

# AI Research at the Artificial Intelligence Laboratory, Massachusetts Institute of Technology

Patrick H. Winston,

*Director of the Laboratory*

THE PRIMARY GOAL of the Artificial Intelligence Laboratory is to understand how computers can be made to exhibit intelligence. Two corollary goals are to make computers more useful and to understand certain aspects of human intelligence. Current research includes work on computer robotics and vision, expert systems, learning and common-sense reasoning, natural language understanding, and computer architecture.

Professor Patrick H. Winston and Professor Marvin Minsky lead research efforts in learning. Professor Robert C. Berwick studies learning in the context of natural language syntax acquisition. Professor Berthold K. P. Horn, Professor Tomaso Poggio and Professor Shimon Ullman lead efforts in computer vision. Dr. J. Michael Brady, Professor John M. Hollerbach, Professor Tomas Lozano-Perez and staff work on other aspects of Robotics. Professor Gerald J. Sussman and Mr. Thomas F. Knight work on the problems of integrated circuit design. Professor Randall Davis and Dr. Howard E. Shrobe work on expert systems that use functional and physical models. Dr. Charles Rich and Dr. Richard C. Waters explore the creation of intelligent programming environments. Mr. William D. Hillis and Mr. Knight explore the development of the connection memory machine, a special-purpose machine for concurrently manipulating knowledge stored in semantic nets. Professor Carl E. Hewitt studies distributed problem-solving and parallel computation. Mr. Richard Greenblatt works on microprocessor networks.

The Laboratory's 120 members include 12 faculty members, 5 academic staff, 45 research and support staff, and 58 graduate students active in research activities funded by the Defense Advanced Research Projects Agency, the Office of Naval Research, the Air Force Office of Sponsored Research, the National Science Foundation, the System Development Foundation, Martin Marietta Corporation and Digital Equipment Corporation.

## **Robotics and Vision**

Professor Tomas Lozano-Perez, Dr. Rodney A. Brooks and Dr. Matthew T. Mason, in collaboration with Dr. Russell H. Taylor of IBM, are designing a software system for automatic synthesis of robot programs from high-level task specifications. The system incorporates results in reasoning about space and object interaction.

Professor Lozano-Perez and Dr. Brooks have extended previous results on obstacle avoidance to deal with paths requiring nearly continuous rotations. Dr. Brooks has developed an efficient obstacle avoidance algorithm for the special case of convex objects. Dr. Brooks also has proposed a method for checking and modifying robot plans to ensure that they will work, given mechanical errors in placement and orientation of objects. The method can be extended to the generation of the robot plans. Dr. Mason has developed a theory of the behavior of objects moving under friction. The

theory has led to powerful methods for accurately positioning objects without the use of sensory feedback.

Professor John M. Hollerbach, Mr. William D. Hillis and Mr. John A. Purbrick has continued to develop high-resolution touch sensors, and emphasis on the use of new, more rugged materials. Professor Hollerbach also designed and implemented a new controller for a three degree-of-freedom tendon-driven shoulder. The controller, based on a model of muscle behavior, achieves human-like control of manipulator stiffness. Professor Hollerbach and Ms. Tamar Flash have studied the interaction forces during human reach movements. Information collected from arm movements is used to compute various contributions to the torques. They have found evidence that humans use trajectory-formation strategies that simplify computation.

Professor Huruhiro Asada completed the design of a direct-drive manipulator. This innovative electro-mechanical manipulator uses no gears in its construction, which makes it very accurate and fast and, hence, ideal for assembly tasks. Professor Warren Seering and his students are investigating novel actuators for manipulators.

Dr. J. Michael Brady has developed a new representation for visual contours based on local symmetry. Dr. Brady and his associates have implemented a software system for processing an image to detect object contours, obtain a representation of the contours, and use the representation to choose where the object should be grasped. This work is aimed at constructing a system that can pick an object out of a bin of similar objects.

Dr. Brady has also developed a new representation for two-dimensional shape based on extending local symmetries to form smooth axes. Dr. Brady and his associates have implemented a software system for processing an image to detect object contours, obtain a representation of the contours, and use the representation to determine its smoothed local symmetries. The representation has been used to choose grasp points on a two-dimensional volumes and relies upon a further representation, called curvature patches, for three-dimensional surfaces. Mr. Demetri Terzopoulos has significantly developed earlier work by Dr. W. Eric L. Grimson, Dr. Brady and Professor Horn on interpolating smooth surfaces. Drs. Brady and Grimson have proposed that the theory of surface interpolation underlies the human perception of subjective contours.

In other vision-related work in Robotics, Professor Berthold K. P. Horn and his staff study the problems associated with the relationship between the shapes of visible surfaces and their gray-level shading. In particular, Professor Horn has formulated the image-irradiance equation, which relates surface orientation to brightness values, and he has developed the reflectance map, which graphically displays the constraint implied by the image-irradiance equation.

In general, the image-irradiance equation has an infinite number of solutions, each corresponding to a surface. Thus, shading alone is not sufficient to determine a surface uniquely. Working with Professor Horn, Dr. Anna R. Bruss has ex-

amined the use of supporting information from silhouettes, edges, and rotational symmetry, showing how to arrive at unique solutions in some cases. In particular, she has determined when the images captured by scanning electron microscopes determine surface shape uniquely. Professor Horn and Mr. Brian G. Schunck have continued to investigate the use of the image-irradiance equation in determining the instantaneous velocity field which has generated on the image plane by movement in a textured environment. One important application is in navigation, where it is desirable to determine camera motion from noisy image sequences.

Another application of the image irradiance equation is in the interpretation of multiple images obtained using differing lighting conditions. Mr. William M. Silver has implemented a multiple-image method to determine the orientation of each patch on the surface of an object. Mr. Philippe Brou presently is using this kind of information to determine the identity of an object and its attitude in space.

In another effort, Professor Tomaso Poggio, Dr. H. Keith Nishihara, and Dr. Kenneth Nielson have developed a theory of spatio-temporal interpolation extending the classical sampling theorem. Work continues on developing efficient algorithms to process time-varying imagery that maintain high spatio-temporal resolution and avoid motion smear. The analysis of time-varying imagery is an important focus of research, since any visual processor with human-like performance must be capable of processing moving objects.

The measurement of the optical flow — the field of velocities in the image — is another important visual task. In collaboration with W. Reichardt and H. Bulthoff of the Max-Planck Institut für Kybernetik in Tübingen, West Germany, Professor Poggio has devised a model that uses intensity information in simple visual systems to compute directional motion and detect discontinuities in the optical flow. Professor Shimon Ullman and Ms. Ellen C. Hildreth have investigated the computation of the optical flow by measuring velocity along contours in the image.

The Marr-Poggio theory of human stereo vision continued to represent a major focus of research. The understanding of stereopsis and its implementation in a real-time stereo machine with human-like performance has important applications especially in Robotics. Furthermore, stereopsis seems to require computations, such as edge detection, that are basic to many other visual tasks. Dr. Nishihara and Mr. Noble G. Larson have developed special hardware for doing some of these computations in real time. Professor Poggio and Dr. Nielson have compared the abilities of various stereo algorithms developed by Dr. Grimson, Dr. Nishihara, Professor Poggio and Mr. Michael Kass.

Dr. Grimson has developed algorithms for deriving a representation of a smooth surface given the incomplete sort of information that stereopsis provides. Professor Shimon Ullman and Mr. Terzopoulos explore the use of multigrid algorithms for efficient surface interpolation and other computations.

Professor Ullman has also investigated the computation

employed by the visual system in the perception of properties and relations such as closure, insideness and connectedness. Understanding these computations is a step toward the development of a visual processor for object recognition.

Finally, Professor Poggio has proposed a new theoretical approach to the computational mechanisms used by the brain to process information, which he calls the Biophysics of Information Processing. Using an extension of cable theory to describe nonlinear interactions, Professor Poggio has shown that a single dendritic tree may perform hundreds of independent analog operations on its inputs. The theory suggests that specific branching patterns reflect the information processing performed by the neuron. A dendritic tree is then close to the ultimate in microelectronics, with patches of membranes representing the fundamental units for elementary computations.

Overall, the Laboratory has expanded work in Robotics considerably. One facet of the enlarged robotics effort has been to host prominent researchers in Robotics for visits ranging from six months to a year. Another facet is to organize and hold intensive three-day workshops on topics in Robotics such as the design of dexterous grippers and the design of high-level programming languages. A third facet of the program is to encourage wider dissemination of results in Robotics. As part of this effort, Dr. Brady and Professor Richard Paul of Purdue University founded "Robotics Research", a new journal published by the MIT Press.

### Expert Problem Solving

Professor Gerald J. Sussman and his associates have made considerable progress in their research into techniques and tools for very large scale integrated circuit (VLSI) design.

In one project, Mr. Gerald L. Roylance is working on a compiler for circuits. A circuit compiler is similar to a programming language compiler except that instead of compiling an abstract language to a concrete machine language, it compiles an abstract circuit design to a concrete one consisting of resistors, capacitors and transistors. For example, an analog circuit design might use an abstract amplifier; the compiler might transform that amplifier into a transistor cascade circuit. In order to do this transformation, the circuit compiler must understand how the parts of the original circuit specification interact to constrain each other. These constraints are found by constraint propagation, a well-established artificial-intelligence technique. While analog circuit design is the domain of the Roylance compiler effort, the ideas about building circuits also apply to VLSI design, digital logic design and other engineering domains such as numerical programming.

In another effort, Mr. Daniel Weise has developed and implemented a system for the interactive and incremental extraction of circuit schematics from hierarchically represented layouts of VLSI circuits. This system enables our design

system to understand and answer questions about circuits represented by layouts created using the design system.

Mr. Neil L. Mayle also has been working on simulation, investigating a method of incremental switch level simulation. A simulator of this sort allows the designer to simulate a cell when it is first designed. The information derived is memorized, thus saving the work of recomputing it during subsequent simulations.

In other VLSI-related work, Mr. James J. Cherry has been investigating circuit simulation algorithms suitable for implementation on parallel/pipelined architectures. Mr. Bruce E. Edwards has been working with Mr. Cherry on writing an electrical simulator.

Professor Randall Davis's work has been focused in two main areas: model-based expert systems and the application of Artificial Intelligence to signal processing.

Professor Davis, Dr. Shrobe and associates are building an expert system that uses knowledge about structure, function, and causality to do fault diagnosis in digital electronics. Previous expert systems have typically been built from large collections of empirical associations. It was clear from the outset that such an approach would not be desirable or practical in this domain. Instead, the new system uses a detailed model of the structure and function of the device under test, allowing it to reason about how the device works, and how it fails, in a manner similar to an experience engineer.

Building and using such a model represents an important advance in the art of expert systems construction, since it provides the system with a more fundamental understanding of the device than is possible using the traditional approach. While the current work is aimed at computer diagnosis because of its familiarity and the significance of the problem, the larger concern is that of reasoning about devices, understanding how they work and how they fail.

A second component of Professor Davis's expert systems work is a project with Mr. Reid Simmons on reasoning about geologic processes. This was inspired by earlier, traditional expert system work on a program for analyzing measurements taken on oil wells. The current effort is aimed at developing a program that understands and is capable of reasoning about geologic processes responsible for formations and deposits.

Professor Davis also has collaborated with Professor Alan V. Oppenheim in exploring connections between Artificial Intelligence and digital signal processing, in the belief that techniques developed in both areas can be modified and combined to produce improved signal interpretation and signal enhancement systems. This work has recently focused on novel varieties of the traditional problems of pitch detection and signal enhancement.

Dr. Charles Rich and Dr. Richard C. Waters study how expert programmers analyze, synthesize, modify, explain, verify and document programs. One of their key observations is that expert programmers rely heavily on a large body of standard implementation methods and program forms. The heart of their research has therefore been to identify

and codify these standard forms. For this purpose, they have developed a language-independent formalism, called the plan calculus, in which these standard forms can be written down in a canonical and abstract way and used by an automatic programming system. Dr. Rich has compiled an initial library of several hundred plans using this formalism.

Recognizing that the long-term goal of totally automatic programming is very far off, Drs. Rich and Waters have directed their research toward the immediate goal of developing an intelligent computer assistant for programmers, called the Programmer's Apprentice. As a first demonstration of how the Programmer's Apprentice can help a programmer, Dr. Waters has implemented a new kind of program editor which understands how a program is built out of standard plans. With existing program editors, a single logical change to a program must often be achieved by many separate changes to different parts of the text. With the new editor, many such logical changes can be achieved by a single command. As their underlying theory develops further, Drs. Rich and Waters expect the capabilities of the Programmer's Apprentice to increase, with a corresponding reduction in the amount of work expert programmers are required to do.

Mr. Kenneth D. Forbus's work is on qualitative simulation of physical systems. The central idea of his theory is that an explicit representation of physical processes is central to common-sense physical reasoning. Processes such as flows and state changes—heating, boiling, moving, and breaking—form a major part of the physical phenomena that must be dealt with, so reasoning about when they occur, the changes they cause, and the ways they end is important for understanding physical reasoning. Mr. Forbus has developed a representation for reasoning about quantities called the quantity space. This representation is particularly good for dealing with the inequalities that mark state changes.

Mr. David A. McAllester has produced a reasoning utility package that is now in use in Mr. Forbus's project and many others. The package has programs for doing simple deductions, recording justifications, tracking down assumptions, and performing modifications as premises are changed.

## Language and Learning

Professor Robert C. Berwick has extended his model of language acquisition so that almost the entire core of English syntax can be learned by a computer program that examines simple English sentences. The extended model makes it possible to test the acquisition program on languages other than English. As a first step towards this goal, two experimental parsers were designed and implemented for languages other than English. One is based on earlier work done at the Laboratory, combined with recently proposed linguistic theories of German syntax. The aim was confirm that the earlier work, based on English, could be extended to German. The other parser analyzes simple Japanese sentences.

Professor Patrick H. Winston's theory of reasoning by

analogy consists of the following parts: an English understanding module, by Mr. Boris Katz, that converts prepared text into relations in a semantic network, a relation-exploiting hypothesizer that searches memory for precedents likely to be similar to a new given situation; a cause-dominated matcher that finds the best possible correspondences according to the causal framework determined by the situations themselves; an analogizing module that reaches conclusions about a given situation by using a remembered precedent; and a rule building that constructs if-then rules.

During the past year, Professor Winston has improved the reliability of the constructed if-then rules. The key idea is to block if-then rules when contraindicated by information already in hand. The blocking mechanism involves censors of the sort suggested by Professor Marvin Minsky.

Censors are an important component of Minsky's society of minds theory, in which intelligence emerges from the interaction of large societies of rather simple individual agents. Because each agent is relatively simple, communication between agents is restricted in amount and in complexity.

Professor Minsky's approach may illuminate the psychological theories of Piaget and Freud, as well as give coherence to artificial intelligence theories that do not consider the kinds of problems inherent in the representation of a personality. The limitations of inter-agent communication make it necessary for the mind to develop hierarchies of control structures that may be similar to developmental stages. The censors and critics of the hierarchy settle conflicts by referring to early-developed self-images.

## Computing Concepts and Distributed Problem Solving

Our connection memory machine is a computer that performs fast, parallel searches through networks, a fundamental, rate-limiting step in many problem-solving programs. Led by Mr. William D. Hillis, work on this machine is a major effort.

During the past year a prototype machine has been designed, and a packaging scheme has been selected. Thus we are almost ready to begin the implementation of prototype hardware. Because the design requires parts to be duplicated in great numbers, yield, testing, packaging and reliability are receiving careful attention. The large scale of the project forces the development of new methodologies and design tools.

Mr. Alan Bawden has written a simulator which simulates the hardwired instruction set of the connection memory machine. This simulator will be used to write and debug programs for use in testing the physical machine. Using the simulator, Mr. Bawden is now writing a truth-maintenance system.

For more advanced applications, such as working with semantic network data bases, we have designed and implemented a prototype high-level language called CMSL. To the

CMSL programmer, all data appears a graph of named vertices connected by edges. In reality, the graph is represented as structures of connection machine cells. Some CMSL operations correspond almost daily to a single connection machine instruction. Others involve complex patterns of messages passed between cells.

Another approach to creating parallel systems is that of Professor Carl E. Hewitt and his associates. They have developed actor theory, a rigorous abstract theory of parallel systems, that may provide a foundation for the construction and analysis of highly parallel problem-solving systems. Several important components have been developed including the Act 1 subsystem to model specialized communicating agents, the Omega subsystem for parallel semantic taxonomic relational networks, implemented by Dr. Gerald R. Barber and Mr. Eugene Ciccarelli, and the Sprites subsystem to model the communications activities involved in processing general goals and assertions, implemented by Dr. William A. Kornfeld.

Working with Professor Hewitt, Mr. Henry Lieberman has completed a preliminary version of the Tinker System for concrete programming. Tinker enables a user at a work station to develop general procedures by abstracting from specific instances of steps performed on concrete examples at the work station. Mr. Lieberman and Professor Hewitt also have developed a real-time garbage collection algorithm based on the life time of objects.

Finally, Professor Sussman and his associates have completed the design of the SCHEME-81 chip, a VLSI device for running SCHEME, a dialect of the LISP programming language. In addition to testing the limits of our automated design aides, the SCHEME-81 chip may be a step toward better LISP-oriented personal computers: preliminary estimates are that SCHEME-81 will interpret simple LISP programs about five times faster than our best current hardware. For large, complex programs, it will do even better.

A special-purpose silicon compiler, by Mr. Philip E. Agre, has been important to the development of SCHEME-81. Given a small program definition, Mr. Agre's compiler produces code specifying the layout of a SCHEME-81 bus-compatible chip to implement that function. The compiler uses traditional techniques and some novel heuristic methods to reason about the tradeoffs involved in writing highly parallel microcode.

Also, Mr. Jonathan D. Taft has been working on building a small SCHEME computer for testing our chips. It uses a Motorola 68000 design module as a front-end processor for performing I/O, for user-level arithmetic, for character manipulation, and for console control and debugging of the SCHEME system.

### The Computing Environment

The Laboratory's computing resources were improved by the installation of a large 20/60 system, a VAX 11/780, and a VAX 11/750, all manufactured by the Digital Equipment

Corporation. The machines complement a variety of existing machines, including nearly two dozen LISP Machines, designed and built by the MIT Artificial Intelligence Laboratory.


All of the machines, together with terminal concentrators, are linked together with an eight-megabit packet-oriented cable system known as the CHAOSNET. The cable system can support as many as 100 communicating computers before reaching intolerable performance deterioration.

(Continued from page 37)

Tatman, J. L. & Smith, J. W., Jr. (1982) Preliminary evaluation of MDX, a medical diagnosis system *Proc 6th Annu. Symp Comput. Appl. Med. Care*, IEEE Computer Society.

Turing, A.M. 1963 *Computing Machinery and Intelligence*. In E. A. Feigenbaum and J. Feldman (Eds.) *Computers and Thought* New York: McGraw-Hill

Yu, V.L., Buchanan, B.G., Shortliffe, E.H., Wraith, S.M., Davis, R., Scott, A.C., & Cohen, S.N. 1979 Evaluating the Performance of a Computer-based Consultant. *Computer Programs in Biomedicine*, 9:95-102.



**Applied Expert Systems, Inc.**

*Providing knowledge-based expert systems to the financial services community*

**Dr. Fred L. Luconi**  
*President*

**Dr. Norton R. Greenfeld**  
*Director — AI Technology*

Five Cambridge Center  
Cambridge, MA 02142  
(617) 492-7322