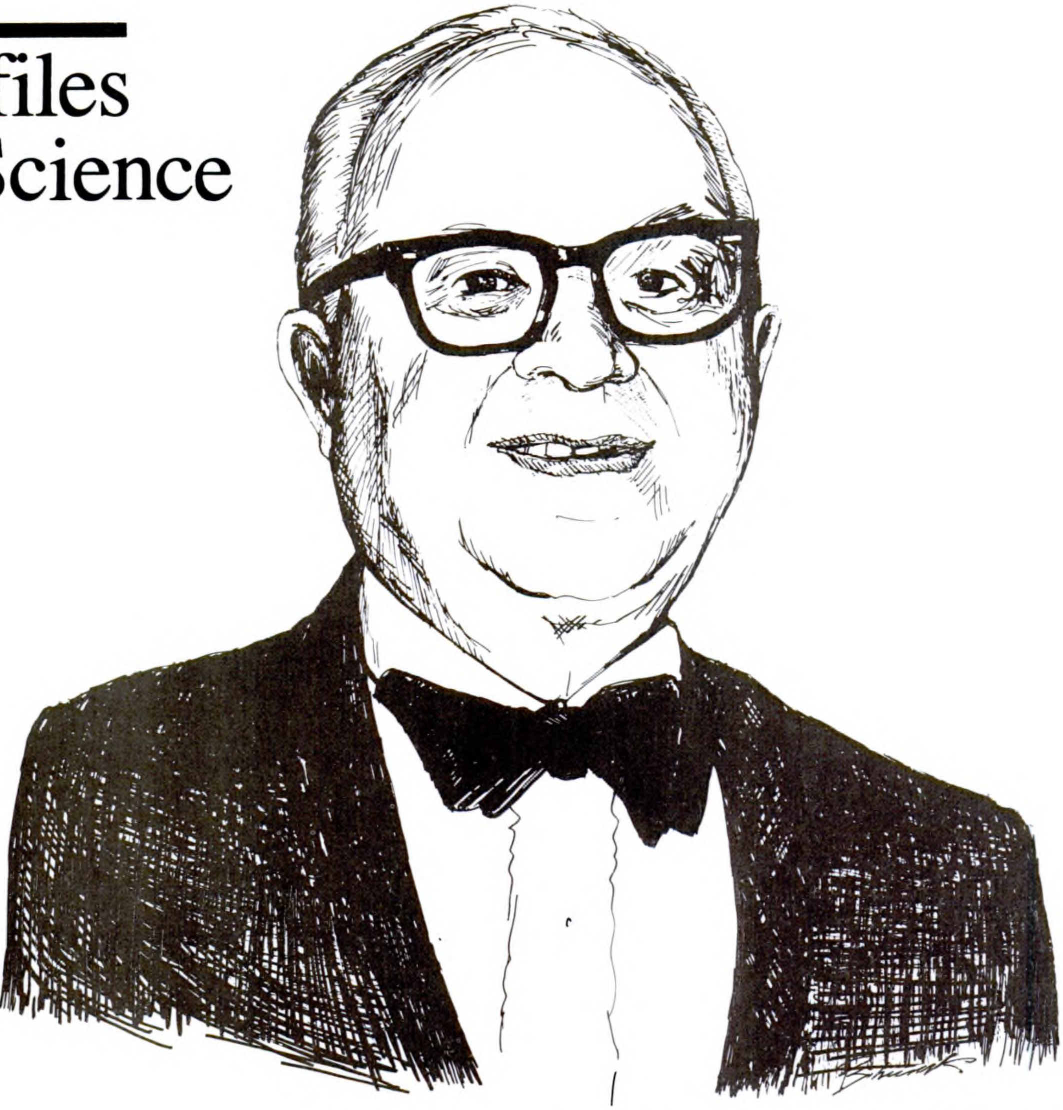

## Captain Charles W. Shilling

MC, USN(Ret.) recently retired after 13 years as Executive Secretary and Treasurer of the Undersea Medical Society where he was a dominant force in advancing the hyperbaric oxygen theory as an offshoot of diving medicine. He has directed numerous diving research projects, three of which are noted because their published reports became "key documents" in the literature: "A Suggested Change in Calculating Decompression Tables," "Quantitative Study of Mental and Neuromuscular Air Pressure," and "Auditory Acuity in Submarine Personnel."

His Navy career began with internship at the U.S. Naval Hospital, Chelsea, Massachusetts in 1928; and for the next 20 years, he assisted in the research and development of safety, salvage, and escape equipment for submarines. He was the medical officer in charge of the rescue in 1939 of the men from the submarine USS SQUALUS. Later, he established and was the first Director of the Naval Medical Research Laboratory at New London, Connecticut, where an effective research program was established in support of the needs of the submarine service.

In 1947, CAPT Shilling was assigned to the staff of the Office of Naval Research where for six years he was responsible for the medical, biological and psychological research programs. For the final two years of this period, he had additional duty as Director of the Research Division, Bureau of Medicine and Surgery. Through out his career he has often been an investigator for ONR.

CAPT Shilling was assigned to the U.S. Naval Academy at Annapolis, Maryland in 1953 where he had concurrent assignments as the Senior Medical Officer of the Academy, Command Medical Officer of the Severn River Naval Command, and Head of the Department of Hygiene. He retired from the Navy in 1955.

For the next years, CAPT Shilling served as Deputy Director, Division of Biology and Medicine, United States Atomic Energy Commission. Here his duties concerned the research, education and training programs as they pertained to the life sciences. Then in 1960, he joined the staff of the American Institute of Biological Sciences as Director of the Biological Sciences Communication Project where he stayed until joining the staff of the Undersea Medical Society. He is a Fellow of the American Academy of Arts and Sciences and American Institute of Biological Sciences as well as President of the Academy of Medicine of Washington, D.C.

# UNDERSTANDING MEMORY PROCESSES



Paul E. Gold
University of Virginia

Joel L. Davis
Office of Naval Research

## Introduction

Most individuals can fairly easily remember where they parked before work that day but cannot recall the spot where they parked last Tuesday. Similarly, they can remember what they had for dinner last night but cannot recall the dinners of last week. Not all memories are this fragile. Individuals are very likely to remember their parking spot if, upon returning to the car, they observed another car hitting and damaging their own vehicle. Other events are also printed indelibly in our memories. For example, depending on the reader's age, he may recall vividly where he was when he heard of the bombing of Pearl Harbor or the assassination of John Kennedy. So, quite evidently we remember important events, and their circumstances, better than other events, even if the less important event occurred more recently. The science of the biology of memory attempts to explain this phenomenon and others by examining events in the nervous system. Recent work in this field funded by the Office of Naval Research has identified a physiological system which appears to control the memory storage process. This article will review the current evidence and will suggest some of the likely directions and consequences of research dealing with physiological regulation of memory storage.

## Neurobiology of Memory

As recently as twenty-five years ago, there was no evidence showing that brains were biologically altered by experience. The best indications that such changes occurred were based on the fact that humans and animals learned and remembered and it was hypothesized their behavioral changes were represented by some type of brain modification. In 1961, the first of several foundation papers appeared which indicated that rates raised in "enriched" environments, complete with toys and other rats, had heavier brains than those of animals raised in "impoverished" environments, i.e., in standard individual cages. Clearly, this important discovery meant that brains could be changed by experience. On the other hand, changes in weight had such vague meaning for brain function that the importance of such changes for learning and memory was obscure. More recently, laboratories at the University of Illinois and at Berkeley have found evidence for much more precise changes in the fine anatomy of the brain, changes which occur not only when animals were reared in different environments but also when they received specific types of training. The changes appear in the extent of dendritic branching of certain cells. The dendrites are a major source of incoming information to the neurons and the findings, therefore, indicate that information transmission between cells may be enhanced by these structural changes. There is also a substantial amount of evidence showing both physiological changes in the transmission across neurons and relatively long-lasting biochemical changes in neurons after training.

Whether one or several of these changes represents the biological basis for memory is unknown. However, the findings are encouraging in the sense that changes similar to these are quite likely to mediate memory storage. Similarly, if we take these brain modifications as the basis for memory, we do not yet know why they occur or what regulates the probability that they will occur.

## Hormones and Memory

Many forms of acute brain damage or trauma (e.g. seizures or concussion) produce a relatively specific form of memory impairment. The impairment is a loss of memory for recent events, usually for experiences which occurred from several minutes to a few days before the trauma. This phenomenon is termed *retrograde amnesia* because of the time-dependent nature of the memory loss.

Comparable findings are readily obtained in animals. In a typical study, rats or mice receive training followed at various times by the administration of a drug or electrical stimulation of the brain. The animals then receive memory tests a day or more after the training-treatment condition. The results of these studies show that treatments which produce amnesia are more effective when administered shortly after training. As the training-treatment interval increases, the extent of the amnesia decreases. The time course after training during which a given treatment impairs memory is termed the *retrograde amnesia gradient.* The gradient length varies substantially with several known factors, including brain stimulation intensity or site, drug dose, type of training, and circadian rhythms. Thus, the major rule underlying the phenomenon of retrograde amnesia is that more recent memories are more susceptible to disruption.

Some time ago, Dr. James L. McGaugh and his colleagues (See Naval Research & Reviews, Vol. 37, No. 4) discovered that some treatments could accomplish the reverse of this phenomenon; many stimulant drugs enhance memory storage when administered at short but not long delays, after training. Somewhat remarkably, then, these foundation experiments showed that memory storage processing can be enhanced by some treatments and impaired by others. Actually, as you will see below, it now appears likely that all treatments which act on memory storage can either enhance or impair memory, depending on factors such as dose and the specific nature of the training situation. With this background, some laboratories began to examine the neurophysiological mechanisms by which various treatments either enhanced or impaired memory. In addition, some investigators selected specific treatments (e.g. protein synthesis inhibitors) to determine whether protein synthesis was necessary for memory storage. Both classes of research used the premise that if one could determine the mechanisms of memory modification, one might learn an important characteristic of memory storage itself.

The principle general finding of the past 10 years has been that recent memories have a most remarkable feature of being sensitive to a wide range of apparently disparate treatments. Why might that be the case? In an evolutionary sense, what advantage is there in a brain design in which memories can be so readily modified?

Most obviously the advantage is that the brain receives additional information which defines whether the recent events are worthy of storage—i.e., whether the experience was important or trivial. For important events, the additional information can direct the brain to store the new memories. We call this directive process *memory modulation* or *regulation.*

One way to characterize the factors underlying memory regulation is to consider arousal as a major contributor toward telling the brain when to make new memories. The possibility that arousal levels interact with learning and memory is not at all new. An old law in psychology states that the relationship between arousal and performance follows an inverted-U function in which low and high arousal levels are related to poor learned performance and moderate arousal levels linked to good performance. However, although the relationship appears to hold true in a wide variety of circumstances, it is difficult to study "arousal" as a mechanism for the neurobiology of memory regulation because it cannot be measured with certainty and because the relationship results in a tautology: important events are arousing and trivial events are boring.

Alternatively, it should be possible to identify physiological responses to experiences which contribute to memory storage processing. In other words, some physiological concommitants of arousal should follow important events and direct storage of information contained in those events. Likely candidates include those hormones often associated with stress (high arousal). As the reader will soon see, there is now a great deal of evidence which suggests that some of the endocrine responses to experiences do indeed regulate memory storage.

The evidence for hormonal regulation of memory storage is best illustrated by describing a typical laboratory experiment. A one-trail behavioral training procedure is particularly useful in these studies. The animals are placed in a two-compartment rectangular box, with a well-lit white start compartment separated by a door from a dimly-lit black compartment which has a metal floor through which a brief electric shock can be administered. The animals are placed in the white compartment, and after a few seconds, the door is opened. Generally, the animals cross into the darker compartment after only a few seconds. As soon as they have fully entered the dark compartment, the door is closed and the animals receive the footshock. They are then returned to their home cages for a day before they are tested for their memory. On the test trail, each animal is placed in the white compartment, as before, and the door is opened. If the animals remember what happened to them before, they demonstrate that memory by being reluctant to enter the shock compartment. The latency to enter the black compartment is then an indication of the strength of the animal's memory.

Not surprisingly, an animal's latency to re-enter the shock depends on the intensity and duration of the shock it received on the training trial. The first question is whether the increase in memory is dependent not only on the actual shock level but, in part, on the different hormonal responses to the shock. To test the latter possibility, animals are trained using a very mild footshock and injected with one of the hormones which would have been released by a more intense footshock. In such an experiment, then, we artificially create a physiological situation in which the animal receives a mild footshock during training but has the hormonal consequences of a stronger shock. When the animal is tested the next day, it behaves as if it had received a more intense footshock during training, i.e., the retention latencies are longer than they would otherwise be.

## Epinephrine Modulation of Memory Storage

Several hormones, such as epinephrine (adrenaline), ACTH, and perhaps vasopressin, have effects such as these. It would be misleading to convey the impression to the reader that epinephrine is the hormone most important to memory storage. One reason that we have more evidence for epinephrine is because there are better procedures for measuring this hormone, and also because there is a wealth of relatively specific drugs with which to manipulate epinephrine actions. As the technology for studying other hormones improves, we will undoubtedly gain further insights into the mechanisms by which they, too, act on memory.

When an animal is trained as described above, an injection of epinephrine enhances memory upon later testing. The specific dose is critical to the phenomenon. A full dose response curve is shown in Figure 1. The important feature is that the shape of the curve is an inverted-U, just as one might predict if the hormonal mileau contributes to the law described earlier. It is one thing to show that a drug has effects on memory, but quite another to try and show that the hormone has endogenous actions on memory. We want to show that animals trained with a higher footshock, and which, therefore, have a better memory for the immediately preceding training experience, release more epinephrine than animals trained with a mild footshock. Furthermore, the optimal epinephrine dose for enhancing memory should, if we are truly mimicking the nature response, produce hormone levels which are comparable to those found in animals trained with a higher footshock. Results confirming these possibilities are shown in Figures 2 and 3.

As you can see in Figure 2, plasma epinephrine levels increase as the footshock training level increases. In particular, the plasma epinephrine levels reach about 1500 pg/ml after the higher footshock intensity. Remember this number as you look at Figure 3. Here, note that the dose optimal for memory enhancement also results in plasma levels which are about 1500 pg/ml. Thus, the dose which the behavioral studies tell us is best for enhancing memory is the same dose that most closely matches the normal physiological response to training.

**Figure 1**

*A full dose response curve.*

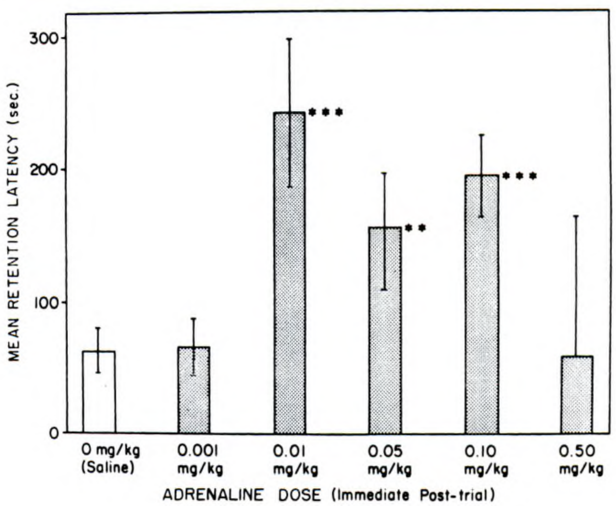


A variety of other experimental strategies indicate that the role of epinephrine in regulating memory storage has remarkably general properties. For example, epinephrine enhances memory for many different kinds of tasks, including tasks which train the animals using rewards as well as those which use punishment. Epinephrine also can be used to ameliorate memory which is otherwise poor under a variety of conditions. In one study, we found that animals could even learn under deep anesthesia if first injected with epinephrine. Anesthetized animals received paired presentations of a tone which ended at the same time that electric shock was presented to the skin. The animals were tested several days after training by presenting

**Figure 2**

*Plasma epinephrine levels.*

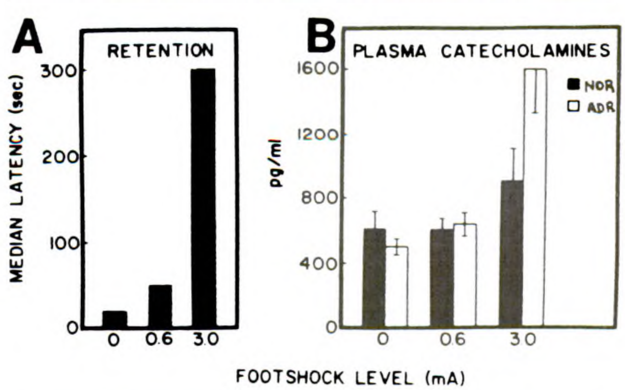

**Figure 3**

*Optimal dose for memory enhancement.*

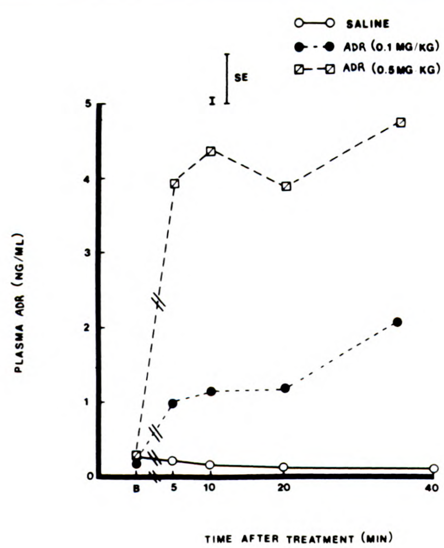


the tone while the animals, which had been deprived of water for a time, were engaged in drinking. Animals which hear a tone they have associated with an aversion stimulus stop drinking. Those animals which had received saline prior to the training under anesthesia showed no evidence of fear when the tone came on. However, the animals which received epinephrine ceased drinking, and in the most dramatic instances, went to the corner of the cage farthest from the speaker and froze in one position throughout the minute the tone was administered.

## Mechanisms of Epinephrine Actions on Memory

The effects of epinephrine on memory indicate that one of the hormone's actions must be to modify brain processes. However, for a long time it has been clear that the hormone itself does not pass from blood into brain cells, at least not in appreciable and measurable amounts. We are faced with an apparent paradox then; the hormone has large effects on brain function but does not act on the brain—at least not directly. How then can epinephrine accomplish its memory enhancing tasks? Over the past year, we tested the possibility that one of the hormone actions outside of the brain might be responsible for epinephrine modulation of memory storage processing. The action we picked is one of the classic physiological effects of the hormone, the release of glucose from liver stores. Reviewing the physiological responses of an animal to stress, we know that epinephrine is released and that one of the consequences of epinephrine release is to initiate the release of glucose from the liver. We, therefore, examined the effects of glucose on memory in rats. The results, shown in Figure 4, demonstrate that glucose too can enhance memory. The dose-response curve is again an inverted-U and the effects are time-dependent; injections delayed by an hour after training have no effect on later retention performance. These are all features in common with the epinephrine effects on memory. There is more evidence, as well, suggesting that plasma glucose may mediate epinephrine's effects on memory and, furthermore, that glucose must then serve to control the neuronal processes responsible for memory storage.

**Figure 4**

*Effects of glucose.*

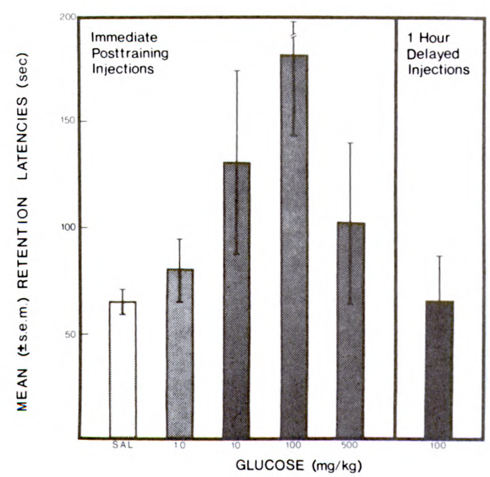


In identifying biological components of memory storage it is becoming more and more evident that a full understanding of the neuronal basis for behavioral changes in response to experience will require the intergration of physiological systems which include not only the brain, but many systems below the neck as well.

## Authors

Paul E. Gold is Professor of Psychology at the University of Virginia. His international reputation as a leader in psychopharmacological research on the biological basis of memory is based, in part, on the authorship of over 85 refereed articles and publications in the open scientific literature.

Joel L. Davis is an ONR scientific officer and Adjunct Associate Professor in the Department of Psychiatry and Biobehavioral Sciences at the U.C.L.A. Medical School. He continues his research on neuropeptides and memory at the V.A. Medical Center in Sepulveda, California.

# The UCSB Free Electron Laser: A Tool for Scientific Research

Gerald Ramian and Luis Elias

Quantum Institute
University of California
Santa Barbara, CA 93106

## Introduction

A unique electrostatic-accelerator-driven Free-Electron-Laser (FEL) is currently in operation at the University of California, Santa Barbara (UCSB) under the sponsorship of the Office of Naval Research (ONR). This laser produces tunable coherent radiation from 400 $\mu$m to 110 $\mu$m at kilowatt power levels and is currently being used in a program of condensed matter physics research.[1] Figures 1 through 4 are photographs of this facility.

The recent award of a five year University Research Initiative (URI) grant, again under sponsorship of ONR, has permitted establishment of a Center for Free-Electron Laser Studies (CFELS) at UCSB. A major part of this program will be operation of a users' facility where FEL radiation is used for research in condensed matter physics, biology, medicine, and materials science. Figures 3 and 4 show the recently completed users' facility building funded by the Strategic Defense Initiative (SDI) Biomedical Program. Other parts of the CFELS program include continued development of novel FEL sources, work in FEL theory, and graduate student education and training.

Unlike conventional lasers using bound electrons in atomic states or molecular vibrational and rotational states, an FEL uses a relativistic electron beam in conjunction with a periodic magnetic field as the lasing medium. Limits on efficiency and non-linear effects, such as thermal depopulation and self-focusing within the lasing medium, are thus transcended. It is important to note that lasing is not limited to a few discreet wavelengths. Continuous tunability is possible by varying electron beam energy.

Development of the UCSB FEL, over the past six years, involved several major accomplishments. First, construction of an ampere level electron gun with near thermally limited emittance in 1981 improved the then current practice by an order of magnitude. Acceleration, recirculation, deceleration, and collection of an ampere level beam at more than 90% efficiency by an electrostatic accelerator permitted high efficiency and high duty cycle operation at relatively low cost. Development of a novel hybrid waveguide resonator permitted a high undulator field strength with the very low losses required for submillimeter lasing. Finally, demonstration of lasing at very narrow linewidth and, perhaps most importantly, demonstration of the FEL as a practical tool for scientific research were accomplished.

**Figure 1**

*The UCSB Free-Electron Laser. Tunable coherent submillimeter radiation at kilowatt power levels is produced. The six megavolt accelerator tank is in the background. Electrons leaving the bottom of the accelerator tank move around the beam line—clockwise in this view, pass through the resonator on the left where a fraction of their energy is converted to radiation, and finally return to the tank where they are decelerated and captured.*

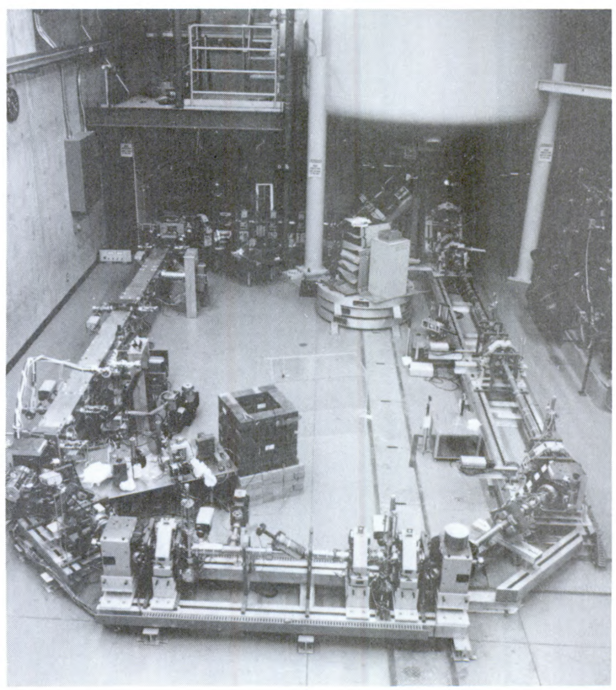

**Figure 2**

*Control Console. Remote control and diagnostics of FEL operation are necessary for personnel x-radiation protection. An LSI-11 based computer control system interfaces all machine functions with the user, permitting machine operation by relatively inexperienced users.*

**Figure 3**

*UCSB FEL laboratory building. The FEL laboratory containing the 6 MV accelerator is seen in the background of this view. The recently completed users' facility laboratory, built around the FEL laboratory, is in the foreground.*

**Figure 4**

*UCSB FEL laboratory building. 2 foot thick walls provide x-radiation shielding. A temporary user facility building will be built in the foreground of this view.*

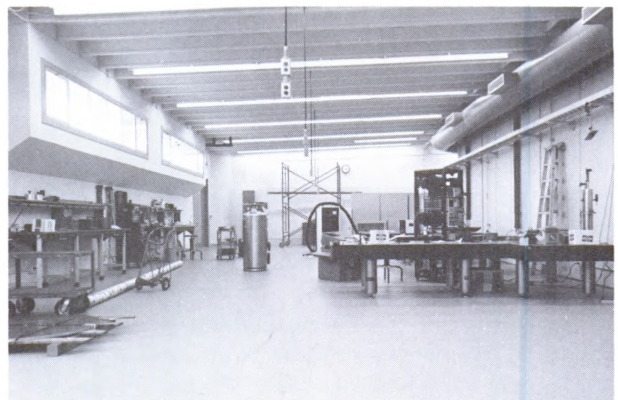

## History

FEL operation was first demonstrated in 1977 at Stanford University under the direction of John Madey,[2] who is due much credit for the considerable theoretical and experimental skill culminating in that demonstration. The authors were participants in that project.

The FEL represents a confluence of microwave and optical technology. In fact the major difference between an FEL and devices such as Phillip's Ubitron[3] is the use of a relativistic electron beam and an optical resonator. Similarly, FEL operation involving a pre-bunching undulator is referred to as an "optical klystron" and is analogous to the microwave klystron. Other types of FEL such as the Smith-Purcell or Čerenkov type achieve amplification using slow-wave structures in analogy with the traveling-wave tube. In all cases, an energy exchange between free electrons and coherent electromagnetic energy takes place.

The theoretical underpinnings of the FEL include the work of Motz and Nakamura in 1951 at Stanford[4] who first suggested the generation of coherent radiation from relativistic beams and the work of Pantell, Soncini, and Puthoff in 1968 at Stanford[5] who suggested the use of Compton scattering to upconvert microwave energy. These precepts and the concept of the undulator as a source of virtual photons under the Weizsacker-Williams approximation permitted Madey[6] in 1971 to calculate gain and efficiency and realize the possibility of the FEL. Rapid progress followed the demonstration of amplification in 1976[7] and especially the demonstration of oscillation in 1977.[2] Today several hundred published papers by an international body of scientists explore the FEL from a variety of theoretical perspectives.

**Figure 5**

*Schematic representation of an FEL. Radiation from a relativistic beam that is subjected to a periodic transverse acceleration by a magnetic undulator is trapped between mirrors in an optical resonator. As intensity builds, bunching of the beam and interference effects produce coherent high power radiation.*

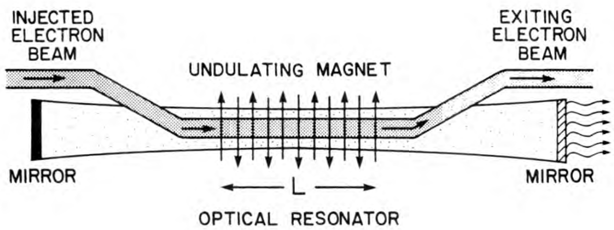


## Basic Principles

Operation of an FEL involves amplification of a coherent optical wave by a relativistic electron beam in the presence of a spatially periodic magnetic field (Figure 5). The radiation process in conventional lasers is inherently quantum mechanical because transitions involved in the creation of a photon involve small numbers of energy levels. Based on this precident, initial theoretical work including Madey's first theoretical analysis of the FEL used quantum-mechanical models such as Compton backscattering or magnetic Bremsstrahlung. It was soon recognized that at typical FEL operating points, transitions cover thousands of energy levels making classical analysis totally adequate. Classical theories are generally divided into two categories depending on electron-beam current density. Operation at high densities ($I_b \approx 1000$ Amp), where space charge instabilities are significant, is refered to as the Raman regime and treatment of the beam as a cold relativistic plasma governed by Boltzmann's dispersion relationships has proven fruitful. Operation at low beam density ($I_b \approx 1$ Amp), typical of the UCSB FEL, is referred to as the Compton regime and self consistent application of Maxwell's wave equations to single particle dynamics has proven most useful.

In the simplest possible terms, the FEL can be described as follows: An electron traversing a periodically reversing magnetic field experiences as transverse acceleration resulting in the familiar dipole radiation pattern on Figure 6a. Since, however, the election is relativistic, this radiation pattern will be compressed into a narrow forward radiation cone (Figure 6b) of half angle $\approx 1/\gamma$, where $\gamma$ is the Lorentz factor $\sqrt{1/(1-\beta^2)}$. Here, $\beta$ is the relativistic velocity $v/c$. The wavelength of the radiation will also be doppler-upshifted by the ratio $\beta/(1-\beta)$. This radiation, refered as "spontaneous emission", is trapped between mirrors in a more or less conventional laser resonator. As optical power builds, the electrons experience a "ponderomotive" force which bunches them longitudinally. As bunching proceeds, optical interference narrows linewidth and further increases intensity in a regenerative manner until saturation is reached. Radiation generated in this manner is referred to as "stimulated emission" and has all the wonderful properties of conventional laser light. Tunability comes about through the variation of $\beta$ with electron-beam energy. Saturation occurs when the optical field builds to such intensity that a significant number of electrons are trapped and overbunched within the ponderomative potential.

**Figure 6a**

*Normal dipole radiation patter of a transversely accelerated election.*

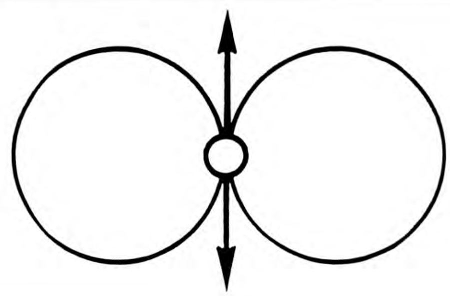

**Figure 6b**

*Relativistically transformed dipole radiation pattern. Radiation is compressed into a narrow forward cone and dopple upshifted in frequency.*

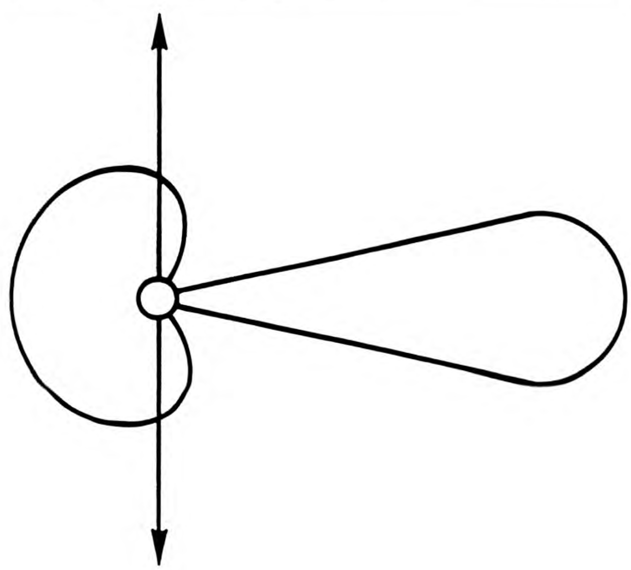

## Electron-Beam Sources

A wide variety of electron-beam sources have been employed in FEL experimental demonstrations. These include RF Linacs, Microtrons, Induction Linacs, Storage Rings, and at UCSB an Electrostatic Accelerater. Each machine has specific advantages and disadvantages and unfortunately the choice has often been determined by the availability of a specific machine to an experimental group. The possibility of using an electrostatic accelerator with recirculation to enhance efficiency was recognized immediately by Madey[6] and later expanded upon by Elias[8] at the beginning of the UCSB program. Recirculation is an important part of this concept for two reasons. First, single-pass efficiency of the FEL is of the order of a few tenths of a percent. An FEL providing kilowatts of average power would consume megawatts if its electron beam were entirely discarded after one pass. Second, the limited available charging current of a few hundred micro-amps would limit generation of an ampere beam to a few micro-seconds and a duty cycle of a few hundredths of a percent. A recirculation level of 99% extends these parameters by two orders of magnitude. Initial recirculation experiments in 1982 achieved this level of recirculation and the FEL has been recently running the recirculation between 90 and 95 percent. An even more important advantage is the generation of a very high-quality electron beam without the RF microstructure characteristic of linacs and microtrons. As a result, single-mode operation of extremely narrow bandwidth and long coherence length is possible and has been demonstrated by the UCSB FEL.[9] This is of considerable adevantage in such diverse areas as spectroscopy, selective excitation of atomic or molecular states, and even radar.

## Applications

As mentioned, the UCSB FEL is tunable and capable of operating at extremely narrow linewidth as well as high average power. Such sources have been heretofore unavailable in the submillimeter region. Historically, rapid progress has followed the development of sources in every region of the electromagnetic spectrum. Examples include the consequences, to communications and radar, of development of the magnetron and klystron. Even more dramatic are the unanticipated applications of lasers in the visible and near infrared that impact almost every aspect of our lives. Among anticipated applications of the FEL are:

*—condensed matter physics;* The unique tunability, high power, and extremely high resolution of the UCSB FEL will permit the exploration of many interesting phenomena. Of particular interest are low-energy excitations such as magnons, phonons, plasmons, and polaritons that occur in most solids in the submillimeter region. Semiconductor band-gap energies also correspond to submillimeter wavelengths opening many applications for the FEL in solid-state physics research. In fact, Bell Labs has on ongoing project to develop an FEL for this purpose.[10] Figure 7 indicates some of the excitations in relation to the present UCSB FEL and the planned two stage FELs. Experiments planned for the near future with the UCSB FEL include studies of the nonlinear excitation of magnons in antiferromagnetic materials, studies of the intrinsic nature of phonon relaxation processes, studies of the non-linear excitation of 2 dimensional electron semiconductor systems, and the study of the dynamics of localized impurity modes.

*—chemistry;* The FEL will find many applications in chemistry. Tunability and high power will premit high-efficiency laser isotope separation. Selective enhancement or degradation of chemical bonds and shifting of reaction equilibrium points will refine many processes that now crudely rely on elevated temperatures. Surface science will also benefit enormously where, for example, selective molecular excitation may help achieve a better understanding of catalytic reactions. Understanding the chemical reactions between metals or polymers and their environments—atmospheric, oceanographic, or biological—could reap enormous economic benefit.

*—biology and medicine;* Just as the laser has found many medical applications in the visible and infrared, the same can be expected in the submillimeter region. Tumor treatment, either through direct decomposition or with a chromophoric agent has received considerable attention. Extension of the techniques to submillimeter wavelengths with their greater penetration depth and the discovery of new agents may make this a major cancer treatment technique. Degradation of plaque in atherosclerotic blood vessels might be a similar application. It is also being recognized that many biological processes are wavelength-dependent. For example, resonances in the submillimeter region may play a part in enzyme reactions or even the unwinding of DNA during replication. Studies of this are scheduled for the UCSB FEL later this year.

*—optical materials;* The UCBS FEL's tunability and high power will permit characterization of materials in the submillimeter region in an unprecedented way. Many materials have been characterized in terms of index of refraction and adsorption coefficient with thermal sources and Fourier Transform Spectrometry, but measurements of electro-optical coefficients or non-linear resonances have not been possible at more than a few discrete wavelengths with heretofore-available sources. Among the UCSB FEL's highest priorities will be the search for, and characterization of, linear materials suitable for windows, lenses, and AR coatings and non-linear materials suitable for mixers, multipliers, and phase conjugators in the submillimeter region.

*—communications;* Atmospheric absorbtion of submillimeter radiation is particularly strong and would seem to limit communication to space-based applications. There, high directionality and moderate antenna size requirements might be useful in providing secure intersatellite communications. Development of heterodyne receivers and

miscellaneous components such as antennas, waveguides, mixers and phase shifters will probably be the FEL's main contribution in this area.

*–submillimeter radar;* The FEL's high power and long coherence length make an attractive source for submillimeter radar. Wide spatial bandwidth and rapid tunability would be especially useful in signature analysis. Again, high atmospheric absorption would limit surface applications to short range such as a terminal defense radar where the high atmospheric absorption would help in providing a clutter-free background. Development of very compact electrostatic accelerator FELs suitable for space applications is a promising possibility.

**Figure 7**

*Condensed matter excitations. The shaded areas indicate the operating range of the present UCSB FEL and the future 2 stage FELs. Dotted lines indicate the range of the present FEL may be extended to 50 $\mu$m with construction of a new undulator and to 1 mm with further beam transport development.*

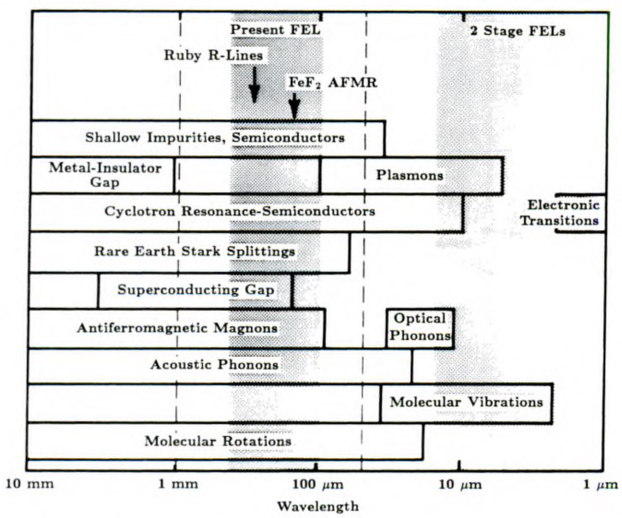


## Current and Future Developments

As one of only two currently working FELs (December 1986) and the only one available for experimental work in the world, the UCSB FEL presents a unique opportunity to accomplish research objectives in a broad range of disciplines as well as to accomplish some direct mission objectives of the defense agencies. Under the CFELS program, emphasis is therefore shifting from pure FEL development to a program balanced between that, users' facility activities, and educational activities. A 3000 ft$^2$ users' facility building funded by the Strategic Defense Initiative (SDI) Biomedical Program has recently been completed and funding has been obtained to partially equip this facility under a Department of Defense (DOD) instrumentation grant. This building provides specialized areas where a number of users can access FEL radiation for research in physics, chemistry, biology, and medicine. The laboratory is also being equipped with a wide range of ancillary instrumentation including a Fourier transform spectrometer, various conventional lasers, cryogenic equipment, and a large selection of optical components to assist the experimenters.

FEL development work includes a two-stage FEL expected to produce 1 $\mu$m radiation. Operation of its first stage, which will generate megawatts of power at 700 $\mu$m, is expected later this year. Figure 8 is a photograph of its helical undulator. Joint work with Hughes Research Corp. on a two-beam two-stage FEL that will optimize generation of radiation at each wavelength and provide tunable range of 3 to 30 $\mu$m is also being undertaken. A program to develop undulators with periods of a few millimeters has resulted in a "micro-undulator" structure (Figure 9) that is currently undergoing testing. Use of this micro-undulator with standard 1 to 2 MV electrostatic accelerators would permit development of very compact submillimeter FELs (Figure 10) or conversely, with commercially available 15 or 25 MV accelerators (Figure 11), development of visible or near infrared FELs with megawatt power output. Funding has been requested to begin development of the necessary technology using a 2 MV test accelerator. An effort to further improve the operating characteristics of the present FEL is anticipated through the use of injection locking techniques. A high resolution electron gun analyzer is also under development as part of the DOD instrumentation program.

**Figure 8**

*2-Stage undulator. The electro-magnetic undulator generates a 200 Gauss helical field with 20 cm period. With a 20 Amp, 6 MV beam, it will generate 700 $\mu$m radiation as the first stage of the UCSB 2-stage FEL.*

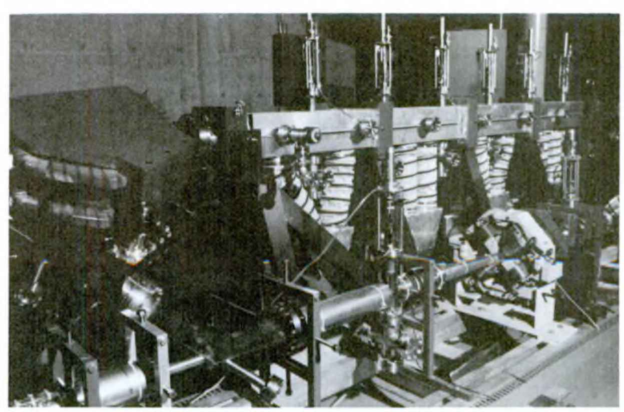

**Figure 9**

*Micro-undulator. Grooves ground in blocks of homogeneously magnetized Neodymium-Iron-Boron generate undulator fields while transcending the mechanical problems associated with conventional construction from individual magnets.*

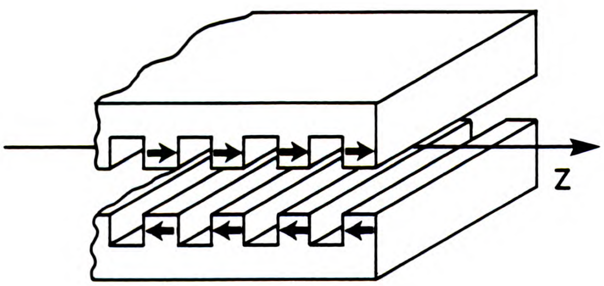


**Figure 10**

*Compact 1 and 2 MV FELs. The small dimensions of a micro-undulator permit its installation within the accelerator. This requires positive terminal operation with the electron gun, collector, and their associated electronics located outside. An additional contribution to small size comes about through the higher insulating gas electric field gradient permitted with positive terminal operation.*

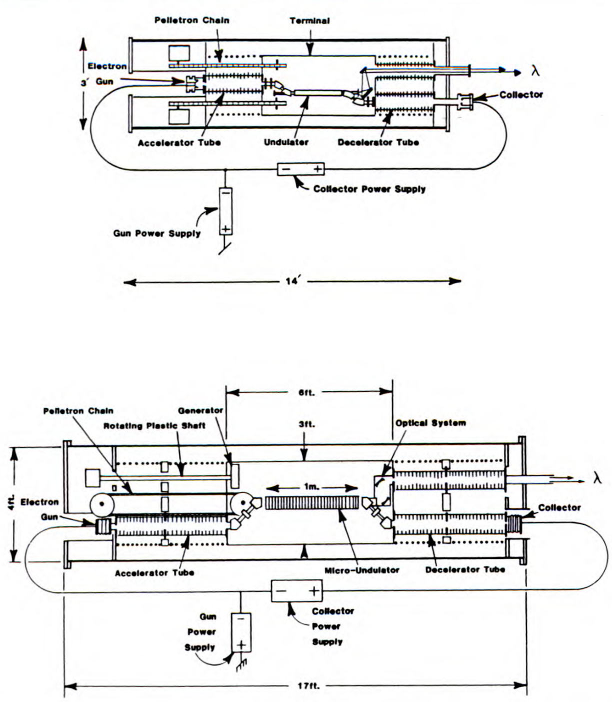


**Figure 11**

*15 and 25 MV FELs. Using the micro-undulators with commercially available 15 or 25 MV electrostatic accelerators would permit operation at near infrared or visible wave-lengths and megawatt power levels.*

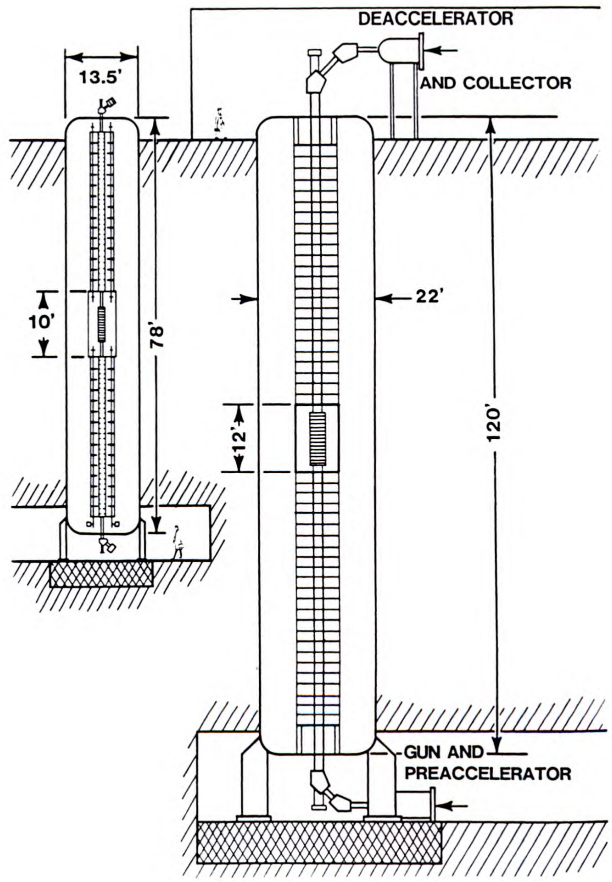




☆ U.S. GOVERNMENT PRINTING OFFICE: 1987 - 176-764 - 814/60310

## Conclusion

The FEL can be thought of as a device bridging the gap between microwave and optical technology and, as such, has opened an important part of the electromagnetic spectrum in the far-infrared and submillimeter wavelengths. Development of techniques to use an electrostatic accelerator with recirculation as a beam source has permitted the FEL's development into a practical research tool as demonstrated at UCSB. Startup of a users' facility at UCSB will shortly permit research in physics, chemistry, biology and medicine. Evolution of the UCSB FEL concepts into very compact submillimeter FELs or, with high voltage accelerators, visible or near-infrared FELs through the micro-undulator concept is anticipated. Coverage of the near-infrared region through the 2 stage FEL concept is expected within the next two years. Improvements in the present UCSB FEL through such techniques as startup injection locking are also anticipated.

## References


**1.** L. Elias, V. Jaccarino, and W. Yen, "Scientific research with the UCSB Free Electron Laser," Proceedings of the Workshop on Applications of Free Electron Lasers, 1984, Castelgandolfo, Italy, p. 439, North-Holland, 1985

**2.** D. Deacon, L. Elias, J. Madey, G. Ramian, H. Schwettman, and T. Smith, "First Operation of a Free Electron Laser," *Physical Review Letters,* v. 38, p. 892, 1977

**3.** R. Phillips, "The Ubitron, a High Power Traveling Wave Tube Based on a Periodic Beam Interaction," *IRE Transactions on Electron Devices,* v. 7, p. 231, 1960

**4.** H. Motz, W. Thon, and R. Whitehurst, "Stimulated Photon-Electron Scattering," *Journal of Applied Physics.* v. 24, p. 826, 1953

**5.** R. Pantell, G. Soncini, and H. Puthoff, "Stimulated Photon-Electron Scattering," *IEEE J. Quantum Electronics,* v. 4, p. 905, 1968

**6.** J. Madey, "Stimulated Emission of Bremsstrahlung in a Periodic Magnetic Field," *Journal of Applied Physics,* v. 42, p. 1906, 1971

**7.** L. Elias, W. Fairbanks, J. Madey, H. Schwettman, and T. Smith, "Observation of Stimulated Emission of Radiation by Relativistic Electrons in a Spatially Periodic Transverse Magnetic Field," *Physical Review Letters,* v. 36, p. 717, 1976

**8.** L. Elias, "High-Power, CW, Efficient, Tunable Free Electron Laser Using Low Energy Electron Beams," *Physical Review Letters,* v. 42, n. 15, April, 1979

**9.** A. Amir, L. Elias, J. Hu, G. Ramian, "Observation of Single Mode Operation in an FEL," *Physical Review Letters,*

**10.** E. Shaw, R. Emanuelson, and G. Herbster, "Far Infrared Free Electron Laser," *Journal de Physique,* C1-1983, p. 385


## Authors


Luis Elias is Professor of Physics at UCSB and Gerald Ramian is a Specialist with the UCSB FEL program. Both have been active in Free Electron Laser research for the past 10 years including participation in the first FEL experiments at Stanford. Both have numerous publications in the field of FEL physics

DEPARTMENT OF THE NAVY
OFFICE OF NAVAL RESEARCH
ARLINGTON, VA. 22217

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, $300

THIRD-CLASS MAIL
POSTAGE & FEES PAID
USN
PERMIT No. GS-9