AI Magazine Volume 3 Number 1 (1982) (© AAAI)

RESEARCH IN PROGRESS

AI and Brain-Theory Research at Computer and Information Science Department University of Massachusetts

Amherst, Massachusetts 01003

Edited by Victor Lesser

Our program in AI is part of the larger departmental focal area of Cybernetics which integrates both AI and Brain Theory (BT). Our research also draws upon a new and expanding interdepartmental program in Cognitive Science that brings together researchers in Cybernetics, Linguistics, Philosophy, and Psychology. This interdisciplinary approach to AI has already led to a number of fruitful collaborations in the areas of cooperative computation, learning, natural language parsing, and vision.

There are nine faculty members, three research computer scientists and fifty graduate students involved in AI and BT research at UMASS. The major research areas include Adaptive Control, Automatic Programming, Distributed Problem-Solving, Expert Systems (Knowledge Representation and Control), Intelligent User Interfaces, Learning, Natural Language Generation and Parsing, Neurophysiology of Vision, Robotics, and Vision. This research has substantial external funding from a variety of sources that include DARPA, DEC, NIE, NIH, NSF, ONR, and the SLOAN Foundation.

Our departmental research computer facilities consist of a network at two VAX 11/780s which have connected to them five high-resolution color graphics devices and an image processing workstation. Most of our AI programming has been done in our own dialect of LISP which is fully integrated into the VMS operating system for the VAX, including access to the command interpreter. This dialect of LISP, which has been operational for three years, has emphasized clarity and uniformity of all conventions and supports directly a full range of data types including arrays. Surrounding this system are a number of support subsystems that include an extensive graphics library and a graph processing language.

The following is a list of research projects in the department. We have organized these projects into three general areas: Expert Systems, Vision and Robotics, and Distributed Problem-Solving.

I. EXPERT SYSTEMS

Intelligent User Interfaces

Bruce Croft and *Victor Lesser*

This project, which has just been initiated, involves the development of a methodology that applies both interpretation and planning to the process of assisting a professional in carrying out computer-based tasks. The approach is to provide an intelligent interface which monitors the professional's command sequences in order to recognize the actions undertaken, correct or warn about possibly incorrect actions based on system predictions, provide information about the current state of the user's actions (over a number of terminal sessions and procedures), and perform task initiation and/or completion automatically. The system should also be able to recognize more difficult situations and go into a mode where it simply provides tools to help the user solve the problem.

The knowledge used by the system to accomplish these monitoring tasks is specified in terms of plans (schemas) that are abstracted descriptions of the sequence of events that need to occur in order for common user tasks to be accomplished. The formalism used for these plans is based on the EDL language [Bates81]. The application of this approach in the software development environment and in an office environment is being investigated.

Natural Language Generation and Parsing

David McDonald and *Robert Moll*

McDonald's work has two foci: basic research on the process of natural language generation, and the rapid transfer and adaptation of major language processing systems developed at other sites. The latter activity will assist other projects in developing facilities for explanation of expert reasoning and communication via natural language.

Previous work in natural language generation resulted in the development of a computer program ("MUMBLE") that produced text for five experimental "micro-speakers" [MCD080]. The program is based on a theory of generation in which the speakers' goals (messages) are interpreted as a program of decisions to be made. These decisions are ordered according to their relative importance and dependencies. An explicit representation of the results of earlier decisions is used to constrain later ones. Goals may be described to the program using any representation the designer chooses provided they can be enumerated as a hierarchy of relations in a fixed vocabulary. Linguistic rigor is strongly emphasized by employing theoretically motivated grammars and by paying close attention to the available psycholinguistic data.

Current work breaks new ground in the areas of planning rhetorical goals and of interactions between conceptual and linguistic knowledge and conventions during processing. The problem domain for this new work is the description of natural scenes that are recognized by a vision understanding system [CONK79].

Moll is working on formal models of language. Research in progress is aimed at extending and clarifying the work by Wexler and his associates on the formal constraints on learnability for more modern theories of grammar. Another project has produced a parser for English compatible with psycholinguistic data on human parsing strategies.

Automatic Programming

Robert Moll

An automatic programming system has been developed which produces LISP code as output when supplied with nonconstructive specifications as input. The system operates by applying a sequence of transformations to specifications. The transformations rewrite the specifications and eventually eliminate al nonconstructive elements. The result is a legal LISP program. Great care has been taken to design domain independent transformations. This has led to the development of a general planning language for program development.

Epistemology and Learning

Edwina Rissland

Understanding the structure and epistemology of complex knowledge domains is the key to the development of tutoring systems and the next generation of expert systems. Past work has included research on the understanding of mathematics [RISS78], the design of intelligent CAI systems (with Elliot Soloway of Yale University) and the tailoring of a conceptual framework for high-school level mathematics.

Current work involves investigating the structure, use and generation of examples in law, linguistics, mathematics, and computer science. As part of this work, a system to generate examples meeting specified constraints has been developed [RISS80a,80b] The CEG (Constrained Example Generator) system is being used not only to study the example generation process, but also to study issues in learning [RISS81a,81b]; examples are critical to any system -- man or machine -- that learns and a system that generates examples ought to be able to learn how to improve its capabilities. There is also investigation into the role of scenarios in tactical planning using ideas from the CEG research.

Rissland, together with McDonald, hopes to integrate a CEG-like system into a software environment and use it to illustrate explanations that the environment provides its users. They plan to build a facility that will be able to generate natural language explanations that incorporate illustrative examples. Related to the use of examples in generating explanations is their use in probing and testing systems [RISS82]. Together with Oliver Selfridge, Rissland is looking at issues in learning and adaptation as they relate to ill-defined systems [SELF82]. CEG-related work may also be applied to research in vision where new scenes can often be generated by modifying already existing scenes, and to medical and legal knowledge bases where the knowledge structure is complex and in which examples play a central role.

II. VISION AND ROBOTICS

Visuo-Tactile Coordination for Robotic Control

Michael Arbib and *Edward Riseman*

We are in the process of establishing a research facility for the study of robot sensing and control, specifically in the area of the integrated visual and tactile guidance of robot arms and hands. As part of our initial work in this area, we have recently completed construction of a tactile sensor with the cooperation of an industrial automation group and have begun to process force data derived from this sensor [OVER81].

The areas to be studied in this research include static and dynamic vision, static and dynamic tactile sensation, the interactions of multimodal sensory data in the creation and maintenance of environmental models (schema-assemblages), and adaptive control structures using a knowledge-directed approach. We expect to draw on our expertise in a range of areas of machine vision, in the areas of neural modelling and physiology, and from studies on animal behavior.

Vision

Allen Hanson and *Edward Riseman*

This project is investigating the construction of integrated computer vision systems. The goal is to provide an analysis of color images of outdoor scenes, from segmentation (or partitioning) of an image through the final stages of symbolic interpretation of that image. The output of the system is intended to be a symbolic representation of the three-dimensional world depicted in the two-dimensional image, including the naming of objects, their placement in three-dimensional space, and the ability to predict from this representation the rough appearance of the scene from other points of view.

This project is attempting to deal in a general experimental manner with the complexities inherent in the interpretation of natural scenes. The VISIONS system [HANS78] represents an empirical approach to the design and implementation of an extremely complex system in which some of the required subprocesses are not fully understood. In a long-range systematic development plan, we have implemented a system of processes and control mechanisms to be used in the interpretation of a variety of types of images. The current emphasis is on knowledge-directed interpretation via structures called schemas [PARM80]. This research focus also includes the development of "low-level" processes that can select the appropriate high-level schema and those that can be biased in their activities as a result of high-level expectations derived from the chosen schema.

Recent applications of VISIONS include work on quantification of diagnostic and disease history variables relevant to the treatment of glaucoma, on the prognosis of melanoma, and on processing of satellite images. We have also begun work on industrial applications of robotics and motion and have an ongoing project in the general area of industrial robotics.

Processing Dynamic Images From Camera Motion

Allen Hanson and *Edward Riseman*

There has been a significant and long-term commitment to motion processing in our department for the past five years. Past work has involved: developing neural networks embodying cooperative algorithms for determining segmentations from stereo and time-varying images [DEV75, BURT75]; developing algorithms for determining optic flow image sequences [PRAG80, LAWT80]; and developing algorithms for effectively extracting environmental and control information from image sequences produced from moving vehicles and robot assembly environments [WILL80, LAWT80].

Our current work addresses the basic issues that must be understood in order to develop computer vision systems for terrestrial and airborne motion. Using a sequence of images obtained from a sensor in motion, we intend to demonstrate the feasibility of determining the changes in the sequence of images and establishing a consistent environmental model over time.

The key scientific issue to be addressed is the recovery and effective representation of information concerning the physical environment, such as surface distance, extent, and orientation, relative to a moving observer. The necessary techniques are being developed using simulated and actual scenes with restricted forms of sensor motion, leading toward analysis of actual scenes with smooth, but arbitrary, motion.

Neurophysiology of Vision

D. Nico Spinelli and *Michael Arbib*

There are a number of projects which address the study of pattern recognition, perception, learning, and visual control of movement both in animals and machines. There are three facets to this overall effort: neurobiological investigations into the development of the mammalian visual system, the modelling of adaptive networks, and neural net simulations of visuomotor coordination and control of locomotion.

The philosophical underpinning of these projects is that truly major advances in our understanding of pattern recognition, learning, the action-perception cycle, and intelligence will come only through a close interaction between experimentation on natural brains and "theory" as realized in BT and computational strategies developed in work on AI. It is a basic assumption that the design and analysis of machines, be they natural or artificial "brains" capable of sophisticated behavior, can benefit from the development of a set of fundamental principles.

An example of this work is that of Spinelli's, which is aimed at discovering how higher levels in the mammalian visual processing system control the way in which the lower level processes information. He has shown that functional properties of single cells in the visual cortex are a direct result of visual experience during development [SPIN72], and also that experience shapes the functional properties not only of visual elements but also of sensory-motor ones [SPIN79]. These functional effects are the result of actual structural changes of how neural elements develop [SPIN80].

Another example is Arbib's work which uses the frog brain as a focus for analyzing neural mechanisms in visuomotor coordination at a level of detail which allows design and simulation of detailed experiments in neurophysiology, neuroanatomy, and animal behavior. This work has provided a seedbed for our current work in robotics as an outgrowth of a concern with sensorimotor coordination, and in processing dynamic images as an outgrowth of studies of human and animal motion perception.

III. DISTRIBUTED PROBLEM-SOLVING

Cooperative Computation in Vision, Brain, and Language

Michael Arbib

This project emphasizes a form of distributed problem-solving, called *cooperative computation*, in which the different processes are embodied as dedicated resources and cooperate to achieve a common goal.

The three subprojects in this effort are:

1. Cooperative computation for dynamic machine vision, tactile processing, and sensory-motor coordination in robot control [ARBI78]. (This project is conducted in collaboration with the Hanson-Riseman vision project.)
2. Cooperative computation in neural circuitry subserving visuomotor coordination and control of movement; models of neural development and plasticity [ARBI81a,81b]. (Many of the simulation problems here are common with those of Barto and Spinelli's research.)
3. Cooperative computation in linguistics [ARBI79]. One subproject (in collaboration with McDonald) studies the interaction of semantic representations of visual scenes with linguistic representations in the production of scene descriptions. A second models cognitive-syntactic interactions in language acquisition. A third studies the interaction of phonetic, syntactic, and pragmatic representations, using a simulation so structured that the effects of subsystem removal can be tested to analyze a variety of hypotheses about the way in which brain damage causes the various aphasias.

One emphasis of this research is on low-level cooperative computations, involving high parallelism of many copies of a few types of processes. The coordination technique among processes is based on the neural control mechanisms of excitation, inhibition, thresholding, and decay. There is also beginning work on high-level cooperative computation structures as a model for interaction among brain regions.

Goal-Seeking Systems of Goal-Seeking Components

Andy Barto and *D. Nico Spinelli*

There are many problems about which insufficient *a priori* knowledge exists to permit the design of adequate problem-solving methods using either traditional mathematical techniques or knowledge-based AI approaches. Adaptive or learning techniques are required for the solution of these problems. This is a study of the feasibility of a distributed or network approach to adaptive problem-solving in which each network component possesses considerably more adaptive power than the components of the adaptive networks typically studied in the past [SPIN79,80] and acts according to its own self-interest. Each component is goal-seeking in the sense that it has a preference ordering over its possible input patterns and employs strategies aimed at causing its input to be as highly preferable as possible.

The thesis in this project is that if the components are sufficiently robust learning systems, they may be able to solve problems collectively that they cannot solve by themselves. The long-term goal is to explore the possibility of designing deeply and pervasively adaptive computational substrates for problem-solving. Issues that arise concern pattern recognition, search, conflicts of interest, cooperativity, and the adaptive construction of local performance evaluation criteria.

One result of this study has been the development of a new kind of associative memory network. Associative memory networks have been proposed as models of biological memory and as potentially useful storage devices. By storing information in distributed form, they show great resistance to noise and damage. When constructed from goal-seeking components, this type of memory retains noise and damage resistance but has the additional ability to use a generate-and-test procedure to search for optimal associations. It does not need a "teacher" to explicitly tell it which associations are desired. This more difficult form of learning permits the use of associative memory networks for problems made difficult by lack of *a priori* knowledge [BART81]. We have demonstrated these capabilities using computer simulations of route finding in spatial environments consisting of detectable landmarks [SUTT81].

Cooperative Distributed Problem-Solving and Organizational Self-Design

Victor Lesser and *Daniel Corkill*

This work is part of a new research area that has recently emerged in AI, called *Distributed AI.* It has relevance to the organization of both physically and logically distributed AI systems. The major focus of current research efforts has been the exploration of a new model for distributed problem-solving systems [LESS80, 81a]. This model differs from conventional approaches in its emphasis on dealing with distribution-caused uncertainty and errors in control, data, and algorithms as an integral part of the problem-solving process.

In experiments with such systems we have come to realize that one of the key issues in effectively exploiting such an approach is how to maintain global coherence among the autonomous, self-directed problem-solvers. Our research advocates organizational self-design as a framework for obtaining this coherence. Organizational self-design is the explicit planning by the system of the pattern of information and control relationships that exist among the nodes in the system and the distribution of problem-solving capabilities (expertise) among the nodes in the system. As the problem-solving environment changes, the distributed system may need to change its organizational structure to maintain its effectiveness [CORK80].

We have developed a distributed interpretation testbed [LESS81b] that simulates a model of a distributed, knowledge-based, problem-solving architecture applied to an abstracted version of a vehicle monitoring task. Using the testbed, we are empirically exploring a wide range of control/communication policies for organizing distributed problem-solving systems. We are also developing a distributed planning system for the testbed that integrates both the coordination of "domain-level" activities and the construction and maintenance of an appropriate organizational structure.

REFERENCES

[ARBI78] Arbib, M A , "Segmentation, Schemas, and Cooperative Computation," in *Studies in mathematical Biology*, Part I (S A Levin, Ed) MAA Studies in Mathematics *15*, (1978), pp 118-115

[ARBI79] Arbib, M A and Caplan, D , "Neurolinguistics Must Be Computational," *Behavioral and Brain Sciences*, 2, (1979), pp 449-483

[ARBI81a] Arbib, M A , "Visuomotor Coordination: From Neural Nets to Schema Theory," *Cognition and Brain Sciences*, 2, (1979), pp 449-483

[ARBI81b] Arbib, M A , "*Rana Computatrix* An Evolving Model of Visuomotor Coordination in Frog and Toad," *Machine Intelligence*, 10, (D Michie and J Hayes, Eds), Ellis Horwood, (1981), (in press)

[BART81] Barto, A G , Sutton, R S ; Brouwer, P S , "Associative Search Network A Reinforcement Learning Associative Memory," *Biological Cybernetics*, 40, (1981), pp 201-211

[BATE81] Bates, P C , Wileden, J C , and Lesser, V R , "EDL: A Basis for Distributed System Debugging Tools," Computer and Information Science Department, University of Massachusetts, Technical Report 81-7, (1981)

[BURT75] Burt, P J , "Computer Simulations of a Dynamic Visual Perception Model," *International Joint Man-Machine Studies*, 7, (1975), pp 529-546

[CONK79] Conklin, J C , "A Scene Description System Preliminary Survey," Master's Project, Department of Computer and Information Science, University of Massachusetts, Amherst, (1979)

[CORK80] Corkill, D D , "An Organizational Approach to Planning in Distributed Problem-Solving Systems," Department of Computer and Information Science, University of Massachusetts, Technical Report 80-13, (May 1980)

[DEV75] Dev, P , "Perception of Depth Surfaces in Random-Dot Stereograms A Neural Model," *International Joint Man-Machine Studies*, 7, (1975), pp 511-528

[HANS78] Hanson, A R and Riseman, E M , "VISIONS: A Computer System for Interpreting Scenes," *Computer Vision Systems*, (A R Hanson and E M Riseman, Eds) Academic Press, (1978)

[LAWT80] Lawton, D T , "Constraint-Based Inference From Image Motion," *Proceedings of the AAAI Conference*, Stanford University, (1980), pp 31-34

[LESS80] Lesser, V R , and Erman, L D , "Distributed Interpretation A Model and Experiment," *IEEE Transactions on Computers*, Vol C-29, No 12, (December 1980), pp 1144-1163

[LESS81a] Lesser, V R and Corkill, D D , "Functionally Accurate, Cooperative Distributed Systems," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol SMC-11, No 1, (January 1981), pp 81-96

[LESS81b] Lesser, V R , et al , "A High-Level Simulation Testbed for Cooperative Distributive Distributed Problem-Solving," Technical Report, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts, (June 1981)

[MCDO80] McDonald, D D , "Natural Language Production as a Process of Decision-Making Under Constraints," Ph D Dissertation, MIT

[OVER81] Overton, K J and Williams, T , "Tactile Sensation for Robots," *Proceedings of IJCAI-81*, Vancouver, British Columbia, (August 1981), pp 791-795

[PARM80] Parma, C C , Hanson, A R , and Riseman, E M , "Experiments in Schema-Driven Interpretation of a Natural Scene," *NATO Advanced Study Institute on Digital Image Processing*, (R Haralick and J C Simon, Eds), Bonas, France; (1980)

[PRAG80] Prager, J M , "Segmentation of Static and Dynamic Scenes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol PAMI-2, (January 1980), pp 16-27

[RISS78] Rissland (Michener), E L , "Understanding Mathematics," *Cognitive Science*, Vol 2, No 4, (1978)

[RISS80a] Rissland, E L and Soloway, E M , "Overview of an Example Generation System," *Proceedings First National Conference on Artificial Intelligence*, Stanford, (August 1980)

[RISS80b] Rissland, E L , and Soloway, E M , "Generating Examples in LISP Data and Programs," *Proceedings International Workshop on Program Construction*, France, (September 1980)

[RISS81a] Rissland, E L , Selfridge, O G , and Soloway, E M , "The Role of Experiences and Examples in Learning Systems," *Proceedings Third Annual Meeting of Cognitive Science Society*, Berkeley, California, (August 1981)

[RISS81b] Rissland, E L , and Soloway, E M , "Constrained Example Generation: A Testbed for Studying Issues in Learning," *Proceedings IJCAI-81*, Vancouver, B C , (1981)

[RISS82] Rissland, E L , "Examples, Adaptation and Learning," to appear in [SELF82]

[SELF82] Selfridge, O G , Rissland, E L , and Arbib, M A , (Eds), "Adaptive Control and Ill-Defined Systems," Plenum, (in preparation)

[SPIN72] Spinelli, D N , Hirsch, H V B , Phelps, R W and Metzler, J , "Visual Experience as a Determinant of the Response Characteristics of Cortical Receptive Fields in Cats," *Exp Brain Res*, 15, (1s, i) 1972, pp 289-304

[SPIN79] Spinelli, D N and Jensen, F E , "Plasticity The Mirror of Experience," *Science*, 203, (1979), pp 75-78

[SPIN80] Spinelli, N D , Jensen, F E , and Viana Di Prisco, G , "Early Experience Effect on Dendritic Branching in Normally Reared Kittens," *Exper Neurol*, 68, (1980), pp 1-11

[SUTT81] Sutton, R S and Barto, A G , "Toward a Modern Theory of Adaptive Networks Expectation and Prediction," *Psych Rev*, 88, (1981), pp 135-170

[WILL80] Williams, T D , "Depth From Camera Motion in a Real World Scene," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-2, (1980), pp 511-516 ■