MACHINE LEARNING AND INFERENCE LABORATORY: AN OVERVIEW OF RESEARCH AND ACTIVITIES: 1994-1995

Edited by

R. S. Michalski
J. Wnek

Reports of the Machine Learning and Inference Laboratory, MLI 96-1, George Mason University, Fairfax, VA, January 1996.

Reports

Machine Learning and Inference Laboratory

AN OVERVIEW OF RESEARCH AND ACTIVITIES
1994 - 1995

P96-1

School of Information Technology and Engineering

George Mason University

Editors:

Ryszard. S. Michalski
Janusz Wnek

For more information about the Laboratory see our web pages at:
http://www.mli.gmu.edu

To inquire about the Laboratory's activities,
the GMU Ph.D. program, or research assistantships
in Machine Learning and Inference/Intelligent Systems contact:

R. S. Michalski
Machine Learning and Inference Laboratory
School of Information Technology and Engineering
George Mason University
4400 University Drive
Fairfax, VA 22030-4444

Phone: (703) 993-1719
Fax: (703) 993-3729
E-mail: admin@aic.gmu.edu

MISSION

One of the central directions of information technology is the development of intelligent systems, that is, systems able to conduct inference, learn and communicate with the external world using a wider range of forms than conventional computer systems.

The mission of the Machine Learning and Inference Laboratory (MLI) is to support education, scholarship and research in the areas of machine learning and inference, as well as their applications to the development of intelligent systems. Special emphasis is placed on the topics of computational models of learning and inference, knowledge discovery, knowledge acquisition and representation, decision support and advisory systems, vision systems with learning capabilities, and the development of experimental intelligent systems for a wide range of engineering and other applications. The MLI Laboratory started in 1988 when Professor R. S. Michalski moved with his research group from the University of Illinois at Urbana-Champaign to George Mason University.

The ideas, methods and systems developed in the Laboratory are used to teach courses in this area, to generate state-of-the-art projects for student research and theses, and to provide students with hands-on experience and training through research. The implemented experimental systems are applied, in cooperation with industry, to a wide spectrum of real-world problems. The Laboratory maintains a close and mutually supportive partnership with various SITE and GMU academic departments and centers. It also interacts and cooperates with national and international organizations working in similar areas.

To support the above mission, specific topics of research concentration include:

- theory and methodology of learning and inference, with special emphasis on multistrategy approaches
- knowledge discovery in databases
- evolutionary computation and genetic algorithms
- applications of learning, inference and discovery systems to areas of special practical significance, such as engineering design, education, fact finding in law, software engineering, development of intelligent software agents and intelligent information systems, information superhighways, networking, C^3I, multimedia, economics, business, molecular biology, biochemistry and medicine
- development of intelligent systems with a special emphasis on engineering design and software engineering
- knowledge representation, acquisition, and visualization
- machine learning in computer vision and intelligent robots
- advisory systems for education, in particular undergraduate education and professional training
- models of plausible and probabilistic reasoning and discovery
- advisory systems and decision support systems with learning capabilities.

The above topics represent a continuation of past research and its extension in new directions, in particular, toward development and applications of intelligent systems to engineering, education, software systems, and communication and information systems.

The stated mission indicates the Laboratory's aspiration to serve three intertwined and mutually dependent functions—to support excellence in education, to conduct basic research, and to apply developed ideas to real-world problems, and by that to facilitate the transfer of university research to industry. Such aims thus fully support the University's goals of excellence in teaching, scholarship and research.

Machine Learning and Inference Laboratory

The Laboratory investigates research issues concerned with the understanding of learning and inference processes, building experimental machine learning and inference programs, and incorporating them into intelligent systems. The Laboratory also studies practical applications of the developed methods and systems, and collaborates with companies and government organizations on practical problems that require machine learning and inference capabilities.

Research projects conducted in the Laboratory cover a wide spectrum of topics, including the inferential theory of learning, multistrategy task-adaptive learning, knowledge discovery in databases, knowledge acquisition, constructive induction, evolutionary computation and genetic algorithms, cognitive models of plausible reasoning, knowledge representation, expert systems with learning capabilities, expert database systems, and applications of machine learning to computer vision. Using grant support and some initial support from the University, the Laboratory has acquired and organized excellent computational facilities. These facilities are built primarily around the SUN SPARC and Apple Macintosh workstations (see *Computing Facilities* on page 41).

The Laboratory has established its own technical publication, *Reports of the Machine Learning and Inference Laboratory*. This publication serves as a mechanism for prompt publishing of the results of the Laboratory's research. It continues the traditions of *Reports of the Intelligent Systems Group*, which Dr. Michalski's group published at the University of Illinois at Urbana-Champaign. The MLI Reports are disseminated to many research centers, academic institutions and industrial organizations around the world. A list of recent MLI publications is at the end of this report.

This report primarily reviews the most recent MLI activities; however, to provide a proper context for them, it also mentions some of the past accomplishments, projects, and collaborators.

Our Web site, http://www.mli.gmu.edu, contains a list of all publications with abstracts for 1995 and a complete list of our past publications. This site also contains the updated information about our projects and researchers.

Closely Related Research

Other research in the Laboratory is conducted by faculty members and collaborators who, though not members of the MLI group, have been interacting or collaborating with MLI members on various joint projects and research grants. Research topics include machine learning in software engineering, information systems, discovery and fact finding, computer vision, cognitive modeling, the development of task-oriented intelligent systems, connectionist systems and neural nets, and applications of intelligent systems to world economics, business, education, communication networks, biochemistry, and agriculture.

Cross-Disciplinary Collaboration

The Laboratory strongly emphasizes and encourages a cross-disciplinary collaboration. In this spirit, the collaborators in the Laboratory have not been limited to one academic department. Although most faculty members are from Departments of Computer Science and Systems Engineering, our GMU collaborators also include faculty from the Departments of Operations Research and Engineering, Information and Software Systems Engineering, Psychology, and School of Business Administration. The MLI Laboratory has also affiliates and collaborators from other Universities, both in US and abroad, as well as from industrial and governmental organizations.

Since 1992, the Laboratory has had a joint ARPA-sponsored project with the University of Maryland to apply machine learning to computer vision. In this project, the Laboratory researchers have been collaborating with the UMD Computer Vision Laboratory, which is one of the leading and the oldest research centers in computer vision.

The collaborating researchers from that Laboratory include Professor Azriel Rosenfeld, a world renowned computer vision pioneer and researcher, and Professor Yiannis Aloimonos, a leading computer vision researcher and recipient of the prestigious National Science Foundation's Young Investigators Award. Recently, Dr. Zoran Duric, former student of Professor A. Rosenfeld, joined our Laboratory as a Senior Research Faculty member and serves as a liaison between both Laboratories.

Our researchers, jointly with the researchers from the UMD Computer Vision Laboratory, organized the NSF/ARPA Workshop on Machine Vision and Learning. The workshop brought together for the first time leading US researchers in these two disciplines and produced the special report on computer vision and learning.

Collaboration with Other Organizations

The nature of research in the Laboratory requires that developed methods and systems be experimentally validated by applying them to real-world problems. Collaboration between the Laboratory and industry is therefore a significant part of the Laboratory's activities.

The Laboratory has been collaborating with various government and industrial organizations, in particular with the Naval Research Laboratory and the Chrysler Corporation. Our major industrial collaborators included at various times also The Analytic Science Corporation (TASC), Computer Sciences Corporation (CSC), the MCI Corporation, and the Science Applications International Corporation (SAIC).

International Collaboration

Since its inception, the Laboratory has been involved in collaboration with various international researchers. There is ongoing interaction with Hugo de Garis from the Brain Building Group in Japan, and Heedong Ko from the Korean Institute of Science and Technology.

We have been collaborating with Professor Iven Van Mechelen and P. Theuns from Catholic University of Leuven in Belgium, and Professor James Hampton from the City University of London. This collaboration has recently produced an edited book *Categories and Concepts: Theoretical Views and Inductive Data Analysis*, which is the first book that relates the research on cognitive theories of concepts and methods for data analysis. There have been short and long-term visits and collaboration on various projects with scientists from Belgium, Egypt, Hungary, Germany, Italy, Japan, Poland, Russia, and Slovenia.

The Laboratory, jointly with the Center for Expert Systems at the School of Business Administration, was awarded a grant from the United Nations to collaborate with the Agricultural Expert Systems Project of the Arab Republic of Egypt.

Within the framework of this collaboration, we have trained two researchers from Egypt in various aspects of expert systems, knowledge acquisition and machine learning. Two of our researchers have been invited to Egypt to conduct joint research projects and provide consultation. An Egyptian researcher (Ibrahim Imam) came to spend six months at the Center under the sponsorship of the United Nations, to collaborate in the area of knowledge acquisition for expert systems. Subsequently, he joined the doctoral studies program at George Mason University and became a MLI member. Dr. Michalski spent several weeks with the Expert System Project at the Egyptian Ministry of Agriculture to consult on the development of agricultural expert systems.

Collaboration with Egypt is an expression of our long-held conviction that intelligent systems technology has a great potential for the developing countries.

We also established links and are developing collaboration with the Institute of Computer Science at the Polish Academy of Sciences in Warsaw. These links have been recently strengthened by an appointment of Ryszard Michalski as an Affiliate Professor of this Institute.

Workshop Organization

Our professional activities include organization of several research workshops. In view of our long standing interests in the study and integration of learning strategies, we organized the First International Workshop on Multistrategy Learning (MSL-91) in Harpers Ferry, WV, on November 7-9, 1991. This workshop brought together leading researchers from U.S. and Europe interested in this very challenging and promising area of machine learning. One of the results that followed this workshop was a publication of the first book on multistrategy learning, edited jointly by R. S. Michalski and G. Tecuci (1994).

On May 27-29, 1993, we organized the Second International Workshop on Multistrategy Learning (MSL-93) held in Harpers Ferry, WV. The workshop was attended by representatives of eleven countries and clearly demonstrated a growing interest in this area worldwide.

In collaboration with the Computer Vision Laboratory at the University of Maryland, we organized the NSF/ARPA Workshop on Machine Vision and Learning (MVL) that was held October 15-17, 1992, also in Harpers Ferry, WV. It brought together for the first time leading US researchers in learning and vision. A significant result of this workshop was a special report on computer vision and learning (Michalski, Rosenfeld and Aloimonos, 1994).

We organized the Intelligent Adaptive Systems (IAS-95) workshop, held on April 26, 1995, hosted by the Florida AI Research Symposium (FLAIRS-95). The organizers and editors of the workshop proceedings were our researchers, Ibrahim Imam and Janusz Wnek.

For copies of the proceedings from these workshops contact Janusz Wnek (jwnek@gmu.edu) or Abhay Kasera (akasera@aic.gmu.edu).

The next section describes the next workshop that we are currently organizing, the Third International Workshop on Multistrategy Learning.

3rd International Workshop on Multistrategy Learning (MSL'96)

The rapid expansion of machine learning methods, approaches and paradigms creates a strong need for investigating their interrelationships, and the development of methods for their synergistic integration. Multistrategy learning workshops (MSL) provide a forum for presenting and discussing research on these topics, and related issues, such as the cross-fertilization of machine learning and cognitive science research, learning with large knowledge bases, goal-oriented learning agents, and multistrategy approaches to knowledge discovery.

In this context, we will organize the 3rd MSL workshop to be held in May 23-25, 1996 in Harpers Ferry, WV. The workshop will explore current issues in multistrategy learning research, with emphasis on the following topics:

- comparative studies of learning strategies, methods and paradigms
- cognitive models of learning, inference and discovery
- intelligent learning agents and complex adaptive systems
- user-oriented learning in distributed information systems (e.g., WWW)
- the role of learning goals and learning in large knowledge systems
- knowledge representation, acquisition and reuse in multistrategy learning and inference systems
- advanced applications of multistrategy learning and knowledge discovery

MSL '96 will be held in the picturesque and historic Harpers Ferry, located at the intersection of Virginia, West Virginia and Maryland.

We are very fortunate that outstanding and well-known researchers have agreed to give invited presentations or serve on the Program Committee for this Workshop.

Invited Speakers

John Anderson, Carnegie Mellon University
Jaime Carbonell, Carnegie Mellon University
Hugo de Garis, ATR, Kansai Science City
Laveen Kanal, University of Maryland & LNK Inc.
Doug Lenat, CycCorp
Doug Medin, Northwestern University
Marvin Minsky, MIT
Michael Pazzani, University of California at Irvine
Luc de Raedt, Catholic University of Leuven
Paul Rosenbloom, University of Southern California
Lorenza Saitta, University of Torino
Claude Sammut, University of New South Wales
Derek Sleeman, University of Aberdeen

Program Committee

John Anderson, Carnegie Mellon University
Francesco Bergadano, University of Messina
Jaime Carbonell, Carnegie Mellon University
Marie desJardins, SRI International
Hugo de Garis, ATR, Kansai Science City
Diana Gordon, Naval Research Laboratory
Kenneth Haase, MIT
Heedong Ko, Korea Institute of Technology
Yves Kodratoff, University of Paris South
Stan Matwin, University of Ottawa
Doug Medin, Northwestern University
Raymond Mooney, University of Texas at Austin
Stephen Muggleton, Oxford University
Michael Pazzani, University of California at Irvine
Luc de Raedt, Catholic University of Leuven
Ashwin Ram, Georgia Institute of Technology
Lorenza Saitta, University of Torino
Claude Sammut, University of New South Wales
Jude Shavlik, University of Wisconsin
Derek Sleeman, University of Aberdeen
Gheorghe Tecuci, GMU and Romanian Academy

Sponsors

American Association for Artificial Intelligence
National Science Foundation
Office of Naval Research

Organizers

Ryszard S. Michalski, Chair
Janusz Wnek, Co-Chair

The workshop information is available on the internet, http://www.mli.gmu.edu/msl96.html.

Colloquia and Seminars

The Laboratory conducts regular colloquia which are advertised around GMU and other Universities in the Washington D.C. area, as well as at various governmental organizations and industrial laboratories. The colloquium speakers frequently are well-known researchers from major universities and research organizations in the US and abroad. These presentations help the research community in this area keep abreast of new research and facilitate inter-organizational collaboration. Colloquia held in the 1994-1995 period are listed later in the brochure. In addition, there are regular seminars given by the Laboratory researchers.

Major Research Funding

The Laboratory has been fortunate to win some of the most prestigious research grants. We have significant grants from such organizations as the National Science Foundation (NSF), the Office of Naval Research (ONR), the Advanced Research Projects Agency (ARPA) and the Air Force Office for Scientific Research (AFOSR). Other major sources of funding received include the United Nations and the Naval Research Laboratory.

Research Projects

The Laboratory's research covers a wide spectrum of topics, and is done through projects, each led by one or more senior researchers. Typically, a project involves a theoretical component, the development of a method, its implementation, and a subsequent testing on some designed and/or practical problems. Our recent research has been concerned with such topics as the inferential theory of learning, multistrategy learning, knowledge discovery in data, genetic algorithms and adaptive systems, knowledge acquisition through multistrategy learning, constructive induction, computer vision with learning capabilities, expert databases, diagrammatic visualization of knowledge and inference, knowledge engineering methods, and several others. The following pages provide a brief description of the current projects.

MAJOR RESEARCH PROJECTS

Inferential Theory of Learning

(Michalski, Wnek, Alkharouf, Bloedorn, Kaufman and Utz)

This project aims at the development of the Inferential Theory of Learning (ITL) that views learning as a goal-oriented process of improving the learner's knowledge by exploring the learner's experience. The theory aims at understanding the *competence* aspects of learning processes, in contrast to the Computational Learning Theory that concerns their *computational complexity*. ITL addresses such questions as what types of inference and knowledge transformations underlie learning processes and strategies; what types of knowledge the learner is able to learn from a given input and from a given prior knowledge; what logical relationships exist among the learned knowledge, possible inputs and prior knowledge, etc.

The theory analyzes learning processes in terms of high level inference patterns called *knowledge transmutations*. Among basic transmutations are generalization, abstraction, similization, generation, insertion and replication. The central aspect of any transmutation is the type of underlying inference. If results of inference are found useful, then they are memorized. Thus, we have an "equation":

Learning = Inferencing + Memorizing

Since learning processes may involve any possible type of inference, the ITL postulates that a complete learning theory has to encompass a theory of inference. To this end, we have attempted to identify and classify all major types of inference. The figure in the right column illustrates the proposed classification. The first criterion divides inferences into deductive and inductive. To explain them in a general way, consider the *fundamental equation for inference*: $P \cup BK \models C$, where P stands for premise, BK for reasoner's background knowledge, $\models$ for entailment, and C for consequent. Deductive inference is deriving C, given P and BK, and is truth-preserving.

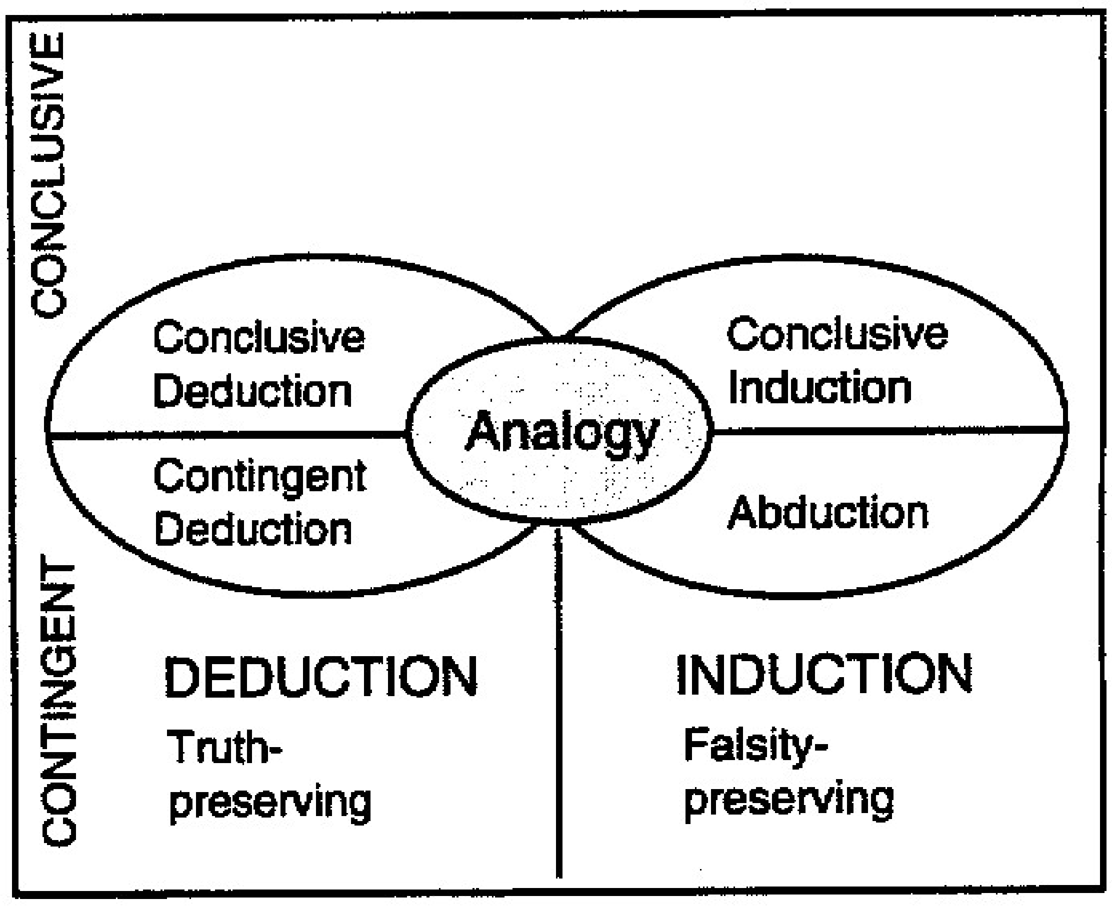

A classification of major types of inference.

Inductive inference is hypothesizing P, given C and BK, and is falsity-preserving. The second classification divides inferences into conclusive (strong) and contingent (weak). Conclusive inferences involve domain-independent inference rules, while contingent inferences involve domain-dependent rules. Contingent deduction produces likely consequences of given causes, and contingent induction produces likely causes of given consequences. Analogy can be characterized as induction and deduction combined, and therefore occupies the central area in the diagram. Using this approach, we have clarified several basic knowledge transmutations, such as inductive and deductive generalization, inductive and deductive specialization, and abstraction and concretion. Generalization and specialization transmutations change the *reference set* of a description, and abstraction and concretion change its level-of-detail.

The support for this research has been provided by the National Science Foundation (NSF), the Office of Naval Research (ONR), and the Advanced Research Projects Agency (ARPA).

Multistrategy Task-Adaptive Learning: MTL

(Michalski, Wnek, Kaufman, Utz, Vafaie, Zhang)

This project is concerned with developing a novel methodology for multistrategy learning, based on the Inferential Theory of Learning. The proposed methodology, called *multistrategy task-adaptive learning* (MTL) integrates a range of learning strategies, in particular, two basic and mutually complementary learning paradigms: *empirical learning and analytical learning* (see Figure beside). Empirical learning assumes that the learner does not have much prior knowledge relevant to the task of learning, while analytic learning assumes that the learner has sufficient knowledge to solve the problem in principle, but that knowledge is not directly applicable or efficient. Empirical learning is based primarily on inductive inference from facts, while analytical learning is based primarily on deductive inference from prior knowledge.

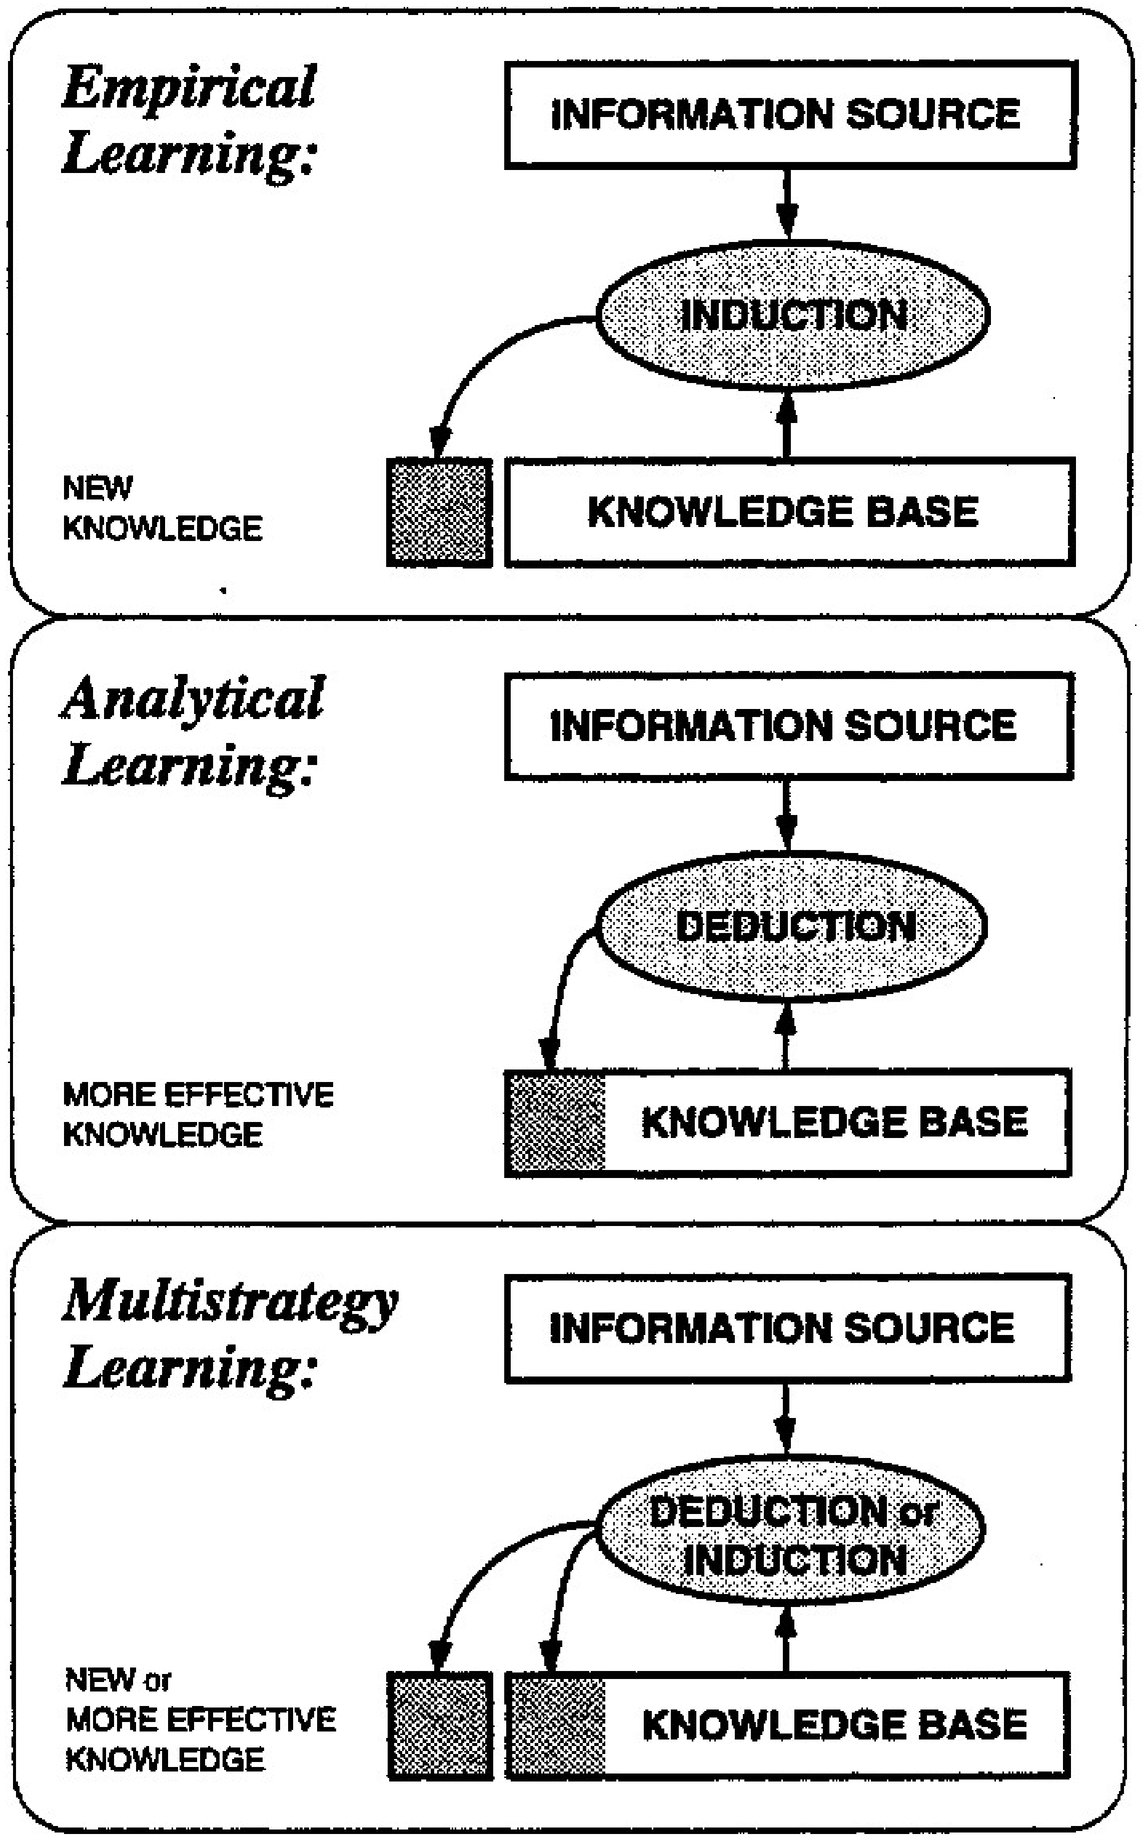

Three paradigms of learning.

Other major learning strategies that are integrated in MTL include *constructive induction*, *analogical learning*, and *abstraction.* Constructive induction employs background knowledge to generate problem-relevant descriptive concepts, and through them derives the most plausible inductive hypotheses. Analogical learning transfers knowledge from one problem domain to another through an analysis of similarities between concepts or problem solving methods. Abstraction transfers a description from a high-detail level to a low-detail and more goal-oriented level.

MTL postulates that the learning strategy, or a combination thereof, should be based on the analysis of the *learning task* at hand. A learning task is defined by the input, learner's prior knowledge and the learner's goal(s). The learning goal(s) are viewed as a central factor in controlling a learning process. This research provides foundations for building advanced learning systems, and applying them to such tasks as knowledge acquisition, planning, problem solving, intelligent robots and knowledge extraction from databases.

The support for this research is provided by the National Science Foundation (NSF), the Office of Naval Research (ONR) and the Advanced Research Projects Agency (ARPA).

Knowledge Discovery in Databases: INLEN

(Michalski, Kerschberg, Wnek, Bloedorn, Imam, Kaufman, Ribeiro, Wozniak)

This project is concerned with the development of a large-scale multi-type reasoning system, called INLEN, for extracting knowledge from databases. The system assists a user in discovering general patterns or trends, meaningful relationships, conceptual or numerical regularities or anomalies in large databases. The volume of information in a database is often too vast for a data analyst to be able to detect such patterns or regularities. INLEN integrates symbolic learning and statistical techniques with database and knowledge base technologies. It provides a user with "knowledge generation operators" (KGOs) for discovering rules characterizing sets of data, generating meaningful conceptual classifications, detecting similarities and formulating explanations for the rules, generating rules and equations characterizing data, selecting and/or generating new relevant variables or representative examples, and testing the discovered rules on new data.

The support for this research is provided by NSF, ONR and ARPA.

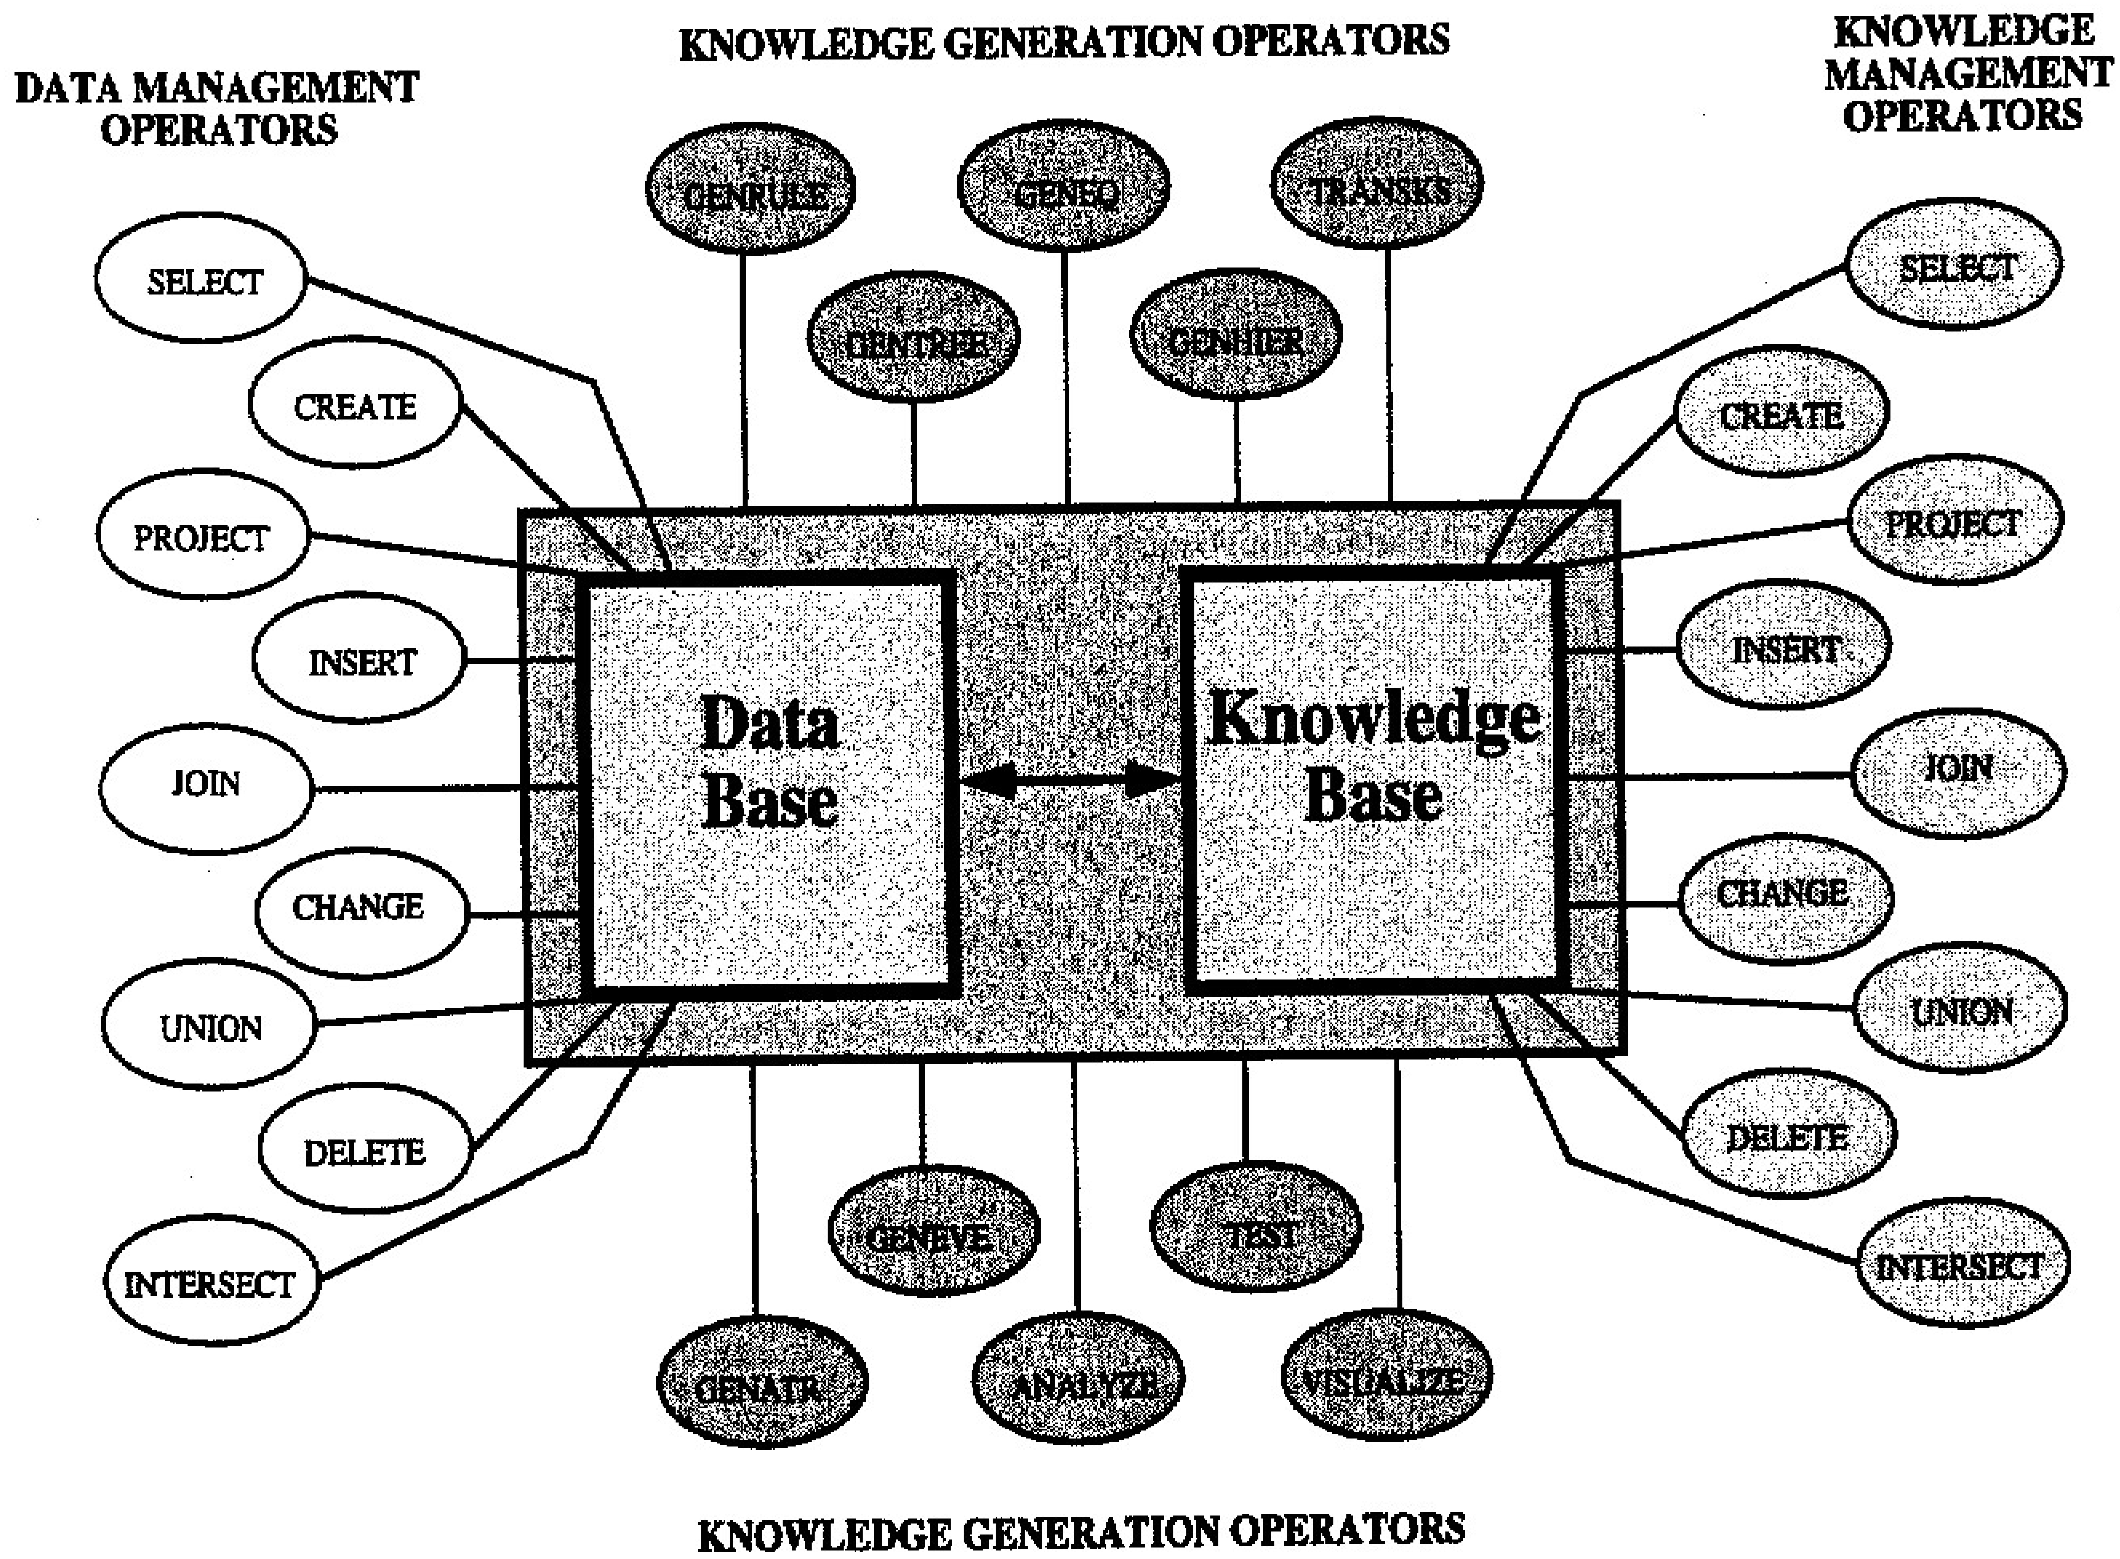

Multistrategy Constructive Induction: Program AQ17-MCI

(Michalski, Wnek, Bloedorn)

Conventional concept learning techniques generate hypotheses in the same representation space in which original training examples are presented. In many learning problems, however, the original representation space is inadequate for formulating the correct hypothesis. This inadequacy can be evidenced by a high degree of irregularity in the distribution of instances of the same class in the original representation space.

We have been developing a methodology and a system, AQ17-MCI, for interpreting a range of strategies for an automated improvement of the knowledge representation spaces.

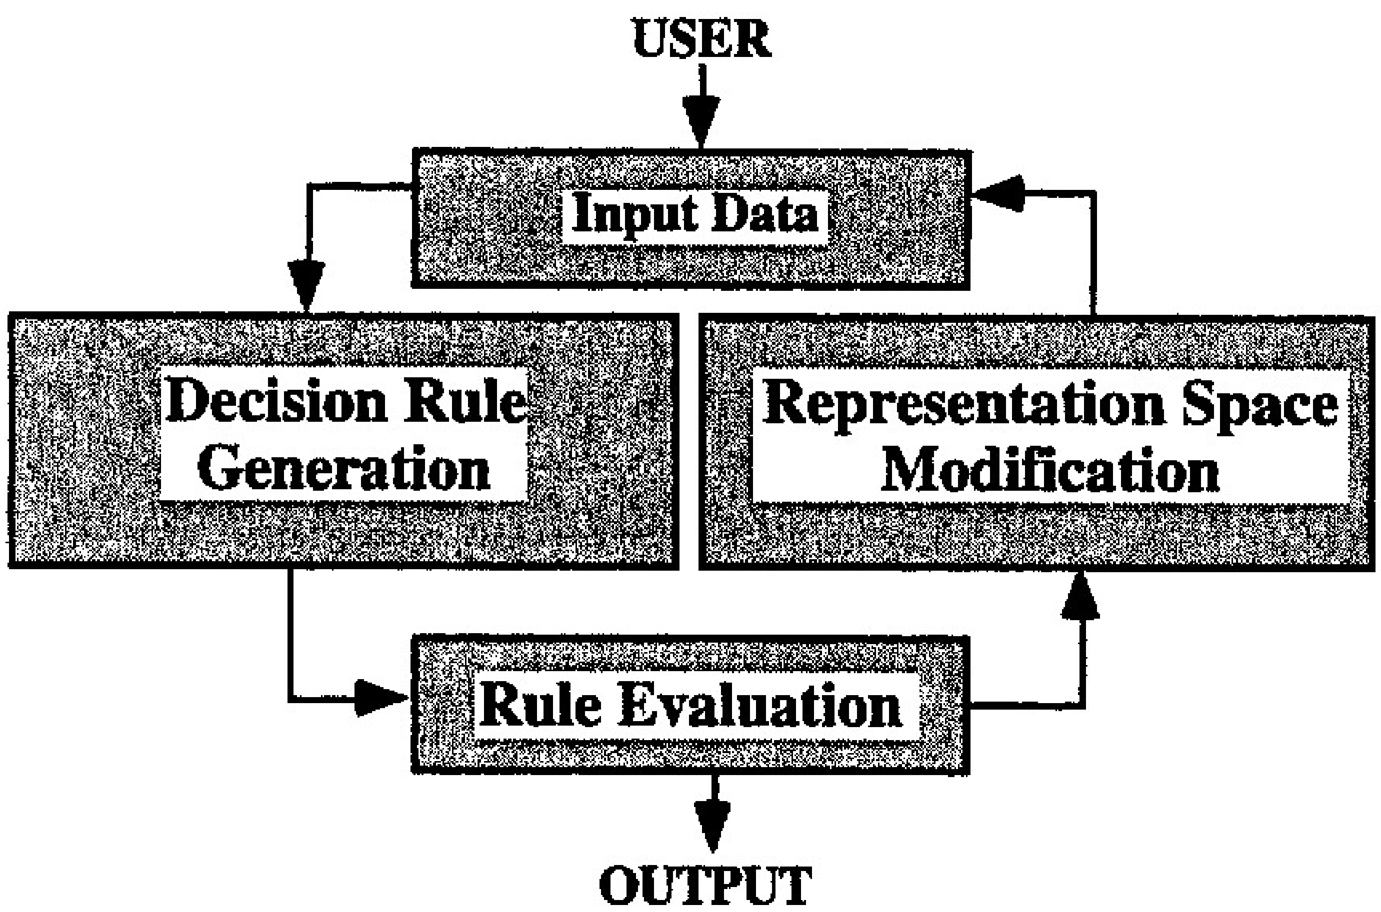

A functional diagram of the AQ17-MCI program.

The system includes three basic mechanisms:
(1) for accepting expert advice about the rules and procedures for generating new attributes;

(2) for analyzing learning examples and generating new attributes as logical or mathematical functions of the original attributes (implemented in AQ17-DCI version, which stands for *data-driven constructive induction)*

(3) for detecting "strong patterns" in the rules generated in one iteration of the rule generation module, and then using these patterns for proposing candidate attributes for a new iteration (implemented in AQ17-HCI version, which stands for *hypothesis-driven constructive induction*).

The attributes generated by these mechanisms are evaluated for their relevance to the problem at hand. If they pass the relevance test, they are used to reformulate original learning examples, and the rule generation module (based on the AQ algorithm) generates new rules. The quality of the rules is determined, and those that pass the quality criterion are stored in the knowledge base.

AQ17-MCI significantly extends current machine learning capabilities, as it is capable for "multi-mechanism" improvement of the original description space. It is a powerful program that represents a new generation of symbolic learning systems, and thus has a potential for important new applications. (See its performance on MONKS' problems, described in the project "A Comparative Study of Learning Methods.")

The support for this research is provided by NSF, ONR and ARPA.

Integrated Learning Systems for Research and Education: EMERALD

(Michalski, De Jong, Kaufman, Bloedorn, Schultz, Wnek)

This project concerns an integrated system of machine learning and discovery programs which serves as a tool in education and research in machine learning and cognitive modeling of learning processes. We have developed the EMERALD system (Experimental Machine Example-based Reasoning and Learning Disciple), which integrates five modules ("robots"), each displaying a capability for some form of learning or discovery:

AQ learns general decision rules from examples of different classes of correct or incorrect decisions made by experts.

INDUCE learns structural descriptions of groups of objects, and determines important distinctions between the groups.

CLUSTER creates meaningful categories and classifications of given entities, and formulates descriptions of these created categories.

SPARC predicts possible future objects or events by discovering rules characterizing the sequence of objects or events observed so far.

ABACUS conducts experiments, collects data, discovers mathematical and logical descriptions of data, and then uses these descriptions for predicting the behavior of some phenomenon.

Each module is represented by a robot figure and employs a different voice (through a voice synthesizer) for communicating with the user.

An earlier and smaller version of the system, called ILLIAN, was a part of the exhibition "Robots and Beyond: The Age of Intelligent Machines," organized by a consortium of eight U.S. Museums of Science (Boston, Charlotte, Fort Worth, Los Angeles, Seattle, Chicago, Minneapolis and Columbus). Support for the development of the exhibit version was provided in part by the Boston Science Museum, the Digital Equipment Corporation, and the University of Illinois at Urbana-Champaign.

EMERALD is the first system of its kind built, which integrates several learning capabilities with natural language processing, voice communication, and a highly user-oriented graphical interface. It enables users to experiment on-line with various learning and discovery programs under a unified control, and to use predefined objects to set different learning tasks for the system. EMERALD was developed under the direction of Professor R.S. Michalski in collaboration with his students and associates.

The system has recently been adapted for SUN workstations, and used in teaching machine learning. It has also been made available to various universities and organizations.

Although EMERALD modules are demonstrated in the context of certain predefined classes of problems, they are not specifically oriented toward these problems and objects. These modules are domain-independent programs that have already been used or have a potential to be used for concept learning and discovering regularities in such fields as medicine, agriculture, engineering, biology, chemistry, plant control, financial decisions, air traffic control, computer vision and intelligent robots.

Machine Vision through Learning

(Michalski, Duric, Maloof, Zhang, Wnek, Bloedorn)

in collaboration with

Computer Vision Laboratory of the University of Maryland

(Rosenfeld, Aloimonos, Davis)

The goal of this project is to develop a methodology and an experimental system that is capable of learning general visual concept descriptions from specific observed objects, and then use these descriptions to efficiently recognize new objects in a visual scene. It is assumed that the system should be able to recognize objects among other objects in a scene under a variety of conditions, such as changing viewpoints, changing illumination, object overlap, and in the presence of noise in the sensory data.

Our approach is based on a two-prong architecture in which the first prong processes the surface information about objects, and the second prong processes the shape information. Learning unique surface characteristics ("surface signatures") involves problem-oriented transformations of the representation space, and an iterative application of an inductive learning program. The input to the system are classified samples of surfaces.

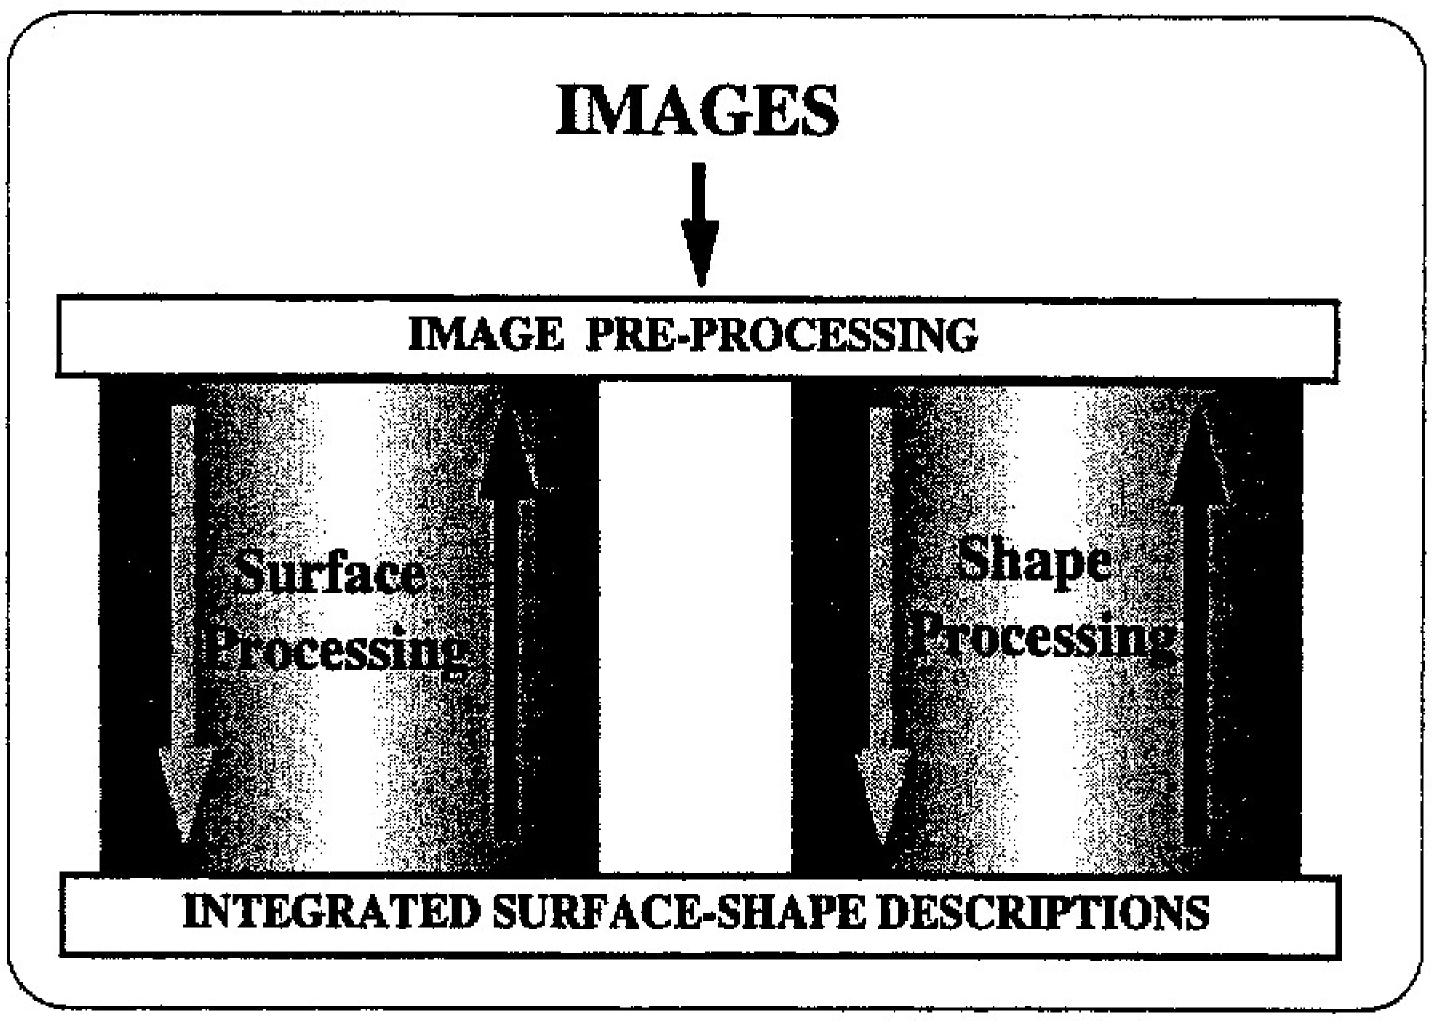

In the object recognition phase, the system applies the learned rules to identify the surface, and then uses this information to generate a set of candidate hypotheses about the object's identity.

These hypotheses are then employed to retrieve specific 3D structural models of the objects from a knowledge base. We use CAD/CAM descriptions of objects to discover their characteristic structural and symbolic features and feature relations, and to learn recognition strategies. These processes are also driven by vision tasks, such as localization, recognition and inspection. Learned models are then used to determine characteristics of objects sufficient to identify the object in the scene. These discriminatory characteristics are determined by a process called *dynamic recognition.*

The research on this project is conducted in collaboration with the Computer Vision Laboratory of the University of Maryland. Jointly, we organized the NSF/ARPA workshop on Machine Vision and Learning which was held in Harpers Ferry, WV, in October 15-17, 1992.

This project is supported by an ARPA grant administrated by the Air Force Office of Scientific Research.

A Methodology for Studying Vision through Learning:
Multilevel Image Sampling and Transformation Methodology (MIST)

(Michalski, Duric, Bala, Maloof, Zhang)

The goal of this project is to develop a system that can learn descriptions of visual objects (images, visual sources, visual scenes) and to use these descriptions to recognize unknown objects. We have developed a general methodology for this purpose, called *multilevel image sampling and transformation methodology* (MIST).

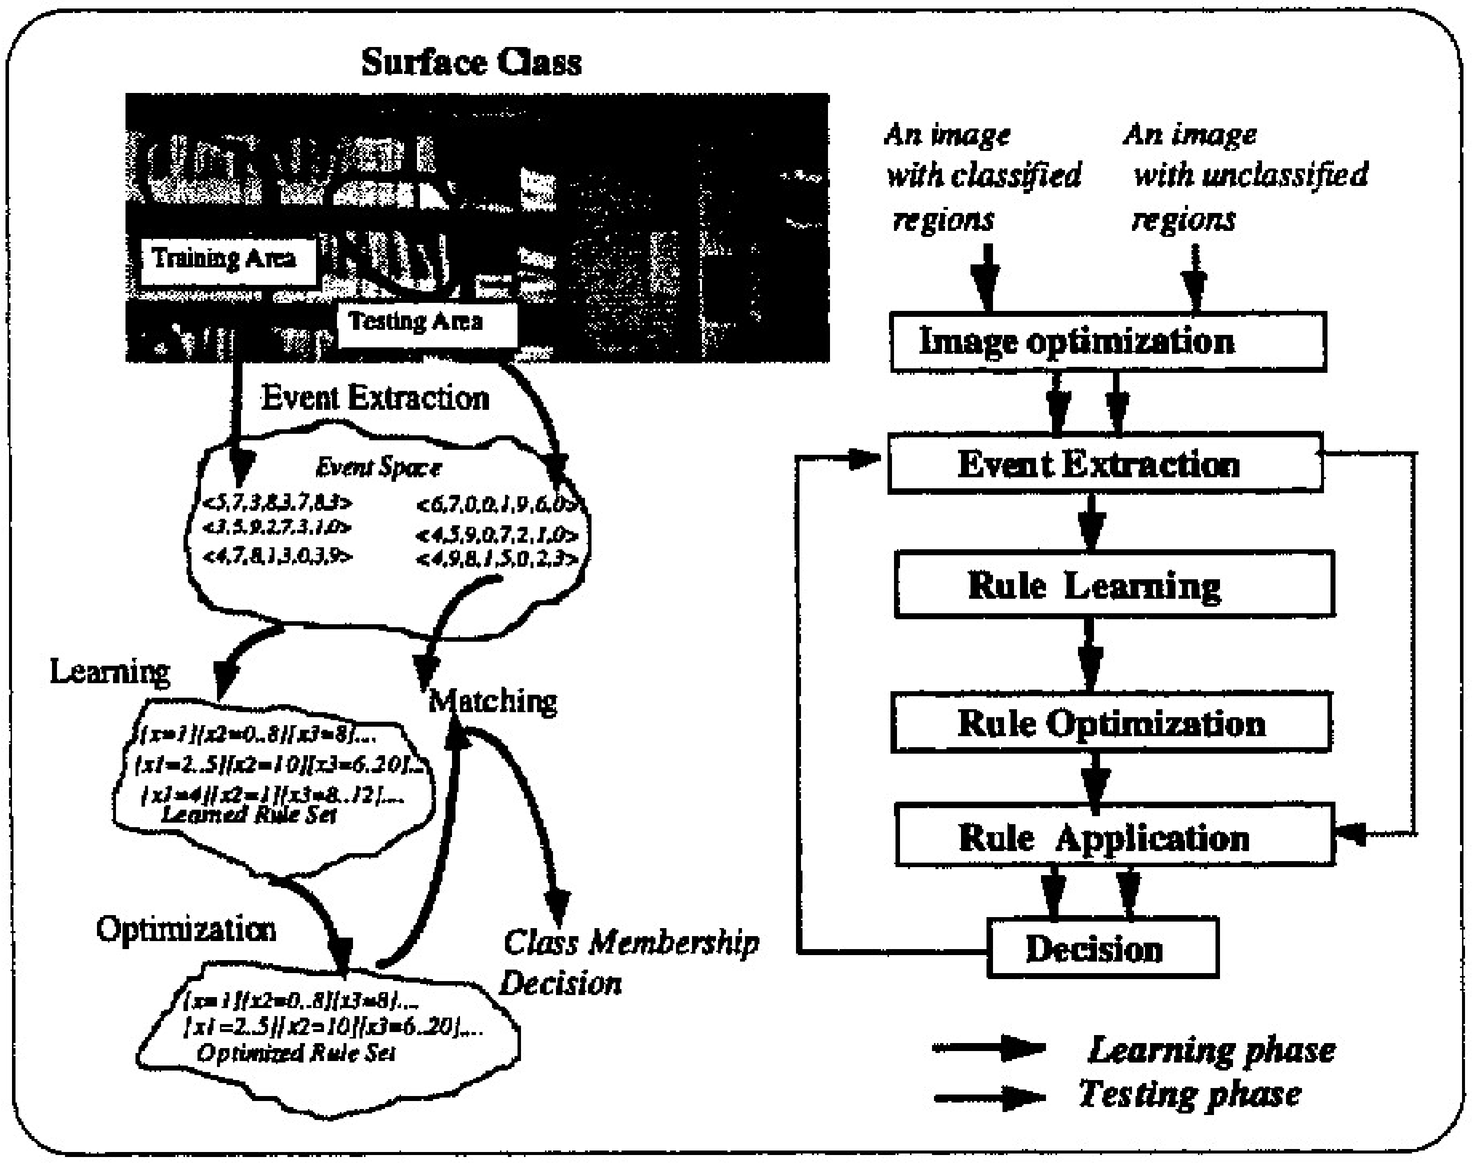

The basic idea under this project can be explained as follows (Figure). Given an image with labeled samples of different surfaces, the learning system determines a sequence of operators that transform the image to a "symbolic" image, in which picture elements are labels of corresponding surface areas. The sequence of operators that produces such a labeling serves as a surface description ("surface signature"). A surface description is a logical expression in disjunctive normal form associated with a decision class (here, a texture class). Each conjunction in this expression together with the associated decision class can be viewed as a single decision rule. The basic operator in the process of generating surface description is an application of a set of logic-style rules to transformed surface samples. The rules can be applied in parallel, and serve as "logical templates" that are applied to "events" (attribute vectors) representing surface samples. To recognize an unknown surface sample, the system matches it with all candidate surface descriptions. This is done by applying decision rules to the events in the sample. For each event, the class membership (surface class) is determined. The assignment of the sample to a given decision class (surface) is based on determining which of the candidate classes gets the majority (or) plurality of votes. Thus, even if some events in the sample are incorrectly recognized, the classification of the sample may be correct.

A series of experiments is conducted with gradually increased complexity of data, increased influence of noise, and under variety of other external conditions. The images used for experiments are divided into two groups: office environment images, and outdoor scene images.

This project is supported by an ARPA grant administrated by the Air Force Office of Scientific Research.

Learning Vision Tasks by Integrating Symbolic and Neural Net Learning

(Michalski, Bala, Bloedorn, Zhang)

The project concerns the development of a novel multistrategy learning methodology that is specifically oriented toward vision learning. The methodology combines symbolic rule learning and neural-based learning strategies in order to achieve high efficiency and accuracy in learning visual object descriptions, and in applying these descriptions to rapid object recognition.

The initially developed vision system has several advantages: it can be easily modified and applied to new problems (due to learning), its learning speed can be at least an order of magnitude faster than neural net learning (due to "symbolic pre-structuring" of the net), it has short recognition times (due to its parallel architecture), and its underlying recognition rules are easy to understand by a human operator (due to the symbolic knowledge representation of the basic decision rules). The developed system was experimentally applied to natural scene recognition.

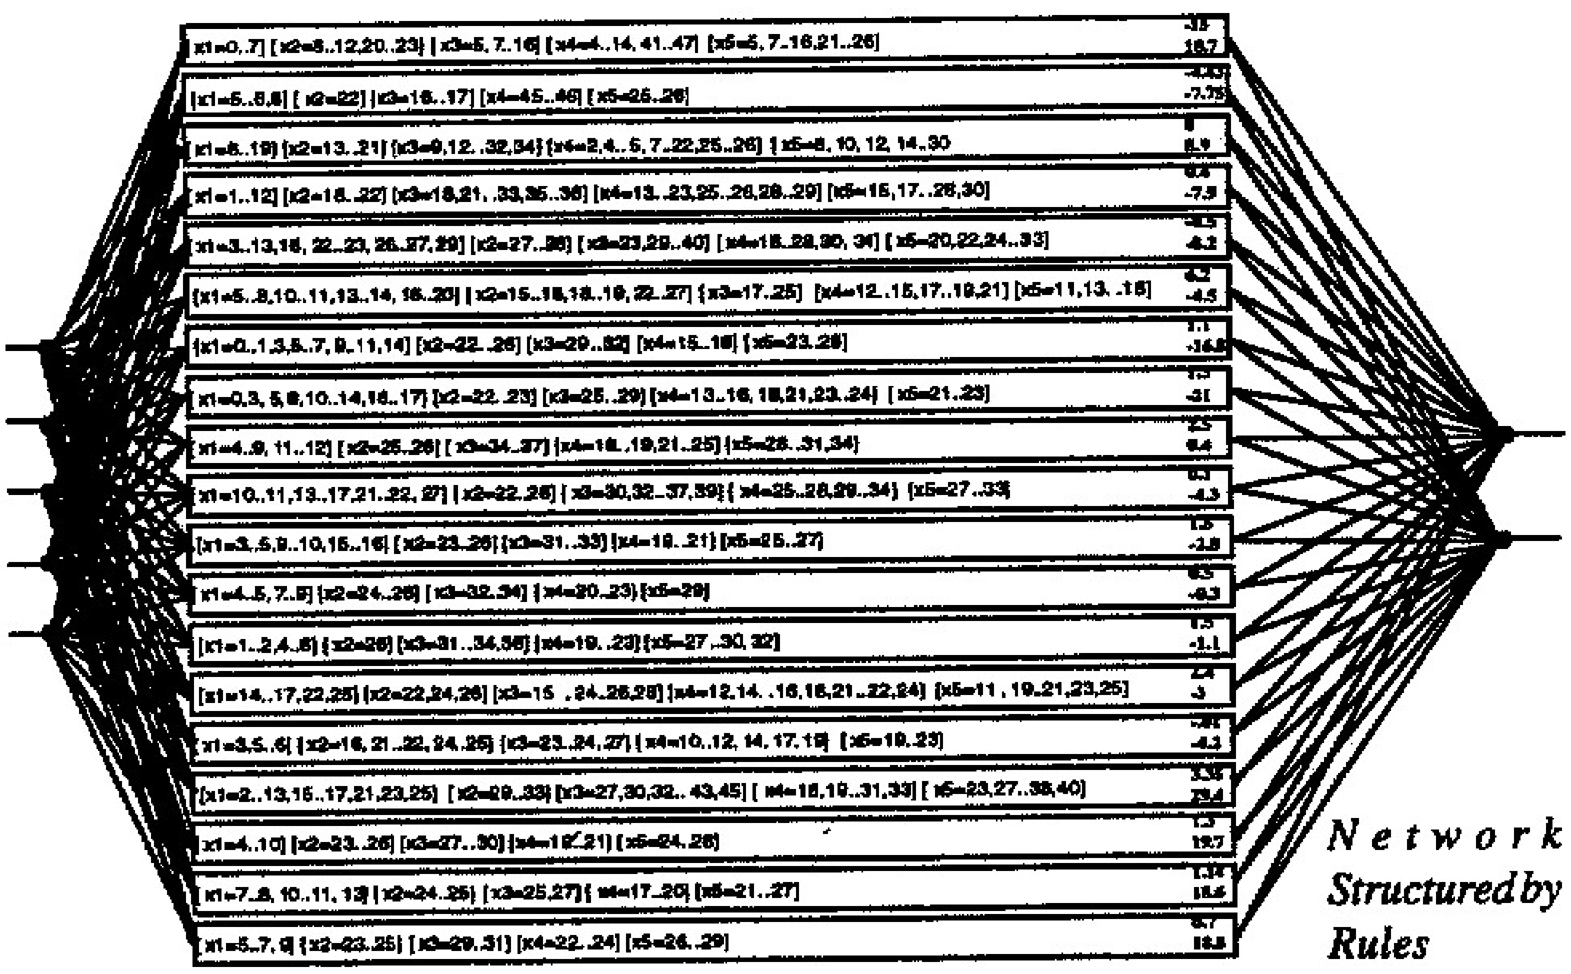

Network Structured by Rules

The method is a special case of the MIST methodology. It works in two stages: 1) *rule learning* using the AQ algorithm. This phase generates rules that generally and approximately describe the training examples, 2) *neural net learning to* determine the final visual concept description. The network is structured according to the rules obtained in stage. Each node in the hidden-layer of the network corresponds to a single rule. The degree of match of an example to the rule represents node activation. This activation value is input to the sigmoid transfer function associated with each node. Weight values for the connections between nodes and outputs are obtained using backpropagation method.

LEARNING APPROACH	TRAINING TIME	RECOGNITION TIME	RECOGNITION ACCURACY	
			Pixel-based	Window-based
Symbolic Learning (AQ-15 program)	6 sec	500 sec.	88%	~100%
Neural Network Learning (BP with no hidden units)	no convergence	N/A	N/A	N/A
Neural Network Learning (BP and 1 layer of hidden units)	5960 sec.	25 sec.	85%	~100%
Integrated Learning (AQ-15 - 1 layer)	120 sec.	20 sec.	92%	~100%

Comparison of Learning /Recognition Times and Recognition Accuracy

Learning to Recognize Shapes

(Michalski, Duric, Maloof)

The goal of this research is to apply inductive learning methods to problems of 2D shape recognition under highly variable perceptual conditions. The *multilevel image sampling and transformation* (MIST) methodology is being used to detect blasting caps in x-ray images of luggage. An intelligent system capable of quickly and reliably performing this task could be used to assist airport security personnel in baggage screening.

We have acquired x-ray images of luggage containing blasting caps that appear at differing degrees of occlusion and at various orientations with respect to the x-ray source. Task-oriented image transformations are used to segment blasting caps and other objects, and to extract training events, which are vectors of attribute values. These training events serve as input to the learning process which induces descriptions of shape that are robust with respect to planar rotation and translation and partial occlusion. Induced shape descriptions can be used to recognize unknown objects.

Various symbolic, non-symbolic and statistical learning approaches are being investigated for acquiring descriptions of shape, including AQ15c, neural networks, and k-nn. These learning approaches are compared using predictive accuracy, and learning and recognition time. Experimental results have demonstrated strong advantages of AQ15c over neural networks and k-nn. AQ15c also has the advantage of producing comprehensible symbolic descriptions that can be optimized by either a human or by a machine process in post-learning phases.

This project is supported by an ARPA grant administrated by the Air Force Office of Scientific Research.

Building An Intelligent Agent for Software Testing

(Rine, Michalski, Imam)

This project is concerned with building an intelligent agent for software testing. The agent assists a user in testing whether a software system satisfies a set of requirements. The agent learns a decision structure (a directed acyclic graph) consisting of nodes, branches and leaves. In this structure each node corresponds to a particular test strategy of the software, each branch corresponds to the possible outcomes of the given test strategies, and each leaf corresponds to a possible general decision of testing the software (e.g. faulty or not).

The process of testing software consists of two steps; first a set of sufficient and inexpensive test strategies must be determined, and then for these testing strategies the best way to perform them is learned. The agent generates a decision structure by selecting the optimal set of testing strategies for testing the software. This optimal set of testing strategies is determined by a cost function, determined by user preferences. The agent uses different criteria for adapting the decision structure to fit a decision making situation.

This agent is designed to overcome the difficulties that face experts in software testing such as selecting the optimal and/or the minimum set of necessary testing strategies for different software systems.

Dynamically Interlaced Hierarchies (DIH)

(Michalski, Alkharouf, Utz)

This project concerns a development of a new type of knowledge representation that facilitates all kinds of inferences and is thus particularly relevant to the development of multistrategy task-adaptive learning. Dynamic Interlaced Hierarchies (DIH) is based on psychological research into human semantic memory structure and utilizes hierarchies as its basic organizational principle. By storing new knowledge as links between hierarchically organized concepts, a conceptual framework is constructed that can represent very diverse and complex forms of knowledge as well as various knowledge transformations.

DIH uses type and part hierarchies of concepts as *background knowledge*, or knowledge considered to be relatively stable and unchanging. Statements or facts are stored as links between concepts and are considered *dynamic knowledge*, as these links are constantly being created and modified, strengthened or weakened. These links have numeric factors (or "merit parameters") attached that affect the strength of the relationship between the various concepts. Rules and dependencies are bidirectional, each with a separate forward and backward "strength".

Inference patterns such as generalization/specialization, abstraction/concretion, and similarity are easily visualized in DIH. Also these inferences are facilitated, since the procedure consists of manipulating links between hierarchies. Creating new links between concepts represents learning. In this way learning builds upon the background knowledge of the hierarchies and the dynamic knowledge already in place.

The support for this project is provided by grants from the NSF, ONR, and ARPA.

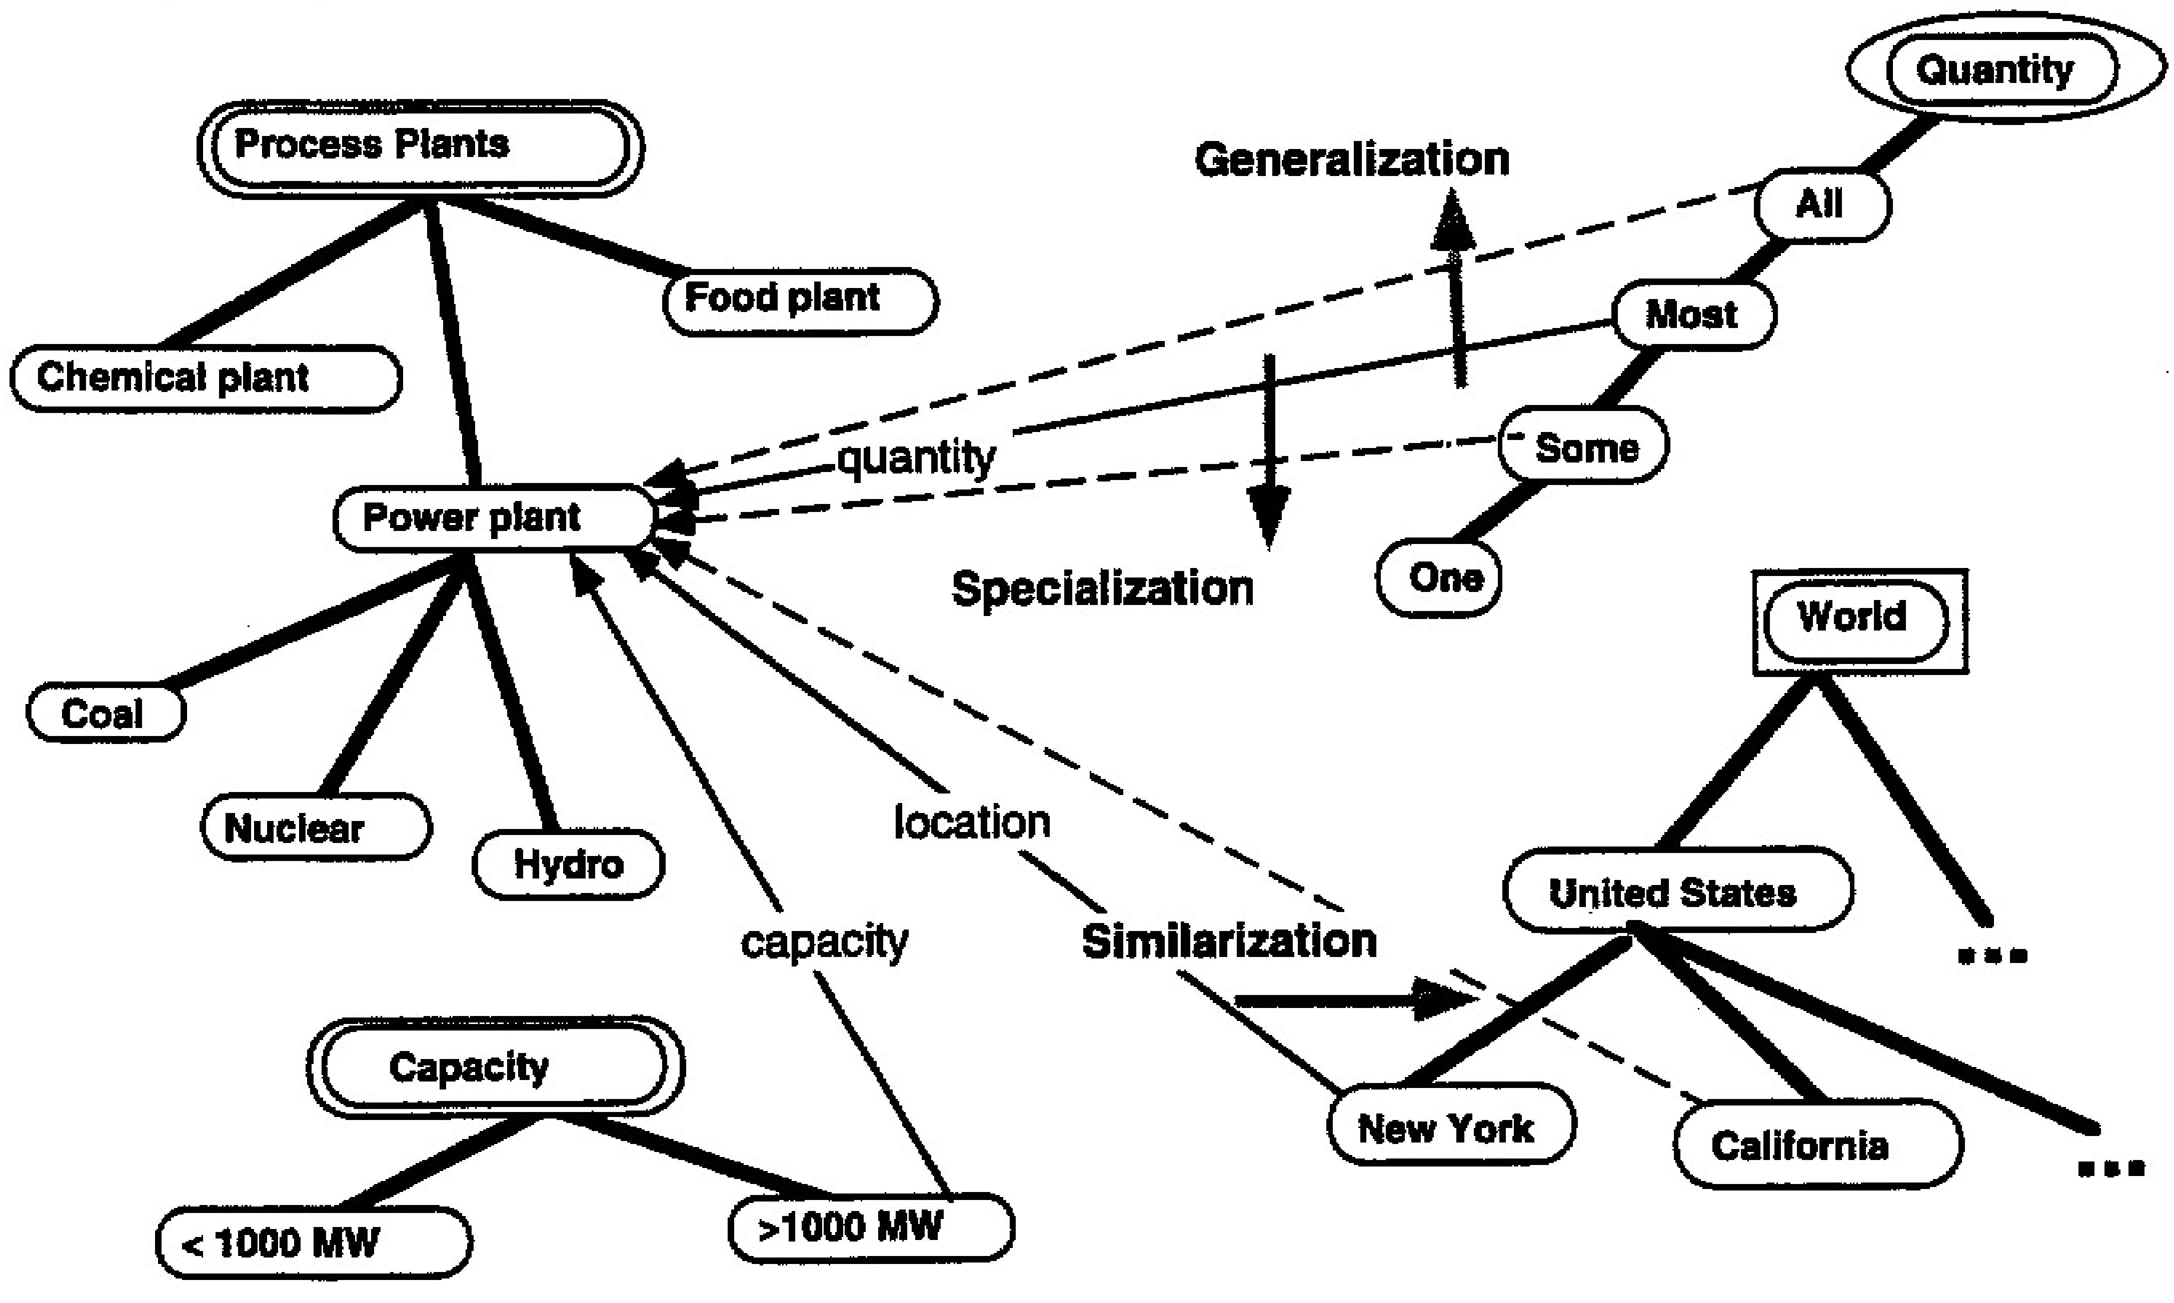

Transmutations of "Most Power Plants in New York have > 1000 MW Capacity."

Genetic Algorithms and Adaptive Systems

(De Jong, Spears, Schultz, Vafaie, Sarma, Potter, Stefanski)

Genetic algorithms explore analogies to evolution as a model for building systems capable of learning and adapting to changes in the environment and in the requirements for performing various tasks. Genetic algorithms are particularly effective for learning in unstructured domains for which strong domain theories are unavailable.

Current research includes the use of genetic algorithms: (1) in the design of adaptive agents capable of improving their performance, in complex and unpredictable environments (2) as a heuristic difficult combinatorial optimization problems and engineering problems with complex, non-linear constraints, (3) as a complement to symbolic learning in multistrategy systems, (4) as a parallel search technique capable of exploiting the power of highly parallel and distributed computing architectures, and (5) as a technique for evolving both the structure and weights of neural networks.

These ideas are being tested in a variety of contexts including: (1) adaptive game playing programs, (2) texture recognition systems which can operate in the presence of noise without strong domain models and are capable of creating and testing new feature recognition sets, (3) adaptive control programs for robot navigation and for performing tasks in changing and uncertain environments, and (4) the use of Connection Machines for massively parallel adaptive search.

The conceptual foundations for this work draw from and require extensions to the areas of genetic algorithms, classifier systems, neural networks, and classical machine learning techniques. Support for this project is provided by grants from NRL, ARPA and ONR.

Cognitive Models of Plausible Reasoning

(Michalski, Sklar)

The ability to reason plausibly, that is to derive useful conclusions from imperfect premises, is one of the most remarkable properties of the human mind, and a key to understanding intelligent behavior. In plausible reasoning, the premises may be incomplete, uncertain, imprecise or only partially relevant to the task. Yet, people are able to make useful conclusions from such premises. The initial core theory of human plausible reasoning was developed by Collins and Michalski. The goals of this research are to develop a computational theory and models of plausible reasoning, to validate the theory by experiments involving the models and human subjects, and to apply it to developing a new approach to knowledge representation, filling gaps in databases, and *dynamic recognition.*

This project is supported by a grant from ONR and NSF.

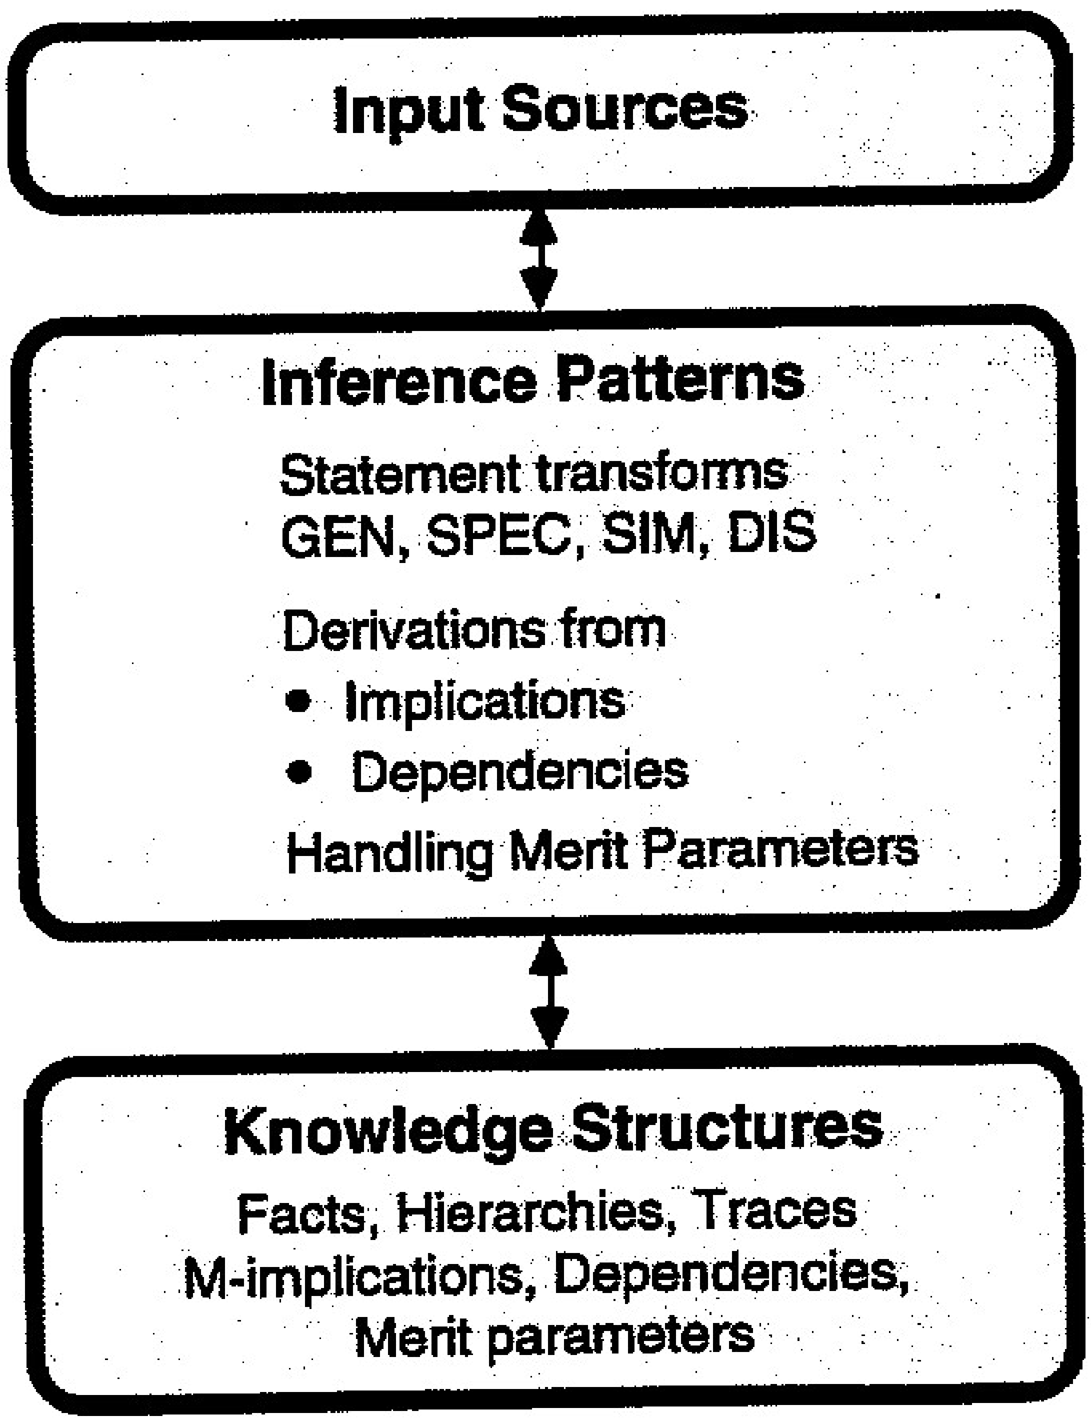

Diagrammatic Visualization of Learning Processes (DIAV)

(Michalski, Wnek)

The DIAV project concerns building a system for visualizing learning processes. The system employs a planar model of a multidimensional space, spanned over a set of discrete attributes. The model is in the form of a diagram, in which each cell represents a unique combination of attribute values. For example, in the figure below, the cell marked by X represents the vector: (HeadShape = square, Holding = sword, JacketColor = red, Smiling = no). Each attribute partitions the diagram into areas corresponding to individual values of this attribute. Conjunctive rules correspond to certain regular arrangements of cells that can be easily recognized visually. The shaded areas marked *Class 1 and Class 2* correspond to the rules stated at the bottom part of the figure.

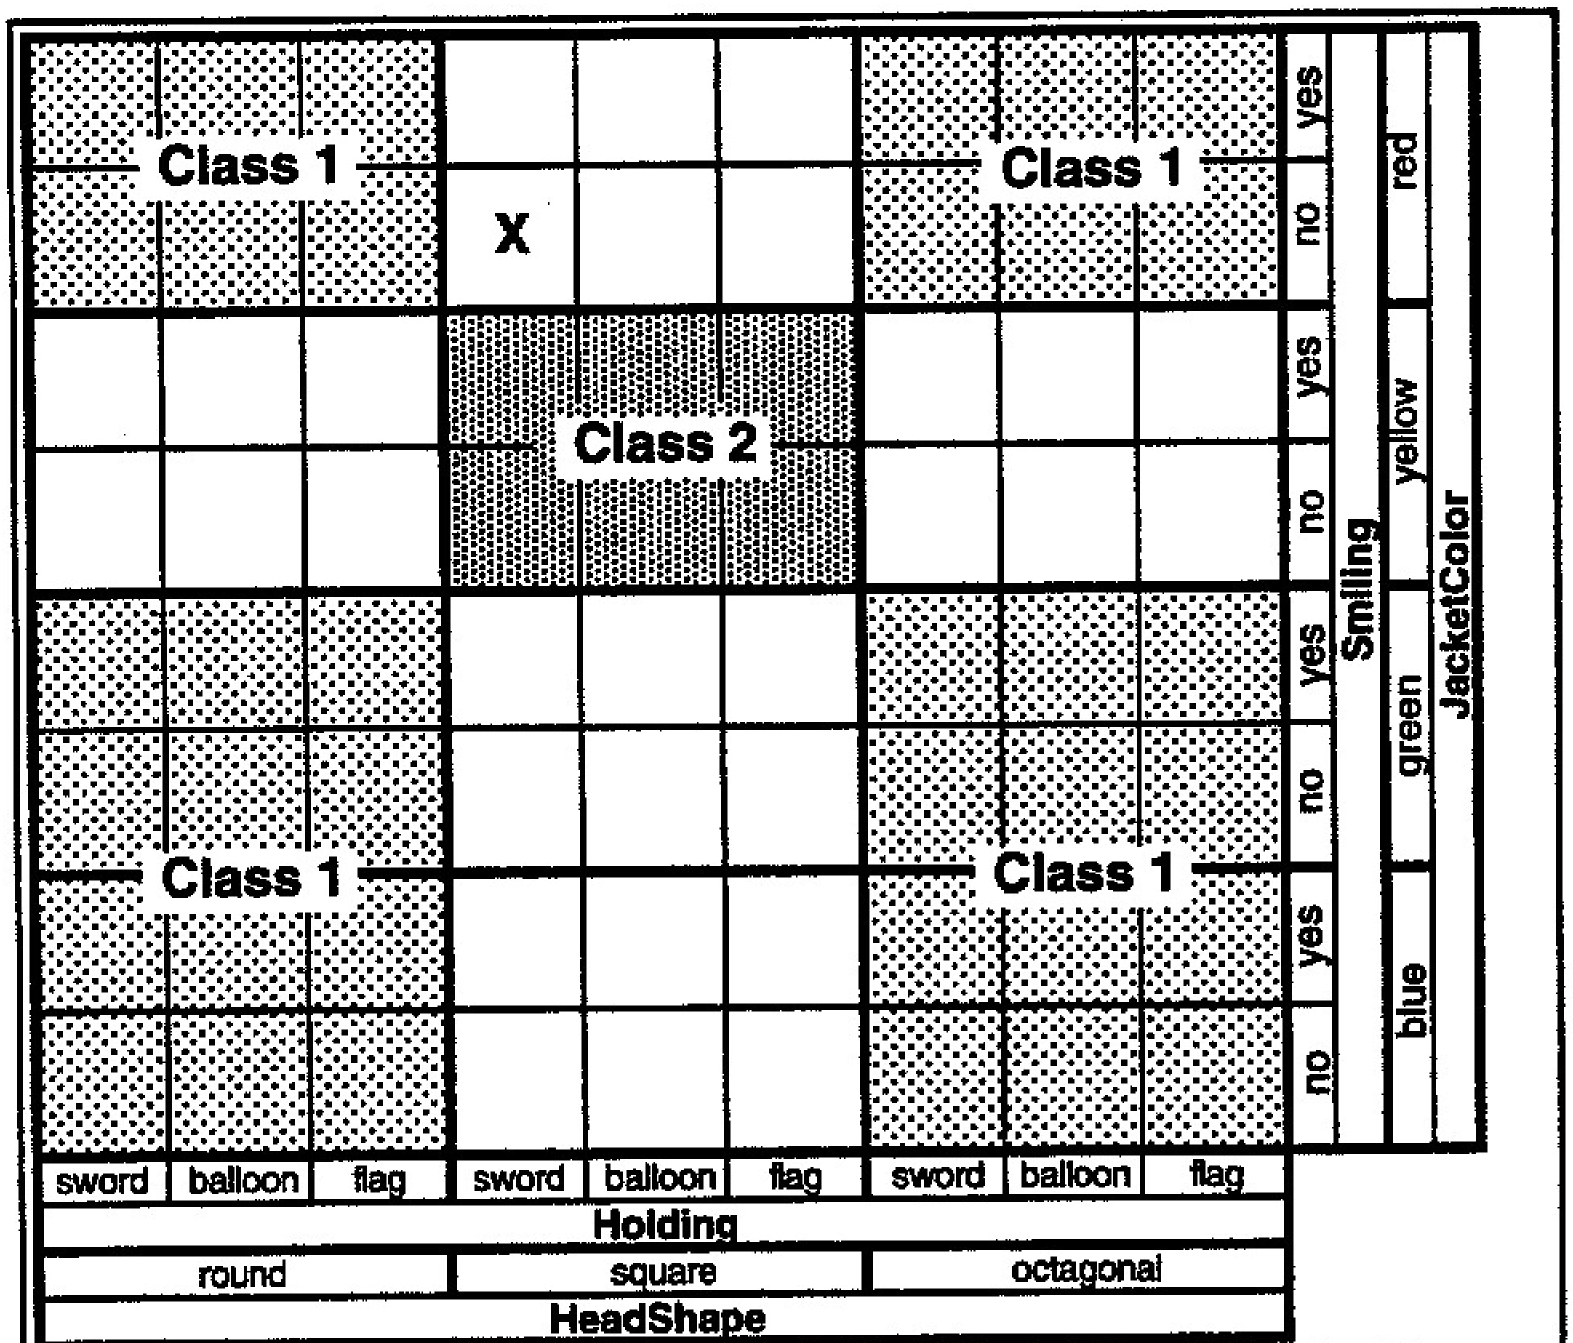

Class 1 <:: **JacketColor** is red, green or blue *and* **HeadShape** is round or octagonal

Class 2 <:: **HeadShape** is square *and* **JacketColor** is yellow

If the diagram is used to represent target and learned concepts, then the difference between them represents the *error area,* which identifies all errors in the learned concept.

The diagram can also display results of any operation on the concept, such as generalization or specialization, or any change of the description space, such as adding or deleting attributes, or their values. Another interesting feature is that it can also visualize concepts acquired by non-symbolic systems, such as neural nets or genetic algorithms. Using the diagram one can directly express the learned concepts in the form of decision rules. Thus, the diagram allows one to evaluate both the quality and the complexity of the results of symbolic, as well as non-symbolic learning. The implemented system, DIAV-2, can display description spaces with up to a million events, i.e., spaces spanned over up to 20 binary variables (or correspondingly smaller number of multiple-valued variables). The system has proven to be very useful for analyzing behavior of learning algorithms. It is available to universities and industrial organizations.

This project is supported by a grant from ONR.

Learning Problem-Optimized Decision Trees from Decision Rules

(Michalski, Imam)

This project is concerned with learning problem-optimized decision trees from rules. A standard approach to determining decision trees is to learn them from examples. A disadvantage of this approach is that once a decision tree is learned, it is difficult to modify it to suit different decision making situations. Such problems arise, for example, when an attribute assigned to some node cannot be measured, or there is a significant change in the costs of measuring attributes or in the frequency distribution of events from different decision classes. An attractive approach to resolving this problem is to learn and store knowledge in the form of decision rules, and to generate from them, whenever needed, a decision tree that is most suitable in a given situation. An additional advantage of such an approach is that it facilitates building compact decision trees, which can be much simpler than the logically equivalent conventional decision trees (by compact trees are meant decision trees that may contain branches assigned a set of values, and nodes assigned derived attributes, i.e., attributes that are logical or mathematical functions of the original ones). The project describes an efficient method, AQDT-1, that takes decision rules generated by an AQ-type learning system (AQ15c or AQ17), and builds from them a decision tree optimizing a given optimality criterion. The method can work in two modes: the standard mode, which produces conventional decision trees, and compact mode, which produces compact decision trees. The preliminary experiments with AQDT-1 have shown that the decision trees generated by it from decision rules (conventional and compact) have outperformed those generated from examples by the well-known C4.5 program both in terms of their simplicity and their predictive accuracy.

AQ15 rules for the MONK-1 problem

Positive rules
1 [x5 = 1]
2 [x1 =3][x2 = 3]
3 [x1 = 2][x2 = 2]
4 [x1 = 1][x2 = 1]

Negative rules
1 [x1 = 1][x2 = 2..3][x5 = 2..4]
2 [x1 = 2][x2 = 1..3][x5 = 2..4]
3 [x1 = 3][x2 = 1..2][x5 = 2..4]

AQ17 rules for the MONK-1 problem

Positive rules

Negative rules

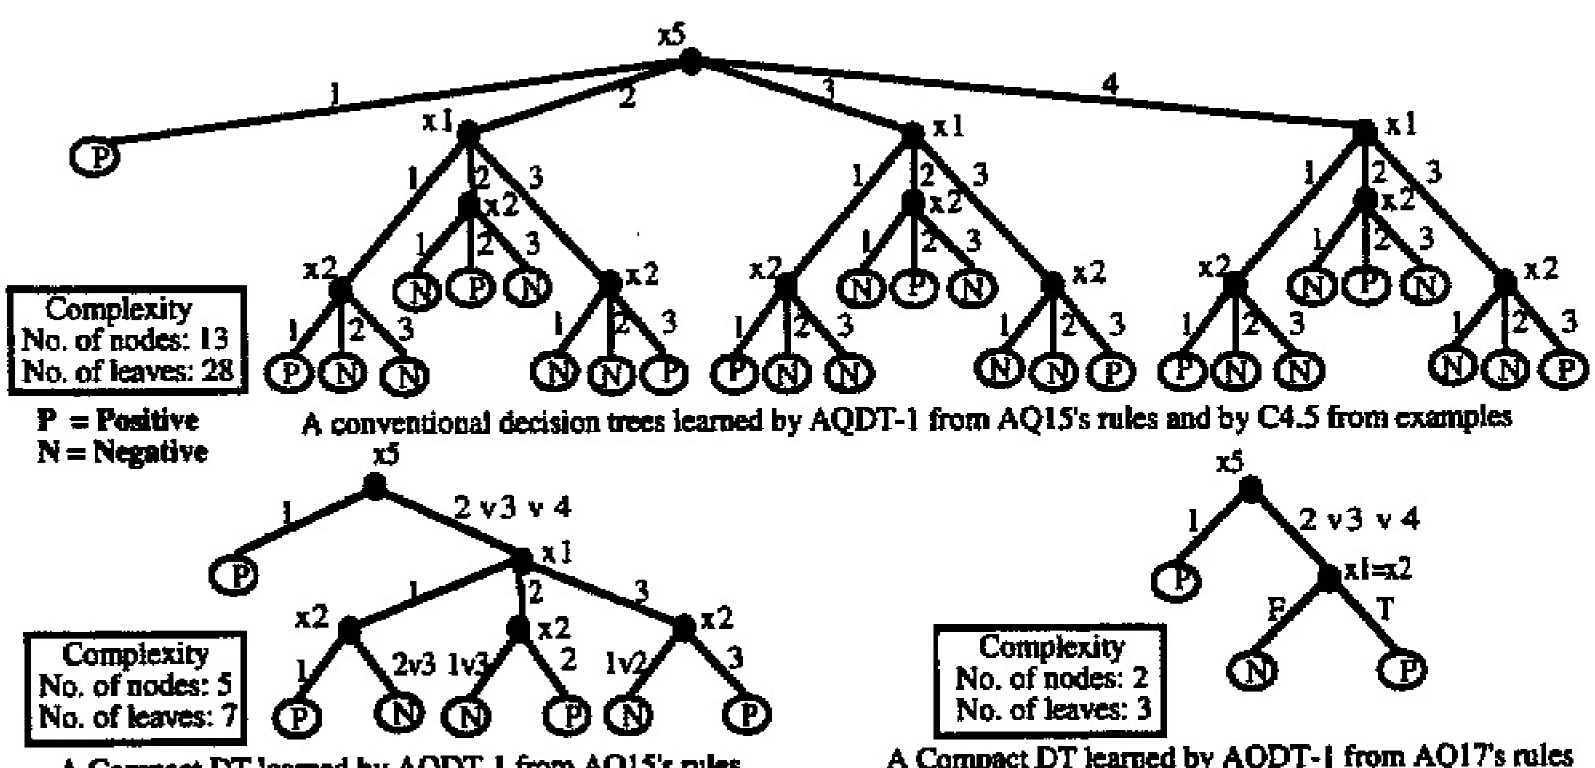

Decision trees learned by different methods from identical input data

Advisory Agents with Learning Capabilities

(Michalski, Kaufman, Imam, Ribeiro)

Standard systems do not have learning capabilities. Their knowledge bases are built entirely by hand-encoding of an expert's knowledge. Such a process is time-consuming and prone to error. This project is concerned with the development of a PC-based expert system shell with learning capabilities. The system incorporates a knowledge base for storing rules and a data base for storing facts and examples. It has a learning program for rule acquisition, and a powerful inference mechanism.

The project is based on our earlier experience with ADVISE and AURORA systems. ADVISE is a large-scale inference system with rule learning capabilities and multiple control schemes. The system served as a laboratory for experimenting with methods for knowledge acquisition, multiple knowledge representation and machine learning. Aurora is a PC-based inference system, and an expert system shell that incorporates a program for incremental rule learning and improvement.

A related project concerns a method for discovering qualitative and quantitative models from data characterizing the behavior of a system. This method builds upon our experience with the ABACUS system for quantitative discovery. The current system is capable of determining a set of equations that fit a given set of datapoints, and a set of symbolic descriptions characterizing preconditions for the application of these equations. ABACUS integrates methods for data-driven quantitative discovery, concept learning from examples and conceptual clustering. This research has applications in building advanced expert systems and discovering quantitative and qualitative regularities in data.

This project is sponsored by ONR and ARPA.

The Role of Learning Goals in Multistrategy Learning

(Michalski, Utz)

Learning arises from an intelligent individual's inability to reason and comprehend with its current knowledge. From prior research, every learning task requires background knowledge, sufficient inputs and a learning goal to achieve success. In multistrategy learning, though, the student faces additional complexity: here, the pupil must derive several learning goals for sequential and possibly parallel application in the learning process.

The aims of this research project are twofold: (1) to experiment with a formalism to specify goals for multistrategy learning and (2) to construct a processing mechanism to generate and apply learning goals appropriately. This project is based on the MTL methodology supporting the Inferential Theory of Learning. Such a formalism must be accurate and complete to enable the processing mechanism to create explicit learning goals to understand the context, direct the procedure, and evaluate the results of multistrategy learning tasks. The formalism must be domain-independent as well as task-adaptive.

Proper specification is essential for success. Learning goals must be specified to enable an examination of any newly acquired or more efficient knowledge. When the examination indicates that the results are implausible or incompatible with the target knowledge, the process must be capable of trying again. It should reselect inputs or learning strategies on the advice of the original learning goals or, where necessary, regenerate alternative learning goals to redirect the learning task.

A paper by R.S. Michalski and Ashwin Ram from Georgia Tech reports some of this work (see http://www.mli.gmu.edu/publications.html or Michalski and Ram, 1995).

Two-Tiered Representation of Flexible Concepts

(Michalski, Bergadano, Bloedorn, Matwin, Zhang)

This project concerns a novel method for representing and learning *flexible concepts*, by which we mean concepts that lack precise definition and whose meaning is context-dependent. The method is based on the *two-tiered* (TT) representation of such concepts. A TT representation consists of two parts: 1) *base concept representation* (BCR) that captures essential and typical concept properties; and 2) *an inferential concept interpretation* (ICI) that defines allowable BCR transformations, and handles exceptional and context-dependent properties. Numerous experiments with learning systems implementing this idea (e.g., AQ15, POSEIDON and FCLS) have shown that it often leads to an order-of-magnitude reduction of the size of the knowledge base, and at the same time to an increase of the accuracy of the knowledge base.

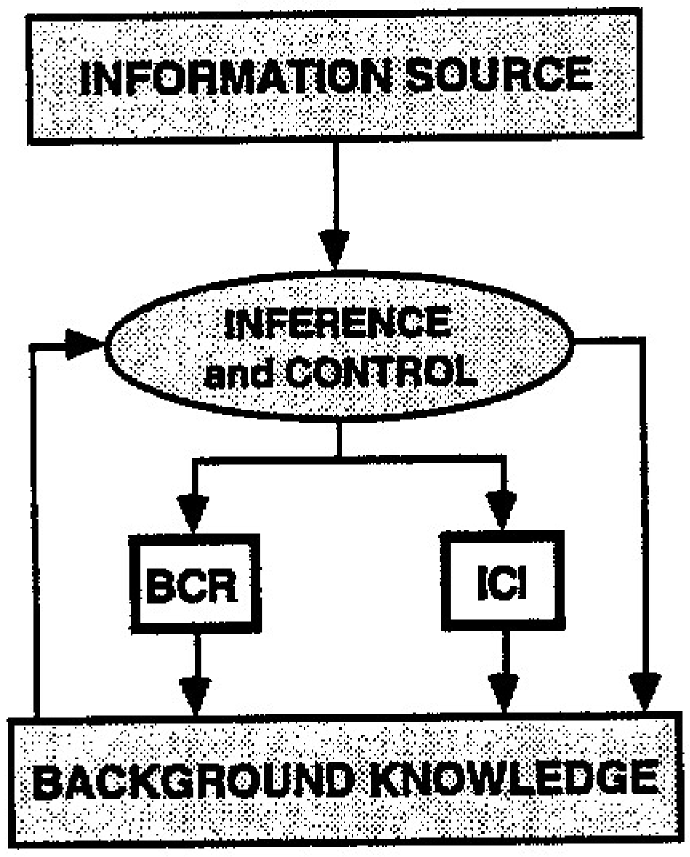

The support for this project is provided by ONR and ARPA.

Dynamic Recognition

(Michalski, Bloedorn, Zhang)

Any recognition process involves making a connection between a concept representation stored in the system's memory and a stream of observational data. Present recognition systems attempt to recognize objects by matching descriptions with the data stream. If the input data satisfies rules characterizing an object, the object is recognized. To implement such a system for practical tasks, a very large number of rules may be required.

This aspect severely limits present recognition systems, as it prevents them from being applied to the recognition of a large number of objects. In contrast to this approach, humans can recognize objects from a great variety of different cues, without "matching" rules. For example, they can recognize a known person from seeing a face, a silhouette, the characteristic way of walking, hearing the person's voice, or even from observing the person's gesticulation or seeing his/her shoes.

The *dynamic recognition* (DR) approach (initially proposed by Michalski in 1986) overcomes this problem by using inductive inference to dynamically determine *discriminant* object descriptions from *characteristic* object descriptions, and this allows the system to avoid matching rules. Only one characteristic description per concept is stored in memory. Potentially, the DR method can efficiently handle a great variety of different practical recognition problems. An initial implementation of the system has strongly supported the theoretical expectations.

The support for this project is provided by ONR, ARPA and AFOSR.

Active Databases for Approximate Consistency Maintenance

(Kerschberg, Seligman)

As collections of distributed, heterogeneous knowledge- and data-based systems enter into federations, there is a real need to define architectures and mechanisms for these systems to communicate and to exchange information relevant to their problem-solving tasks. This project involves a novel approach to supporting approximate consistency maintenance between primary and secondary copies of data in such an environment. The approach uses an intelligent interface to active databases called a "Mediator for Approximate Consistency" (MAC). The MAC has several unique features: (1) it permits applications and federation components to specify their consistency requirements declaratively, using a simple extension of a frame-based representation language, (2) it automatically generates the interfaces, rules, and other database objects necessary to enforce those consistency requirements, shielding the application developer from the implementation details of consistency maintenance, and (3) it provides an explicit representation of consistency constraints in the database, which allows them to be queried and reasoned about. Because of the need to maintain various kinds and degrees of consistency, which may be short of 100% consistency, a hardware model of cache consistency is not appropriate. What is required is a quasi-cache, which contains quasi-copies, which are cached copies whose values are allowed to deviate in controlled ways from the primary copies of those objects. This approach automatically generates the database objects necessary to enforce the consistency constraints specified in a declarative form. The figure below illustrates the operation of the MAC in its interactions with a single active DBMS. The MAC is composed of two major submodules: the translator, which handles communication from the application to the active database, and the mapper/message handler, which handles communication from the active database to the application. The translator accepts the declarative specification of a given class' consistency requirements as it appears in a quasi-cache definition and translates it into the following: queries to be executed immediately, rules for monitoring the future state of the database, and data definition language commands which result in the creation of and updates to consistency constraint objects in the database. The queries which are to be executed immediately are used to populate the quasi-cache with those instances of the newly defined class for which the selection-conditions are satisfied at quasi-cache initialization time. The rules are of three types: *selection-rules*, which are used to monitor the database for future occurrences of the selection-conditions, *retraction-rules*, which are used to monitor the database for the retraction-conditions, and *consistency-rules*, which are used to monitor the database for conditions which require refreshing of quasi-copies. The MAC has been implemented using CLOS and the Ithaca object-oriented database system. Novel features include the definition of a Class called Rules so that the rules are treated as objects in this framework. Dr. Len Seligman successfully defended his doctoral dissertation on Sept. 28, 1994.

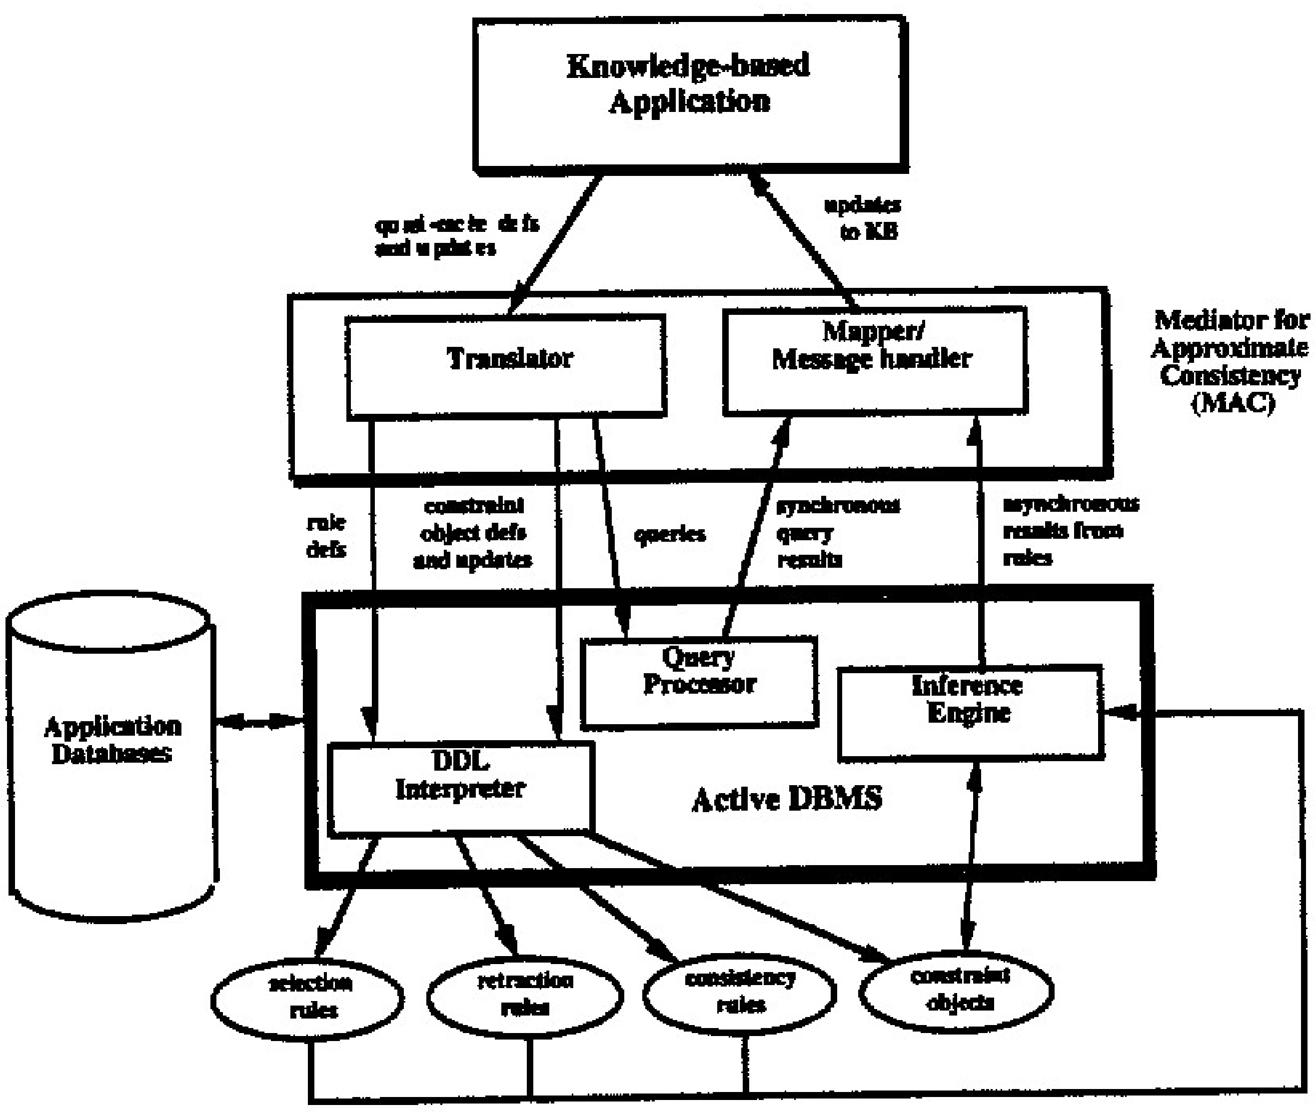

Mediator Architecture for Approximate Consistency.

A Comparative Study of Learning Methods

(Michalski, De Jong, Bala, Pachowicz, Wnek, Bloedorn, Kaufman, Vafaie)

The goal of this project is to experimentally compare and analyze different machine learning methods and approaches to get an insight to their practical applicability. The study involves testing the performance and evaluating properties of concept descriptions learned by different systems on the same set of problems. One study involved a comparison of simple exemplar-based descriptions, the complete and consistent rule-based descriptions obtained by AQ15, the "top rule" descriptions, and two-tiered (TT) descriptions. In the study, the TT descriptions outperformed all other descriptions (Bergadano et al., 1992).

Another study involved a comparison of a neural net, a genetic algorithm, a decision tree (C4.5), and a rule based (AQ15 and AQ17-HCI) methods, as applied to learning logic-style descriptions. The best results were obtained by AQ17-HCI (see Wnek and Michalski, P94-6).

Recently, MLI researchers participated in a study, performed by various research groups from USA and Europe, on the application of their learning methods, symbolic and/or subsymbolic, to the so-called MONKS' problems (so named because they were created in Priory Corsendonk in Belgium, where the Summer-91 Machine Learning School was held). The first problem, M1, involved learning a standard DNF description, and was expected to be easy for symbolic algorithms, such as AQ (rule learning), and ID-type (decision tree learning). The second problem, M2, involved learning an m-of-n relation that is very difficult to describe in the form of a DNF expression that uses only given attributes. Such problems are supposed to be easy for neural net learning. The third problem, M3, involved learning a DNF type concept from data with 10% noise.

PROGRAM	PROBLEM		
	M1	M2	M3
Assistant Professional	100%	81.3%	100%
ID5R	81.7%	69.2%	95.2%
IDL	97.2%	66.2%	-
ID5R-hat	90.3%	65.7%	-
TDIDT	75.7%	66.7%	-
ID3	98.6%	67.9%	94.4%
ID3, no windowing	83.2%	69.1%	95.6%
AQR	95.9%	79.7%	87.0%
CN2	100%	69.0%	89.1%
PRISM	86.3%	72.7%	90.3%
CLASSWEB 0.10	71.8%	64.8%	80.8%
ECOBWEB	82.7%	71.3%	68.0%
mFOIL	100%	69.2%	100%
Backpropagation	100%	100%	93.1%
Cascade Correlation	100%	100%	97.2%
AQ17	100%	100%	100%
AQ17-DCI	100%	100%	94.2%
AQ17-HCI	100%	93.1%	100%
AQ15-FCLS	100%	92.6%	97.2%
AQ14-NT	100%	-	100%
AQ15-GA	100%	86.8%	100%

A summary of preliminary results of the study is presented in the table beside. The last six rows represent results obtained by various versions of AQ programs developed by our Laboratory. AQ17 programs for constructive induction were described on page 15; AQ15-FCLS combines symbolic and subsymbolic learning; AQ14-NT is a noise-tolerant version, and AQ15-GA combines AQ with a genetic algorithm.

The results show that AQ programs were among the most successful of all competing programs, and that AQ17 achieved *the top 100% predictive accuracy for all three problems.* They also produced very simple and understandable concept descriptions. For details, see the report by Thrun et al. 1991 and the book chapter by Wnek and Michalski (P94-6).

This research is supported by grants from ONR, ARPA and NSF.

The EAST-WEST Challenge:
2nd International Competition of Machine Learning Programs

(Michalski, Wnek, Alkharouf, Bloedorn, Imam, Kaufman, Maloof)

To give some sense of how far machine learning methods for generating classification descriptions from examples had progressed in the last 20 years, and to compare the methods with human classification abilities, the "East-West Challenge" was organized by a group based at Oxford University. It was the second international competition of concept learning systems (the first was organized by CMU researchers on the occasion of the Summer School on ML in Priory Corsendonk in Belgium in 1992, and involved the so-called MONKS problems). The new competition received 65 entries from the US, Canada, Australia and Europe, including an entry from a team at the GMU MLI Laboratory.

The East-West Challenge consisted of three problems dealing with classification in the "trains" domain, as originally defined by R.S. Michalski over 20 years ago. In this domain, trains of variable-length, each with cars of different shapes and load types were classified into "Eastbound" or "Westbound". The goal of the "Challenge" was to discover the simplest rules that govern the classification of the trains (measured by Prolog program). The three problems were as follows:

Competition 1

Michalski's original trains example consisted of five Eastbound and five Westbound trains. Five more trains in each class were added to the set in such a way that the rules from the original example are no longer consistent with the entire set of trains. The task was to find a new "Law", as simple as possible, that correctly classifies all twenty trains into their proper groups.

Competition 2

Allocate Eastbound/Westbound labels to 100 trains provided by the creators of the problem. These trains come unaccompanied by any classifying rule or formula. The best rules (based on the bottom quartile of complexity scores) from Competition 1 serve as oracles for the classification of the trains in the test set of 100.

Competition 3

It was a similar to Competition 1, except for using randomly generated and classified trains. For five separate specific problems, each of which involved five Eastbound and five Westbound trains, the challenge was to find the simplest rules that correctly classify the ten trains into their proper groups. The MLI team used three programs developed by their research group: INDUCE, AQ17-HCI and AQDT. The team placed strongly in all three competitions (the sixth best result in the first competition, the three-way tie in the second competition after initially announced first position, and the fourth best result in the third competition). These scores were the best to come from a US-Based team.

These competitions provided a unique challenge in that they were organized from within the European inductive logic programming community, and as such, the intended rules for the first two competitions were oriented toward a Prolog-style representation. As the Challenge progressed, a healthy discussion arose over what constitutes a "simple" rule; after all, a simple description in one representation may be extremely complex in another.

The MLI team proposed an approximate measure of "cognitive or C-complexity" rather than "Prolog or P-complexity" (since rules scoring high on the P-complexity may score low on the C-complexity). A follow-up challenge is being planned that may be shaped in part by the ideas proposed by the MLI team.

Computer-Assisted Fact Investigation

(Schum)

Information accumulates very rapidly in criminal investigation and in discovery-related activities in many other contexts. This project rests on the premise that the more skillful we are in marshalling or organizing information we do have the better are our chances of discovering what we do not have in the way of new possibilities and evidential tests of them. Methods of computer assistance in the process of discovery must provide more than archival capabilities if they are to be maximally useful. Such methods should also act in various ways to stimulate the imaginations of the persons who use them. Discovery, in common with so many other human tasks, is multiattribute. Part of our research has involved determining what are the essential attributes of discovery. This has provided a basis for identifying various forms of evidence marshalling attuned to these attributes. Since there are many intellectual requirements in the process of discovery, we did not imagine that any single scheme for marshalling thought and information would satisfy all of these requirements.

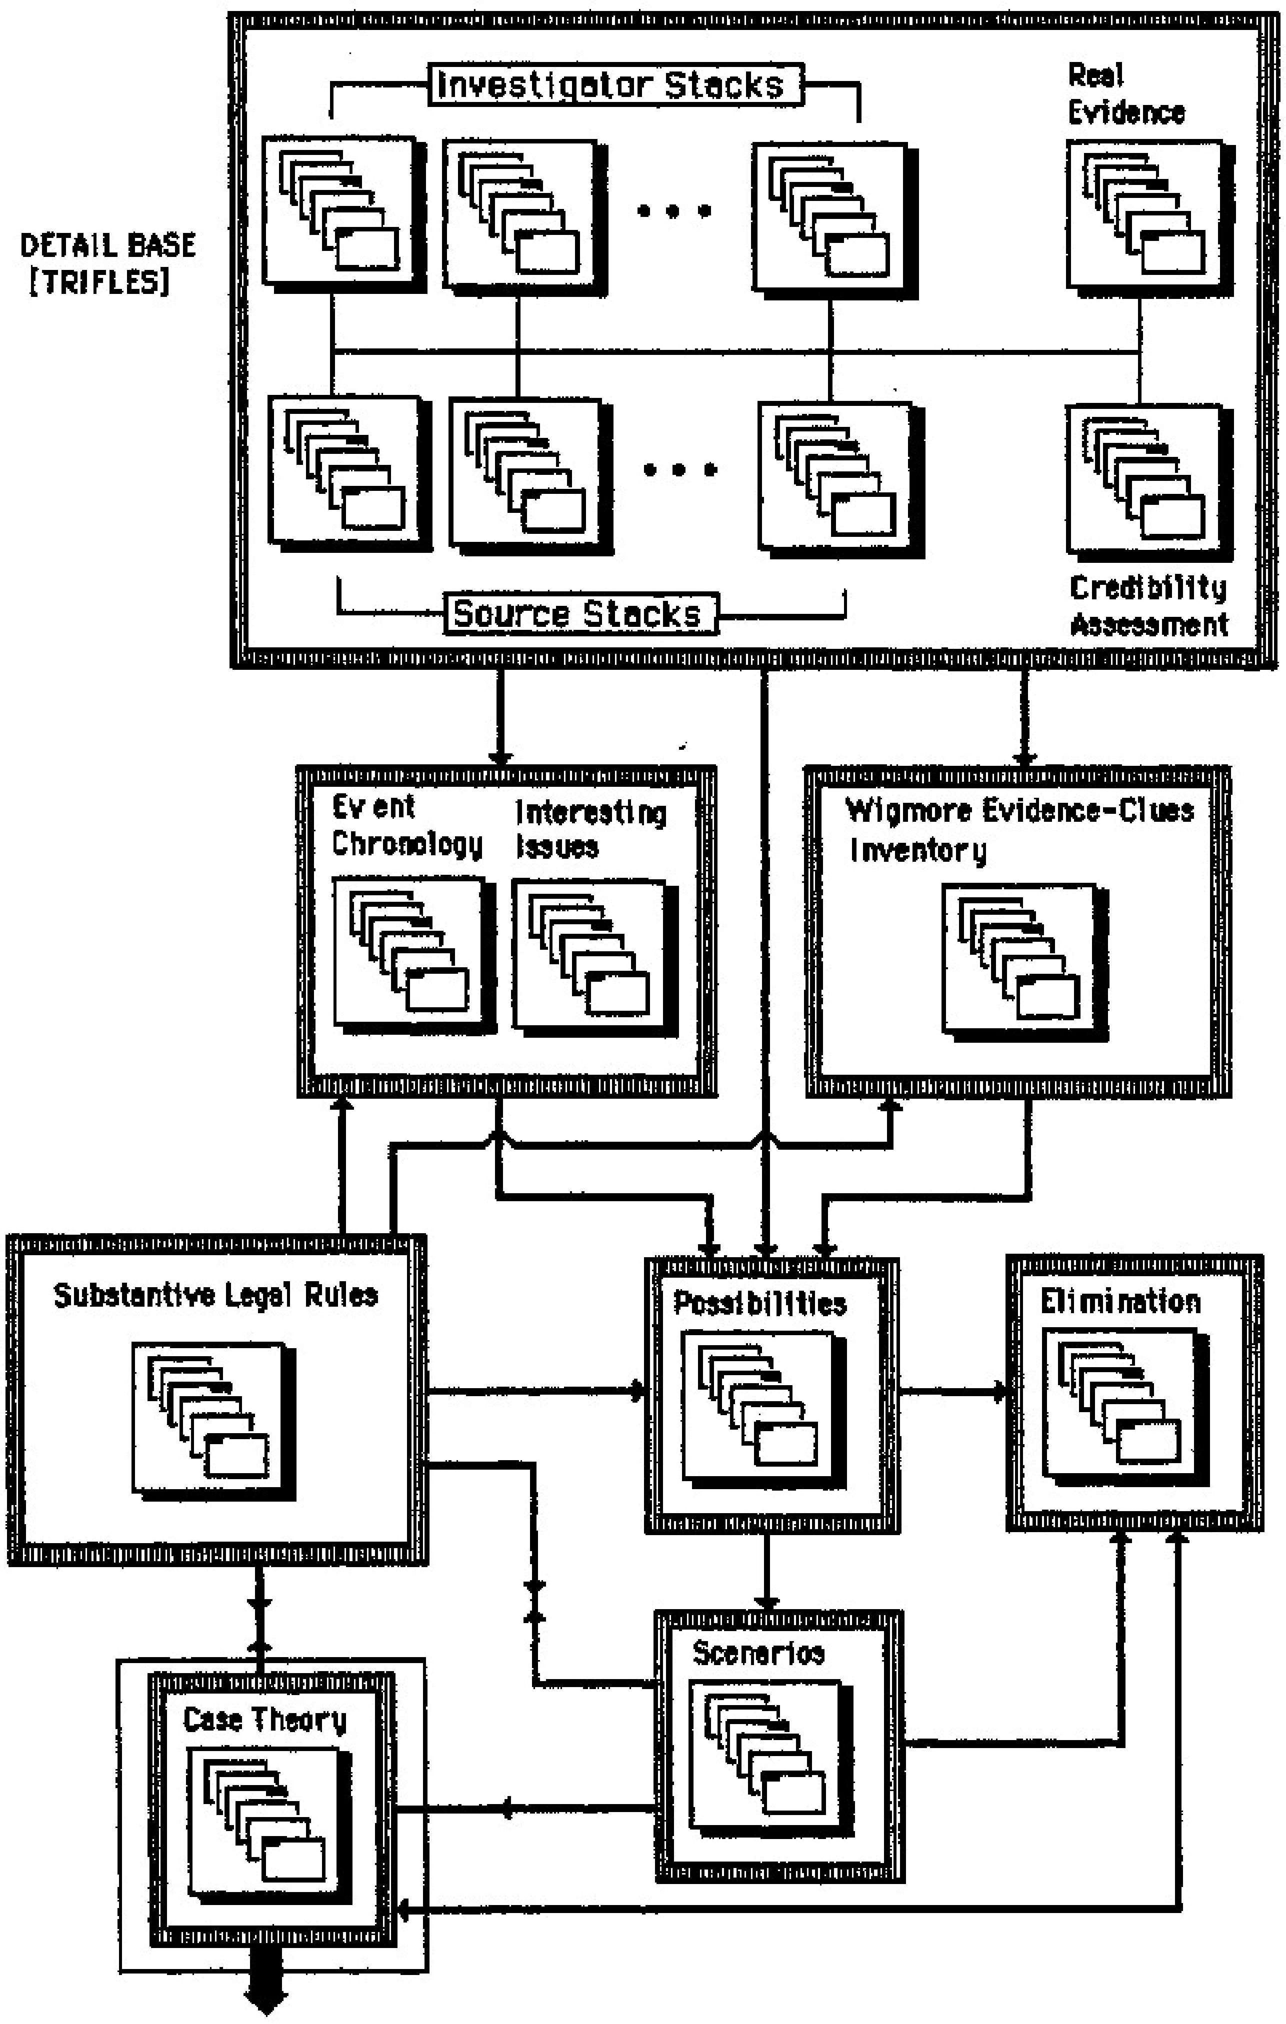

Shown in the figure is a network of evidence-marshalling operations we have identified on the basis of elaborate simulations of episodes of criminal investigation. Each of the nodes on this network corresponds to a particular form of evidence marshalling that seems required at one stage or the other in discovery-related activities. For each node we have developed prototype computer-based systems for organizing evidence in a particular way. Observe that the nodes are linked together in various ways. Our prototype system also incorporates these linkages so that the user can easily navigate from one part of the network to another. Several of these nodes have been designed to exploit ideas from alternative formal systems of probabilistic reasoning.

This research, now in progress at GMU, has involved the active collaboration of Professor Peter Tillers [Cardozo School of Law]. Support for this project has been provided by the National Science Foundation.

Learning Engineering

(Arciszewski, Michalski, Wnek, Doulamis, Szczepanik)

The ultimate objective of this project is to develop learning engineering - a new subarea of knowledge engineering which deals with the methodological aspects of using learning systems in knowledge acquisition. Learning engineering encompasses the evaluation and selection of learning systems, the methodology of automated knowledge acquisition, and the verification of knowledge. Its development is crucial for practical applications of machine learning which have been significantly delayed because of the present gap between the theoretical research on learning in computer science and the needs of potential users of learning systems. For the users, these systems are simply new knowledge acquisition and decision support tools which cannot be applied as long as their performance cannot be evaluated and methodology of their use is unknown.

Presently, research is concentrated on the performance-based evaluation of learning systems, and a method for such evaluation has been developed, and was presented in Arciszewski et al. (1992) in the Special Issue on Machine Learning of the Journal of Knowledge Acquisition "Heuristics." The method includes an evaluation procedure, an evaluation model based on the multiattribute utility theory, and a classification system for evaluation criteria in the form of various empirical error rates. Also, all basic methods assumptions were justified considering experimental results obtained for the actual engineering examples from the areas of structural design and construction safety.

Brain Builder Project

(de Garis)

This project is being undertaken by the Brain Builder Group of the Evolutionary Systems Department, Human Information Processing Division, of the ATR Labs, near Kyoto, Japan. The project leader is Dr. Hugo de Garis, a senior research affiliate of the MLIL who frequently interacts and visits the Laboratory. The ATR (Advanced Telecommunications Research) Labs are financed by the Japanese Ministry of Posts and Telecommunications, who made windfall profits in recent years by privatizing the national telephone company NTT. Some of these profits have been ploughed back into fundamental research. The Evolutionary Systems (ES) Department is the new "Artificial Life" branch of ATR, and aims, among other things to "build a brain" by the end of the decade. ATR salary levels are so competitive, that the establishment of this new department with some "gaijin" (foreign) superstars has been possible. Two of these gaijins have had articles written about them in Scientific American. The aim of the Brain Builder Group is to evolve neural circuits directly in hardware, using special devices called "Darwin Machines." These circuits are intended to control the behaviors of "biots" (biological robots). Initially these biots will have 10 behaviors, later 100 and more. The group will contain specialists in nanotechnology (molecular scale engineering) hoping to carry over the lessons learned from electronic Darwin Machines to molecular Darwin Machines, which can self assemble in embryonic ways to build sophisticated nervous systems. The ES Department consists of 15 people, half of whom are gaijins. It is possible that the personnel structure of this project may attract a lot of attention in Japan because it aims to combine the growing wealth and commitment of the Japanese government to fundamental research, with international talent and creativity. To have such a concentration of gaijins in one team is a new concept in Japan. ATR itself is only 3 years old. The ES Department was formed in 1992, collaborating with several other institutions, both governmental and private.

Pattern Representation of Knowledge

(Vamos)

Current machine learning methods can successfully solve problems that concern relatively well-structured, and conceptually well-defined knowledge. The uncertainty of knowledge is handled in these systems by various models and related methods of uncertainty calculations. These methods combine and propagate statistical estimates, and/or measures of human beliefs, and are based on several different hypothetical mechanisms.

The brain works in a different way – akin to the connectionist view it accepts and selects coherent clusters of information on a specific problem and connects them to other information entities stored by evolutionary and cognitive knowledge successfully used in the solution of the problem. The ongoing research mirrors this mechanism and is oriented to the representation and processing of unstructured or weakly structured knowledge by using selective learning and coherence of information as well as human supported discovery of analogical and discriminatory relations among information patterns. The learning process is more difficult than the usual visual pattern recognition, because the space of these knowledge patterns is non-metric. This feature is due mostly to non-represented tacit knowledge, the most valuable feature of highly skilled expertise. A man-machine system, using methods of pattern recognition is provided with a tuning feature, making it possible to include a relevant part of this expertise into the problem solver. The method was successfully used in the highly complex task of medical diagnosis related to brain development. Another experiment involved an analysis of a legal-sociological decision task.

The objective of this research is to combine the research of the MLI Laboratory in structuring knowledge and discovering relevant features with pattern-oriented methods, enabling the extraction of more subjective and hidden knowledge. This envisaged experimental field deals with the discovery of patterns in the world economy, in the context of new processes that are not yet sufficiently understood.

Integrated Knowledge-Based Systems for Automotive Design

(Arciszewski, Ardayfio, Michalski, Wnek)

The Chrysler Design for Manufacture and Assembly program at Chrysler Technology Center in Auburn Hills, Michigan, has combined resources with George Mason University and the National Science Foundation to initiate the proof of concept research in the application of machine learning to car design knowledge acquisition in the context of the development of Integrated Knowledge Based Systems for car design. The initial stage of the project involved the analysis of design configurations of automotive suspension crossmembers from the point of view of their manufacturability. Next, various car key design parameters are investigated to determine the relationships between them and the car's functional characteristics. The research is in an early stage, but its initial results are promising. The PI in this unique joint University-Industry project is Tomasz Arciszewski, and Dr. David Ardayfio, Senior Quality Executive at Chrysler's Technology Center is a Co-PI. Other members of the research team are R.S. Michalski, J. Wnek, and graduate students.

Constructive Induction in Engineering Design

(Arciszewski, Michalski, Wnek, Bloedorn)

The ultimate objective of this project is to develop a class of constructive induction methods for the applications to engineering design and a practical methodology for their use. A feasibility study has been completed and its results presented in the research report (Arciszewski et al 1992) published at the Machine Learning and Inference Laboratory at George Mason University and in the ASCE Journal of Computing in Civil Engineering (Arciszewski et al, P94-16). The study was conducted in the area of conceptual design of wind bracings in steel skeleton structures of tall buildings.

Design rules were learned from a collection of 336 examples of minimum weight (optimal) designs of wind bracings. Constructive induction was used to produce design rules which explain how design requirements can be optimally (in terms of minimum steel weight) satisfied through the proper selection of individual components of a wind bracing structural system. All examples were prepared under identical design assumptions for a three-bay skeleton structure of a tall building in cooperation with practicing structural designers. Actual minimum-weight designs were produced using SODA, a computer system for structural optimization, analysis, and design of steel structures. The design rules obtained were divided into four classes corresponding to the value of the decision attribute: recommendation, standard, avoidance and infeasibility rules.

Two types of constructive induction have been used in the study: data-driven and hypothesis-driven constructive induction. The performance of both learning systems was formally measured by two empirical error rates: 1. the overall empirical error rate, 2. the omission error rate in accordance to the method of evaluation of performance of learning systems developed at the Laboratory and published in Arciszewski et al (1994). These error rates were calculated for the entire collection of examples using the leave-one-out resampling method. The error rates for constructive induction were compared with rates for the "traditional" induction, based on the use of the AQ15 algorithm. The individual error rates are shown in the table below. There is a significant improvement in performance (more than 50%) between the system based on the "traditional" induction and systems based on constructive induction. The difference in performance between two constructive induction-based systems is insignificant (less than 5%), but this may change as the research progresses.

Recently, the collection of examples was extended with additional 50 examples. The experiments are now being conducted using 386 examples.

The continuation of the project is supported by the National Science Foundation.

		CONSTRUCTIVE INDUCTION	
	AQ15 (Mini-Mode)	Data-Driven	Hypothesis-Driven
Overall Error Rate (%)	4.5	3.3	3.3
Omission Error Rate (%)	4.2	3.2	3.1

Constructive Induction Approach to Growing Neural Networks

(Sazonow, Wnek)

With most symbolic machine learning methods, if the given knowledge representation space is inadequate then the learning process will fail. This is also true with neural networks learning based methods. To overcome this limitation, a method for automatically "growing" neural network is being developed.

The BP-HCI method is a hypothesis-driven constructive induction for neural networks trained by the backpropagation algorithm. The method determines topology of a neural network and the initial connection weights based on patterns in the behavior of the neural network. The behavior of the neural network is captured by concepts called ACCORD and ANXIETY of a neural network.

The method was successfully applied to ten problems including such problems as learning the "exclusive-or" function, MONK2, parity-6BIT and inverse parity-6BIT.

This project is supported by the Machine Learning and Inference Laboratory.

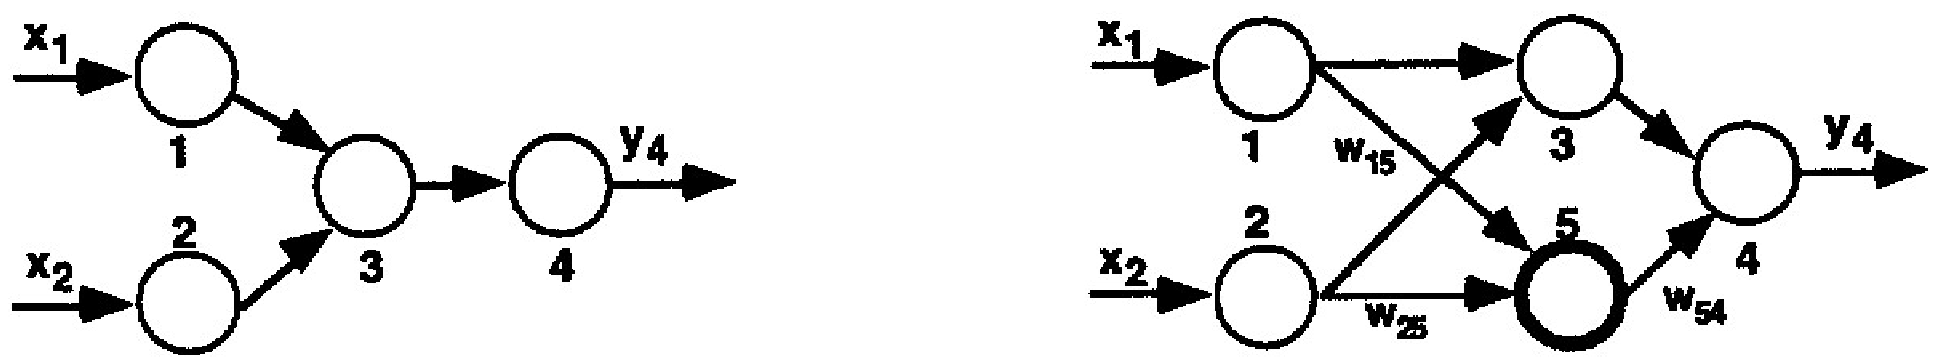

The figures above illustrate the process of growing a neural network. An additional unit is added according to a specially developed algorithm.

Noise Tolerant Learning

(Michalski, Bala, Wnek, Bloedorn)

The focus of this project is to develop highly noise tolerant learning methods and tools, and to apply them to a variety of practical knowledge acquisition and engineering tasks. Our aim is to make learning processes less sensitive to both, the attribute noise and the misclassification of the input data. The developed methods use different principles, such as:

(1) the optimization of learned concept descriptions by truncation or modification of their less significant components, specifically, the "light" rules (the TRUNC method); or the "light" rules, conditions, and references (the TRUNC/SG method implemented in the POSEIDON system);

(2) the application of constructive induction methods that generate various new "problem-oriented" attributes from the given attributes.

In our research we have been concentrating on methods for learning complex concepts (e.g., acceptable labor management contracts), and noise tolerant learning in computer vision. The investigation and testing of these methods is done using images of textured objects. The results obtained show a very high level of noise tolerance of our programs.

The support for this project is provided by grants from ARPA and ONR.

Machine Learning in Bridge Design (MLBD)

(Chen, Arciszewski, Michalski, Wnek, Kaufman)

The long-term goal of this project is to study the applicability of machine learning to large-scale, complex engineering design problems and the development of a general methodology for such application.

Conventional machine learning techniques are "diagnosis-oriented" rather than "design-oriented" in nature. In "diagnosis-oriented" tasks, such as the diagnosis of the diseases of living systems, the assessment of malfunction of artifacts, and the evaluation of the design solutions (instead of generation of these design solutions), the ratio of the number of decision variables to independent variables is usually small, the interdependency between decision variables usually viewed as unimportant, and the number of available examples for training and testing can be quite large. In "design-oriented" tasks, however, these task characteristics partially or totally disappear, dependent on the specific engineering problem. This fact results in great difficulties as well as new opportunities in both, the methodological and application aspects of machine learning research.

The MLBD project consist of two stages: 1) the feasibility study of the application of machine learning to complex problems of engineering design 2) the development of a methodology for such an application.

We chose the preliminary design of cable-stayed bridges as our testing domain not only because the cable-stayed bridge is one of the most complex and representative engineering structures, but also because we have made great effort in the development of a knowledge-based system for that purpose. This system facilitates the definition of the description space, the preparation of training and testing examples, the evaluation of the quality of design rules obtained from machine learning systems, and the engineering interpretation of the generated knowledge. Experiments are conducted using AQ15c and AQ17 learning systems in order to answer the following questions:

A. How to transfer or re-represent a "design-oriented" learning task into a set of or a sequence of "diagnosis-oriented" learning tasks which can be coped with by using current machine learning techniques ?

B. How to achieve enough accuracy of learning processes with limited number of available training examples, which is typical in civil engineering structure designs such as the long-span bridges, the high-rise buildings, and the hydraulic facilities over great rivers ?

C. How to incorporate symbolic heuristic design rule learning and numerical experience-based design formula learning in a learning task (two of the most common knowledge representations used by human design experts) ?

D. How to tune up the "design-oriented" learning task by using the constructive induction ?

An important part of this research is an evaluation of the quality of design rules obtained from machine learning systems. This process utilizes two alternative and complementary approaches:

a) Knowledge verification by using various empirical error rates;
b) Knowledge comparison by using the knowledge base of design rules acquired manually.

This project is partially funded by the Chinese Academy of Sciences, the National Science Foundation and the Machine Learning and Inference Laboratory.

Conceptual Design of Production and Operations Systems

(Karni, Arciszewski, Michalski)

The aim of manufacturing engineering is to provide the right product to the right customer at the right quality, quantity, price, time and location. Thus engineering design has as its purpose to detail a product and concurrently to develop capabilities for production, quality and field support. To this end, three major categories of engineering systems can be distinguished:

(a) hardware (product-related) systems - the product and its manufacturing processes;

(b) human activity (operations-related) systems - the production, logistics and field support regimes;

(c) management (information-related) systems - the information and decision support environment for management and control.

These three categories are usually associated with product engineering, industrial/systems engineering, and information/systems engineering respectively.

Extensive design theories and procedures have been developed for the first and third categories. Almost nothing has been suggested for operations-related systems design. Compared to other engineering disciplines, the design methodology for industrial and systems engineering is relatively primitive and nonrigorous. Much effort needs to be directed toward the development of better design procedures. One way in which this may be done is based upon the Inferential Theory of Learning (ITL). This provides the design with a goal-oriented process which improves his knowledge by explaining his experience in a systematic manner. The Theory analyzes the learning process in terms of high level inference patterns called knowledge transmutations. Analogously, we provide the design with a set of *design knowledge transmutations*, which guide the evolving design.

We thus envisage the development of a computer-based tool for operations-related systems design, based on man-machine interaction. The human deals with the relevance of the problem characteristics, the meaningfulness of the resultant evolving design concepts, and the introduction of new or modified design knowledge. These are derived from cognitive interaction with the process, or associations stimulated by conjunctions of characteristics or properties suggested by the process. The machine deals with problem and concept manipulation for prompting the human during knowledge search and concept creation. Our major objectives are:

(a) to formulate a methodological foundation for the conceptual design of operations-related systems, based on the use of a set of generic *design knowledge transmutations* developed at the MLI Laboratory;

(b) to formulate a methodological foundation for designer/system interfacing and interacting, in order to operationalize the conceptual design process;

(c) to develop a prototypical design tool for implementing and testing the proposed methodologies.

1a. Generalization	1b. Specialization	9a. Insertion	9b. Deletion
2a. Abstraction	2b. Concretion	10a. Replication	10b. Destruction
3a. Similization	3b. Dissimilization	11a. Sorting	11b. Unsorting
4a. Association	4b. Disassociation	12a. Retention	12b. Extension
5a. Selection	5b. Generation	13a. Framing	13b. Reframing
6a. Agglomeration	6b. Decomposition	14a. Implosion	14b. Explosion
7a. Characterization	7b. Discrimination	15a. Symmetry	15b. Anti-symmetry
8a. Reformulation	8b. Randomization	16a. Explanation	16b. Prediction

Design Knowledge Transmutations

FACULTY & ASSOCIATES

Tomasz Arciszewski is Associate Professor of Urban Systems Engineering in the Systems Engineering Department, George Mason University. He received his M.S. (Summa Cum Laude) and Ph.D. degrees from the Warsaw University of Technology in 1970 and 1975, respectively. He taught at Wayne State University in Detroit, Michigan, at the Warsaw University of Technology, and at the University of Nigeria in Nsukka. Dr. Arciszewski gained practical experience in engineering design in Poland and Switzerland. His research interests include engineering design, and applications of machine learning to various areas of engineering, including environmental, construction, transportation and structural engineering. He has also been working on "learning engineering," a new subarea of knowledge engineering dealing with the methodological aspects of using learning systems for knowledge acquisition. His research has been supported by several grants from the National Science Foundation and the State of Michigan. He has authored or co-authored over sixty publications in the areas of structural engineering, design methodology and artificial intelligence, including several book chapters. His research on innovative design led to the development of a conceptual design method which was implemented in a computer program and its use produced three inventions in the area of structural engineering which were patented in Canada, Poland, and the USA. He is one of the two Editors of the recently published monograph "Knowledge Acquisition in Civil Engineering" and the Technical Editor of the American Society of Civil Engineers Journal "Computing in Civil Engineering." Presently, he is the Vice Chair of the Expert Systems and Artificial Intelligence Committee of the same society. At the MLI, he has been collaborating with R.S. Michalski and several other researchers in the areas of learning engineering and constructive induction in engineering design (research sponsored by the National Science Foundation and Chrysler Corporation).

Jerzy W. Bala is Research Affiliate at the MLI Laboratory. He received his M.S. degree in Electrical and Computer Engineering in 1985 from the AGH Polytechnic University, Cracow, Poland, and Ph.D. in Information Technology from George Mason University in May, 1993. Prior to immigrating to the United States he worked in a digital system and microprocessor software laboratory of a leading microcomputer plant. In 1987 he was invited by Dr. Michalski, at that time the director of the Artificial Intelligence Laboratory at the University of Illinois at Champaign-Urbana, to conduct research in machine learning. His stay at the UIUC was supported by a grant from the Kosciuszko Foundation. In January 1988 he moved with Dr. Michalski's machine learning group to George Mason University, where he was pursuing his Ph.D. studies in the MLI Laboratory. His dissertation dealt with the development of a symbolic machine learning system that learns descriptions of visual concepts from low-level vision data. This research was in the novel and rapidly developing area of Learning in Vision, which investigates problems of incorporating learning capabilities in computer vision systems. His other research interests are in the following areas: robust learning, integration of symbolic and numerical learning methodologies, genetic algorithms, and pattern recognition. His research has led to over 20 publications. He presented his research at important conferences in his field, such as International Conferences on Machine Learning, International Symposium on Methodologies for Intelligence Systems ISMIS-91, IEEE Conference on AI Applications 1991, First International Multistrategy Workshop MSL-91, Tools for Artificial Intelligence TAI-90, TAI-91 and TAI-92. He is a member of the ACM, AAAI, and IEEE.

Quin Chen is Professor at the Research Institute of Engineering Mechanics in Dalian University of

Technology of P.R.China and Visiting Research Associate Professor in the GMU Systems Engineering Department. He received his B.S. in Engineering Mechanics in 1982, M.S. in 1984 and Ph.D. in Computational Mechanics from the Dalian University of Technology in P.R. China in 1988. He has been General Secretary of the Chinese Association of Computational Mechanics since 1990. His biography was included in MAN OF ACHIEVEMENT, 16th Edition of International Biographical Center in England, in 1994. Dr. Chen presented a computational model of knowledge-based engineering design process; developed a knowledge-based preliminary design system for practical design purposes of cable-stayed bridges; and formulated an optimal control model for construction of the flexible engineering structures. Dr. Chen's research interests include Information-Driven Engineering Systems (IDES), such as knowledge-based expert systems, neural networks, fuzzy set theory approaches, case-based reasoning, systems engineering, operations research, optimal control, mathematical programming, software engineering, intelligent CAD, and recently machine learning; as well as many IDES aided engineering applications, such as structural optimization; nonlinear structural vibration; optimal control of structural assembling process; preliminary design of large-scale complex structures; safety assessment and diagnosing of the pressure vessels in service; reliability-based optimization of structures subjected to disaster loading conditions; and recently automatic knowledge acquisition for bridge design. Dr. Chen is the author of the monograph "AI Aided Decision-making in Engineering" and one of the authors of the "Handbook of Mechanical Engineering" published by Chinese Society of Mechanical Engineers.

Hugo de Garis has been a Senior Research Affiliate at the MLI Laboratory since 1990, after having spent six months at the Laboratory as a graduate student in 1989. He obtained his Ph.D. at Brussels University in January 1992 with a thesis entitled "Genetic Programming: GenNets, Artificial Nervous Systems, and Artificial Embryos". In 1992, Dr. de Garis was a Science and Technology Agency (STA) postdoctoral fellow at the Electrotechnical Laboratory in Tsukuba Science City, in Japan. In February 1993, he moved to his current position as senior scientist at the new Evolutionary Systems Department that he helped conceive and establish at the ATR Labs near Kyoto, Japan. This new department aims to "build a brain", i.e. artificial nervous systems for "biots" (biological robots), by evolving neural circuits directly in hardware using special devices called "Darwin Machines." Hugo's research interests focus on the idea that evolutionary algorithms (e.g. Genetic Algorithms) can be used to build/evolve complex systems whose dynamics and structures are too complex to analyze but are nevertheless functional. He has published more than 30 papers in the fields of Neural Networks, Genetic Algorithms, Artificial Life, Machine Learning, Data Analysis, and Symbolic Learning. His current interests are concerned with Artificial Life, Nano-technology, and VLSI for Darwin Machine design.

Zoran Duric is Assistant Research Scientist at the Machine Learning and Inference Laboratory and at the Center for Automation Research at the University of Maryland. He received his M.S. in Electrical Engineering from the University of Sarajevo, Bosnia and Herzegovina, Yugoslavia, in 1986, and his Ph.D. in Computer Science from the University of Maryland at College Park, in 1995. From 1982 to 1989 he was a Member of the Research Staff at the Energoinvest, Institute for Control and Computer Science, Division for Vision and Robotics, Sarajevo. He was also affiliated with the Electrical Engineering Department of the University of Sarajevo. From 1990 to 1995 he was a Graduate Research Assistant at the Center for Automation Research at the University of Maryland at College Park. In the summer of 1993 he was a Summer Research Assistant at the NEC Research Institute, Princeton, New Jersey. Dr. Duric's recent major achievements include a method for image sequence stabilization using normal flow and a method for analyzing the functionality of hand tools from motion sequences using normal flow and weak perspective projection. He is a member of the IEEE Computer Society.

Ophir Frieder is Professor of Computer Science at George Mason University. He received his B.Sc. in Computer and Communications Science in 1984, M.Sc. in 1985 and Ph.D. in 1987 in Computer Science and Engineering, all from the University of Michigan. From 1987 to 1990, Dr. Frieder was a Member of Technical Staff in the Applied Research Area of Bell Communications Research. His research focussed on the design and development of parallel and distributed systems aimed at improving modern-day telecommunication systems. He joined George Mason University in 1990. While at GMU, Dr. Frieder has also served as a staff consultant at the Federal Bureau of Investigations (1991-1993) and at the Institute for Defense Analysis (1992-1993). Since 1993, Dr. Frieder has been a staff consultant to IBM FSC (now Loral Federal Systems). Dr. Frieder has published over 60 refereed publications, was granted two patents, and has received research support from several government and industrial grants including four from the NSF. In 1993, he was a recipient of the International Information Science Foundation Award from Japan and the NSF National Young Investigator Award. Dr. Frieder's research interests include parallel and distributed database and information retrieval systems and biological and medical data processing architectures. He is the Area Editor for Databases of IEEE Computer. Dr. Frieder is a member of Phi Beta Kappa and a Senior Member of the IEEE.

Ibrahim F. Imam is Research Affiliate at the GMU Machine Learning and Inference Laboratory. He received his B.Sc. in Mathematical Statistics from the Department of Mathematics, Cairo University, Egypt, in 1986. He received a Graduate Diploma in Computer Science and Information from the Department of Computer Science, Cairo University, Egypt, in 1989. He received his M.S. in Computer Science from the Department of Computer Science, George Mason University, Fairfax, Virginia, in 1992. He received his Ph.D. in Information Technology and Engineering from George Mason University, in Fall 1995. Ibrahim is the founder and co-chair of the first FLAIRS-95 international workshop on Intelligent Adaptive Systems (IAS-95), Melbourne Beach, Florida, 1995. He served on the organizing committee of the Florida Artificial Intelligence Research Symposium, FLAIRS-95; the program committee of Florida Artificial Intelligence Research Symposium, FLAIRS-96. He was also involved or volunteered in organizing many other conferences and workshops including: the AAAI-93 and AAAI-94 conferences; the first and second Workshops on Multistrategy Learning (MSL-91 & MSL-93); and the IEA/AIE-94 conference. Ibrahim's dissertation, titled "Deriving Task-oriented Decision Structures from Decision Rules" was supervised by Professor Ryszard S. Michalski, and supported by the Laboratory. Ibrahim's research interests focus in the areas of machine learning, intelligent agent and adaptive systems, knowledge discovery in databases, hybrid classification, and knowledge-base systems.

Kenneth A. De Jong is Associate Professor of Computer Science at George Mason University. He received his Ph.D. degree in Computer Science from the University of Michigan in 1975. Since then he has been a faculty member in the Department of Computer Science at the University of Pittsburgh and a visiting scientist at the Navy Center for Applied Research in Artificial Intelligence. Dr. De Jong's research interests include adaptive systems, machine learning, expert systems, and knowledge representation. He is an active member of the Genetic Algorithms research community with a variety of papers, Ph.D. students, and presentations in this area. He is also responsible for many of the workshops and conferences on Genetic Algorithms. He is currently chair of the executive council of the Society of Genetic Algorithms and Editor-in-chief of the journal, Evolutionary Computation. Dr. De Jong is involved in directing a number of research projects involving the use of Genetic Algorithms on NP-hard problems and the use of Genetic Algorithms to learn task programs in domains such as robotics, diagnosis, navigation, and texture recognition. He is also interested in theoretical properties of genetic and other evolutionary algorithms, including convergence properties, non-random selection methods, the evolution of structure, and self-adaptive evolutionary systems.

Reuven Karni is an Associate Professor in the Faculty of Industrial Engineering and Management at the Technion - Israel Institute of Technology - located in Haifa, Israel; and a Visiting Research Associate Professor at the GMU MLI Laboratory. He received his B.Sc. (Eng). and M.Sc. (Eng.) in Chemical Engineering at the University of the Witwatersrand, Johannesburg, South Africa; and his D.Sc. in Industrial Engineering at the Technion (1978). Before joining the Technion in 1978, Dr. Karni worked for ten years as a senior consulting chemical and systems engineer at the Industries Development Corporation, and for five years as a systems engineer at IBM (Israel). His research focusses on empowering the industrial engineer within the modern engineering environment, particularly in the applications of artificial intelligence in the fields of industrial engineering design, planning and management. His particular interest is in developing a methodology for the conceptual design of production and operations systems, based on inferential learning.

Larry Kerschberg is Professor and Chairman of the Department of Information and Software Systems Engineering. He holds Ph.D. in Systems Engineering from Case Western Reserve University, M.Sc. in Electrical Engineering from the University of Wisconsin-Madison, and B.Sc. degree in Engineering Science from Case Institute of Technology. His research is in the areas of data and knowledge models, database design, active data dictionaries, distributed query processing, object-oriented systems, knowledge discovery in databases and expert database systems. His current research projects include Information Integration and Interchange: A Federated Systems Approach and Knowledge Discovery in Databases, both sponsored by ARPA, Sustaining Engineering-Life Cycle Support for Evolutionary Software Development, sponsored by NASA. Dr. Kerschberg serves as an Editor-in-Chief of the International Journal of Intelligent Information Systems. He served as General Chair of the 1993 ACM SIGMOD Conference held in Washington, DC. He organized and has served as Program Chairman of both the 1st and 2nd International Conferences on Expert Database Systems. He is past Chairman of the IEEE Computer Society's Technical Committee on Data Engineering.

Yves Kodratoff is Research Director at the University of Paris-South and Adjunct Professor at the GMU MLI Laboratory. He received his M.S. in Chemistry, Physics and Applied Mathematics, and his Ph.D. (1967) from the Institute of Research on Catalysis, at the University of Lyon, France. Dr. Kodratoff is one of the most distinguished researchers in AI and CS in Europe. His work in automatic programming and more recently in machine learning is widely known. He has been Chairman of "Automatic Programming" at the European Congress on AI, Chairman of the Program Committee of ECAI-88, invited speaker at numerous machine learning conferences, and a member of the editorial board of Machine Learning and Expert Systems Journals. Dr. Kodratoff was also the head of Machine Learning and Knowledge Acquisition research project in Europe, and Vice-President of the European Coordination Committee for AI. He has authored or co-authored over 40 publications and thirteen textbooks in automatic programming and AI. His research group is building an AI tutoring system using machine learning techniques. His domains of research interests span computer science, chemistry and psychology. By holding joint appointments at GMU and the University of Paris he is building technical brigdes and contributing to mutual collaboration between the two universities.

Nikolai Lyashenko is Vice President of Empirical Inference Corporation, Schenectady, NY, and a Visiting Professor of the Laboratory since January 1993. He received his M.S., Ph.D. and Dr. of Math. and Phys. Sci. degrees at Leningrad University, USSR. He was Professor at Leningrad Polytechnical Institute and Director of Information Processing Laboratory at CS Institute of the Soviet Academy of Science. He began his research activities in Probability Theory and Data Analysis. He is the inventor of The Empirical Inference Technology, which provides efficient tools of model-free data analysis. Dr. Lyashenko participated in twelve large applied projects in the USSR, including the system for diagnostics of bronchial asthma (under the supervision of the World Health Organization), the environmental monitoring system in the area of the Angara river, the design of the security circuit at the Leningrad Nuclear Power Plant. In the US and Canada he has performed three industrial projects: "Production Line Improvement" (for GE), "Analysis of Plutonium Ingress in Reactor Pressure Tubes" (for Atomic Energy Control Board, Ottawa) and "Magnetographic Diagnostics of Myocardial Infarct and Ventricular Tachycardia" (for GE). He has published over 80 research papers in Probability Theory, Data Analysis, Algorithm Theory and System Theory.

George S. Michaels is Associate Professor of Computational Biology and Deputy Director of the Institute for Computational Sciences and Informatics. He received his Ph.D. in Biochemistry and Molecular Biology from the University of Florida College of Medicine, in 1985. He recently came to GMU from the NIH. His major research interests are in the areas of molecular biology and bioinformatics. His research interests include modeling of DNA/protein interactions of transcriptional regulatory factors, protein design, protein folding, advanced computational methods for genome analysis, metabolic modeling and organism systematics for remediation. This work involves the application of advanced computational and AI methods for pattern analysis, linguistic approaches for analysis of molecular sequence data, molecular simulation and visualization methods for biologically active macromolecules, development of prototype databases optimized for flexible query methods; distributed databases and collaborative internet research environments; and development of complex data visualization systems. He directs the Bioinformatics Laboratory whose goal is to develop a new collection of advanced computational tools to analyze, archive, and curate biological data.

Ryszard S. Michalski is Planning Research Corporation Chaired Professor of Computer Science and Systems Engineering, and Director of the Center. He received his B.S. (equivalent) from the Warsaw Polytechnic University, M.S. from the St. Petersburg Polytechnical University, and his Ph.D. from the University of Silesia in Poland in 1969. Before emigrating to the United States in 1970, he was a Research Scientist at the Polish Academy of Sciences. From 1970 to 1987 he was a Professor of Computer Science and Medical Information Science, and more recently Director of the Artificial Intelligence Laboratory at the University of Illinois at Urbana-Champaign. He moved with his research group to GMU in January 1988. Dr. Michalski is a co-founder of the field of machine learning and a pioneer in a number of research areas in artificial intelligence and cognitive science. He originated research on constructive induction and conceptual clustering; developed a computational theory of inductive learning; invented variable-valued logic; introduced the concept of variable precision logic (with Patrick Winston from MIT); and developed a computational theory of human plausible reasoning (with Allan Collins from Cambridge, MA). Collaborating with a plant pathologist at the University of Illinois, he developed the first agricultural expert system, and the first expert system that learned its rule-base from examples. Dr. Michalski is co-founder of the Journal of Machine Learning, and a co-organizer of the first several international machine learning conferences. He has lectured extensively worldwide, and held visiting professorships at major universities in the U.S., including MIT, CMU and the University of Wisconsin, as well as abroad, specifically, in Belgium, Great Britain, Italy and France. Dr. Michalski's research interests include machine learning and inference, cognitive modeling, applications of machine learning to computer vision, and autonomous intelligent robots. He co-edited the series of books *Machine Learning* which received world-acclaim, and has published over 250 publications in his areas of interest.

David C. Rine is Professor of Computer Science, and Information and Software Systems Engineering. He received his Ph.D. in 1970 from The University of Iowa. He has been in the computing field for thirty years, working in university teaching and research, consulting, and developing software systems for businesses, industry, and government agencies. He joined GMU in 1985 and was the founding chair of the CS Dept., as well as a co-developer of the School of Information Technology and Engineering. He presently guides masters and doctoral level students in the area of software engineering.

He also works with students in the Institute for Computational Sciences and Informatics, as well as with students in the MLI Laboratory. In software engineering he has been specializing in research and teaching in the areas of object-oriented development, metrics, software maintenance and reuse, and software development environments, and their applications. He has also maintained an interest in applications of logics, such as multiple-valued logics and fuzzy logics, to software development. Within the span of his career in computing he has published over 150 papers in the areas of theoretical computer science, logic design, information systems, computer applications, computer education, and software engineering. Dr. Rine has received numerous awards from computer science societies and associations including the IEEE Centennial Award, The IEEE Computer Society Pioneer Award, The IEEE Computer Society Meritorious Service Award, and the IEEE Computer Society Special Award. He has also been a multiple-time recipient of the IEEE Computer Society Honor Roll Award and the IEEE Computer Society Certificate of Appreciation Award. He has chaired a number of computer science and engineering conferences, including those in software development and multiple-valued logic. He has also co-authored a number of computer oriented texts and books, including those in multiple-valued logic, fundamentals of computer science, information systems, and object-oriented systems and applications. He has guest-edited a number of technical journals and is/has been on the editorial boards of several journals, including COMPUTER of the IEEE Computer Society. In the area of fuzzy logics, Dr. Rine has published and written papers dealing with foundations of fuzzy logic and their applications to databases, expert systems, security, object-oriented design, adaptive systems, as well as logic design.

Vladimir N. Sazonov is Senior Member of the Technical Staff of Computer Science Corporation, and a member of GMU Mathematical Science and Computer Science Departments. He received his Ph.D. in 1985. He worked in areas of transfer of polarized radiation from 1968 to 1975, interaction of laser radiation with polyatomic molecules and condensed matter from 1976 to 1986 and artificial intelligence from 1987. His research interests include industrial systems, inductive logic, self expanding neural networks. Before immigrating to the U.S.A. in 1992 he was Professor of Mathematics and Chairman of the Department of Higher Mathematics at Moscow Machine Tool Institute. Dr. Sazonow published ver 90 publications in his areas of interest. He is a member of AAAI and IEEE Computer Society.

David A. Schum is Professor in the Department of Operations Research and Engineering. He received his Ph.D. in Psychology from Ohio State University in 1964. Following a two-year postdoctoral appointment at Ohio State, he joined the faculty of Rice University, Houston, Texas in 1966. He remained at Rice until 1985, holding the rank of Professor in the Department of Psychology and of Mathematical Sciences. He was also a member of the adjunct faculty at Baylor College of Medicine. In 1985, he joined the GMU faculty. His major research interests have involved both formal and empirical studies of the process of drawing conclusions from masses of evidence. Other interests include studies of the process of discovery as it concerns the generation of new possibilities and evidential tests of them. He is a co-author of probability text, author of a work on evidence and inference, and has published in a variety of behavioral, legal and engineering journals.

Tibor Vamos is Professor of the Technical University of Budapest and Distinguished Visiting Professor of the Center for the 1992 Fall semester. He is Chairman of the Board of the Computer and Automation Institute of the Hungarian Academy and Head of the AI Laboratory of the same institution. He received his M.S., Ph.D. and Dr.Sc. degrees in Budapest, is a Full Member of the Hungarian Academy of Science, a Fellow of IEEE, and Honorary President of the John v. Neumann Computer Science Society. He was the President of the International Federation of Automatic Control (IFAC) for the 1981-84 term and is Life Time Advisor of the IFAC, and is an editorial board member of 7 international periodicals. Dr. Vamos began his research activities with computer process control, followed by visual pattern recognition for manufacturing and robotics. He is now involved in AI problems related to the pattern concept and epistemic issues of computer science. He has published over 200 papers. His last book, *Computer Epistemology,* was published by World Scientific, and translated into Italian. His investigations were pioneering in uncertainty problems of process control. Later he, with his group, developed a new fast algorithm for edge detection and fast on-line algorithms for vision controlled industrial robots. Dr. Vamos has also created a new expert system based on pattern principles which actively helps the early diagnostics and rehabilitation of children born with central neural system defects. In his book, he analyzed uncertainty and logic-based AI methods from the epistemic point of view. His current research focuses on the connection between two-tier knowledge extraction methods with the pattern representation concept of knowledge.

Lev M. Vekker is a Consultant in Information Science Technology at the BDM Corporation and Adjunct Research Professor of the MLI Laboratory. He received his M.S. in Physics from the University of St. Petersburg, U.S.S.R, in 1940, and his M.S. and Ph.D. in Psychology both from the University of St. Petersburg in 1947 and 1951, respectively. He also received a Post-doctoral Degree in Psychology from the University of St. Petersburg in 1964. Prior to immigrating to the U.S., he was, for many years, a Professor of Psychology at the University of St. Petersburg. He was also a Professor of General Psychology at Leipzig University in Germany, and Professor of Psychology at the Institute of Technology in Novgorod, U.S.S.R. His current research interests are in the areas of cognitive psychology, robotics, linguistics and philosophy of science. He has written over 100 publications, a book on *Perception and Its Technical Modeling*, and a three-volume monograph, *Psychological Processes: Sensation and Perception (I), Intellectual Functions and Thought Processes (II), Emotions, Actions and Personality (III).*

Janusz Wnek is Assistant Research Professor and Assistant Director for Research Management at the MLI Laboratory. He received his M.S. in Computer Science from the Jagiellonian University, Cracow, Poland, in 1983, and his Ph.D. in Information Technology from George Mason University, in 1993. Before emigrating to the United States in 1988, he was Director of a Software Engineering Laboratory at the KFAP company, where he developed and patented a variety of unique software. Dr. Wnek conducted his dissertation research in the Machine Learning and Inference Laboratory from 1989 to 1993. His recent major achievements include a novel method for constructive induction (HCI), and the diagrammatic visualization system (DIAV) that supports analysis and development of machine learning methods. He presented his research at some of the most important conferences in his field, such as the International Conferences on Machine Learning 1991 and 1992, the International Joint Conferences on Artificial Intelligence in 1991, 1993, and 1995, the International Workshops on Multistrategy Learning in 1991 and 1993, and the International Symposium on Methodologies for Intelligent Systems 1990. He was an active contributor to the organization of the First and the Second International Workshops on Multistrategy Learning. He was an instructor at the International Summer School on Machine Learning at Urbino, Italy. He published three chapters in the books Machine Learning: An Artificial Intelligence Approach, Vol. III and Machine Learning: A Multistrategy Approach, Vol. IV, and has also published ten journal, and conference papers. He is a member of the AAAI. His interests include artificial intelligence, machine learning, and the application of machine learning methods to engineering and economics.

MLI Laboratory Administration and Technical Staff

Ryszard S. Michalski, Director
Janusz Wnek, Assistant Director, Research
Abhay Kasera, Budget Manager
Eric Bloedorn, Software Manager
Naval Deshbandhu, Assistant Systems Manager
Ken Kaufman, Facilities Manager
Mark Maloof, Colloquium Chair
Dig Singh, Technical Assistant

Visiting Professors and Affiliate Scientists

Mahdi Abdelguerfi (Algeria; Summer 90)
Giuliano Armano (Italy; 89)
Francesco Bergadano (Italy; Spring 88)
Claudio Carpineto (Italy; Fall 88 - Sum 89)
Quin Chen, (China; Spring 94 - Spring 95)
Tom Fermanian (Univ. of Illinois; Spring 88)
Hugo de Garis (Belgium; Spring 89 - 91)
Rei Hamakawa (Japan; Fall 88 - Fall 89)
Hesham Ahmed Hassan (Egypt; Spring 90)
Reuven Karni (Israel; Fall 95)
Heedong Ko (Korea; 88-90)
Nada Lavrac (Yugoslavia; Fall 88)
Stan Matwin (Canada; Spring 88)
Soliman A. Mohamed (Egypt; Spring 90)
Igor Mozetic (Yugoslavia; 88 - Spring 89)
Daniele Nardi (Italy; Fall 95)
Tadeusz Piotrowski (Poland; Spring 90)
Gheorghe Tecuci (Romania; Spring 90)
Anke Rieger (Germany; Fall 90 - Summer 91)
Lev Vekker (USSR, 90 -)
Ashraf Wahab (Egypt; Fall 88 - Spring 91)
Dan Tufis (Romania; Summer 91)
Jianping Zhang (Utah State U-ty; Summer 91)
Zenon Kulpa (Poland; 91-92)
Thomas Arciszewski (Spring 92)
Tibor Vamos (Hungary; Fall 92 - 93)
Nikolai Lyashenko (Russia: Spring 93 -)
Vladimir N. Sazonow (CS Corp.; Spring 93 -)

Research Assistants and Collaborators

Nabil Alkharouf (RA; 93 -)
Sung W. Baik (TA; 94 -)
Jerzy Bala (RA; 88 - 93)
Eric Bloedorn (RA; 90 -)
Khanh Bui (92 - 93)
John Doulamis (RA; 94 -)
Tomasz Dybala (RA; 91 - 93)
Kejitan Dontas (RA; 89 - 91)
David Duff (RA; 92 - 93)
Shrinivas Gutta (RA; 94 -)
Scott Fischthal (RA; 94 -)
Susanne Furman (RA; 88 - 91)
Bob Giansiracusa (RA; 89 - 92)
Ali Hadjarian (RA; 93 -)
Raza Hashim (RA; 89 - 92)
Michael Hieb (RA; 90 - 93)
David Hille (RA; 92 - 93)
Isaac Huang (RA; 91 - 93)
Ibrahim Imam (RA; 91 -)
Chris Jardine (88 - 92)
Ken Kaufman (RA; 88 -)
James Kolodziej (93 -)
Ockkeun Lee (RA; 92 -)
Akhtar Lodgher (88 - 91)
Mark Maloof (RA; 92 -)
Jiqian Pan (RA; 91 - 93)
Jim Ribeiro (RA; 91 -)
Alan Schultz (RA; 88 -)
Larry Sklar (RA; 95 -)
Pawel Stefanski (RA; 88 -)
Witold Szczepanik (RA; 94 -)
Bradley Utz (RA; 89 -)
Haleh Vafaie (RA; 88 -)
Janusz Wnek (RA; 89 - 93)
Jianping Zhang (RA; 88 - 91)
Qi Zhang (RA; 95 -)

Computing Facilities

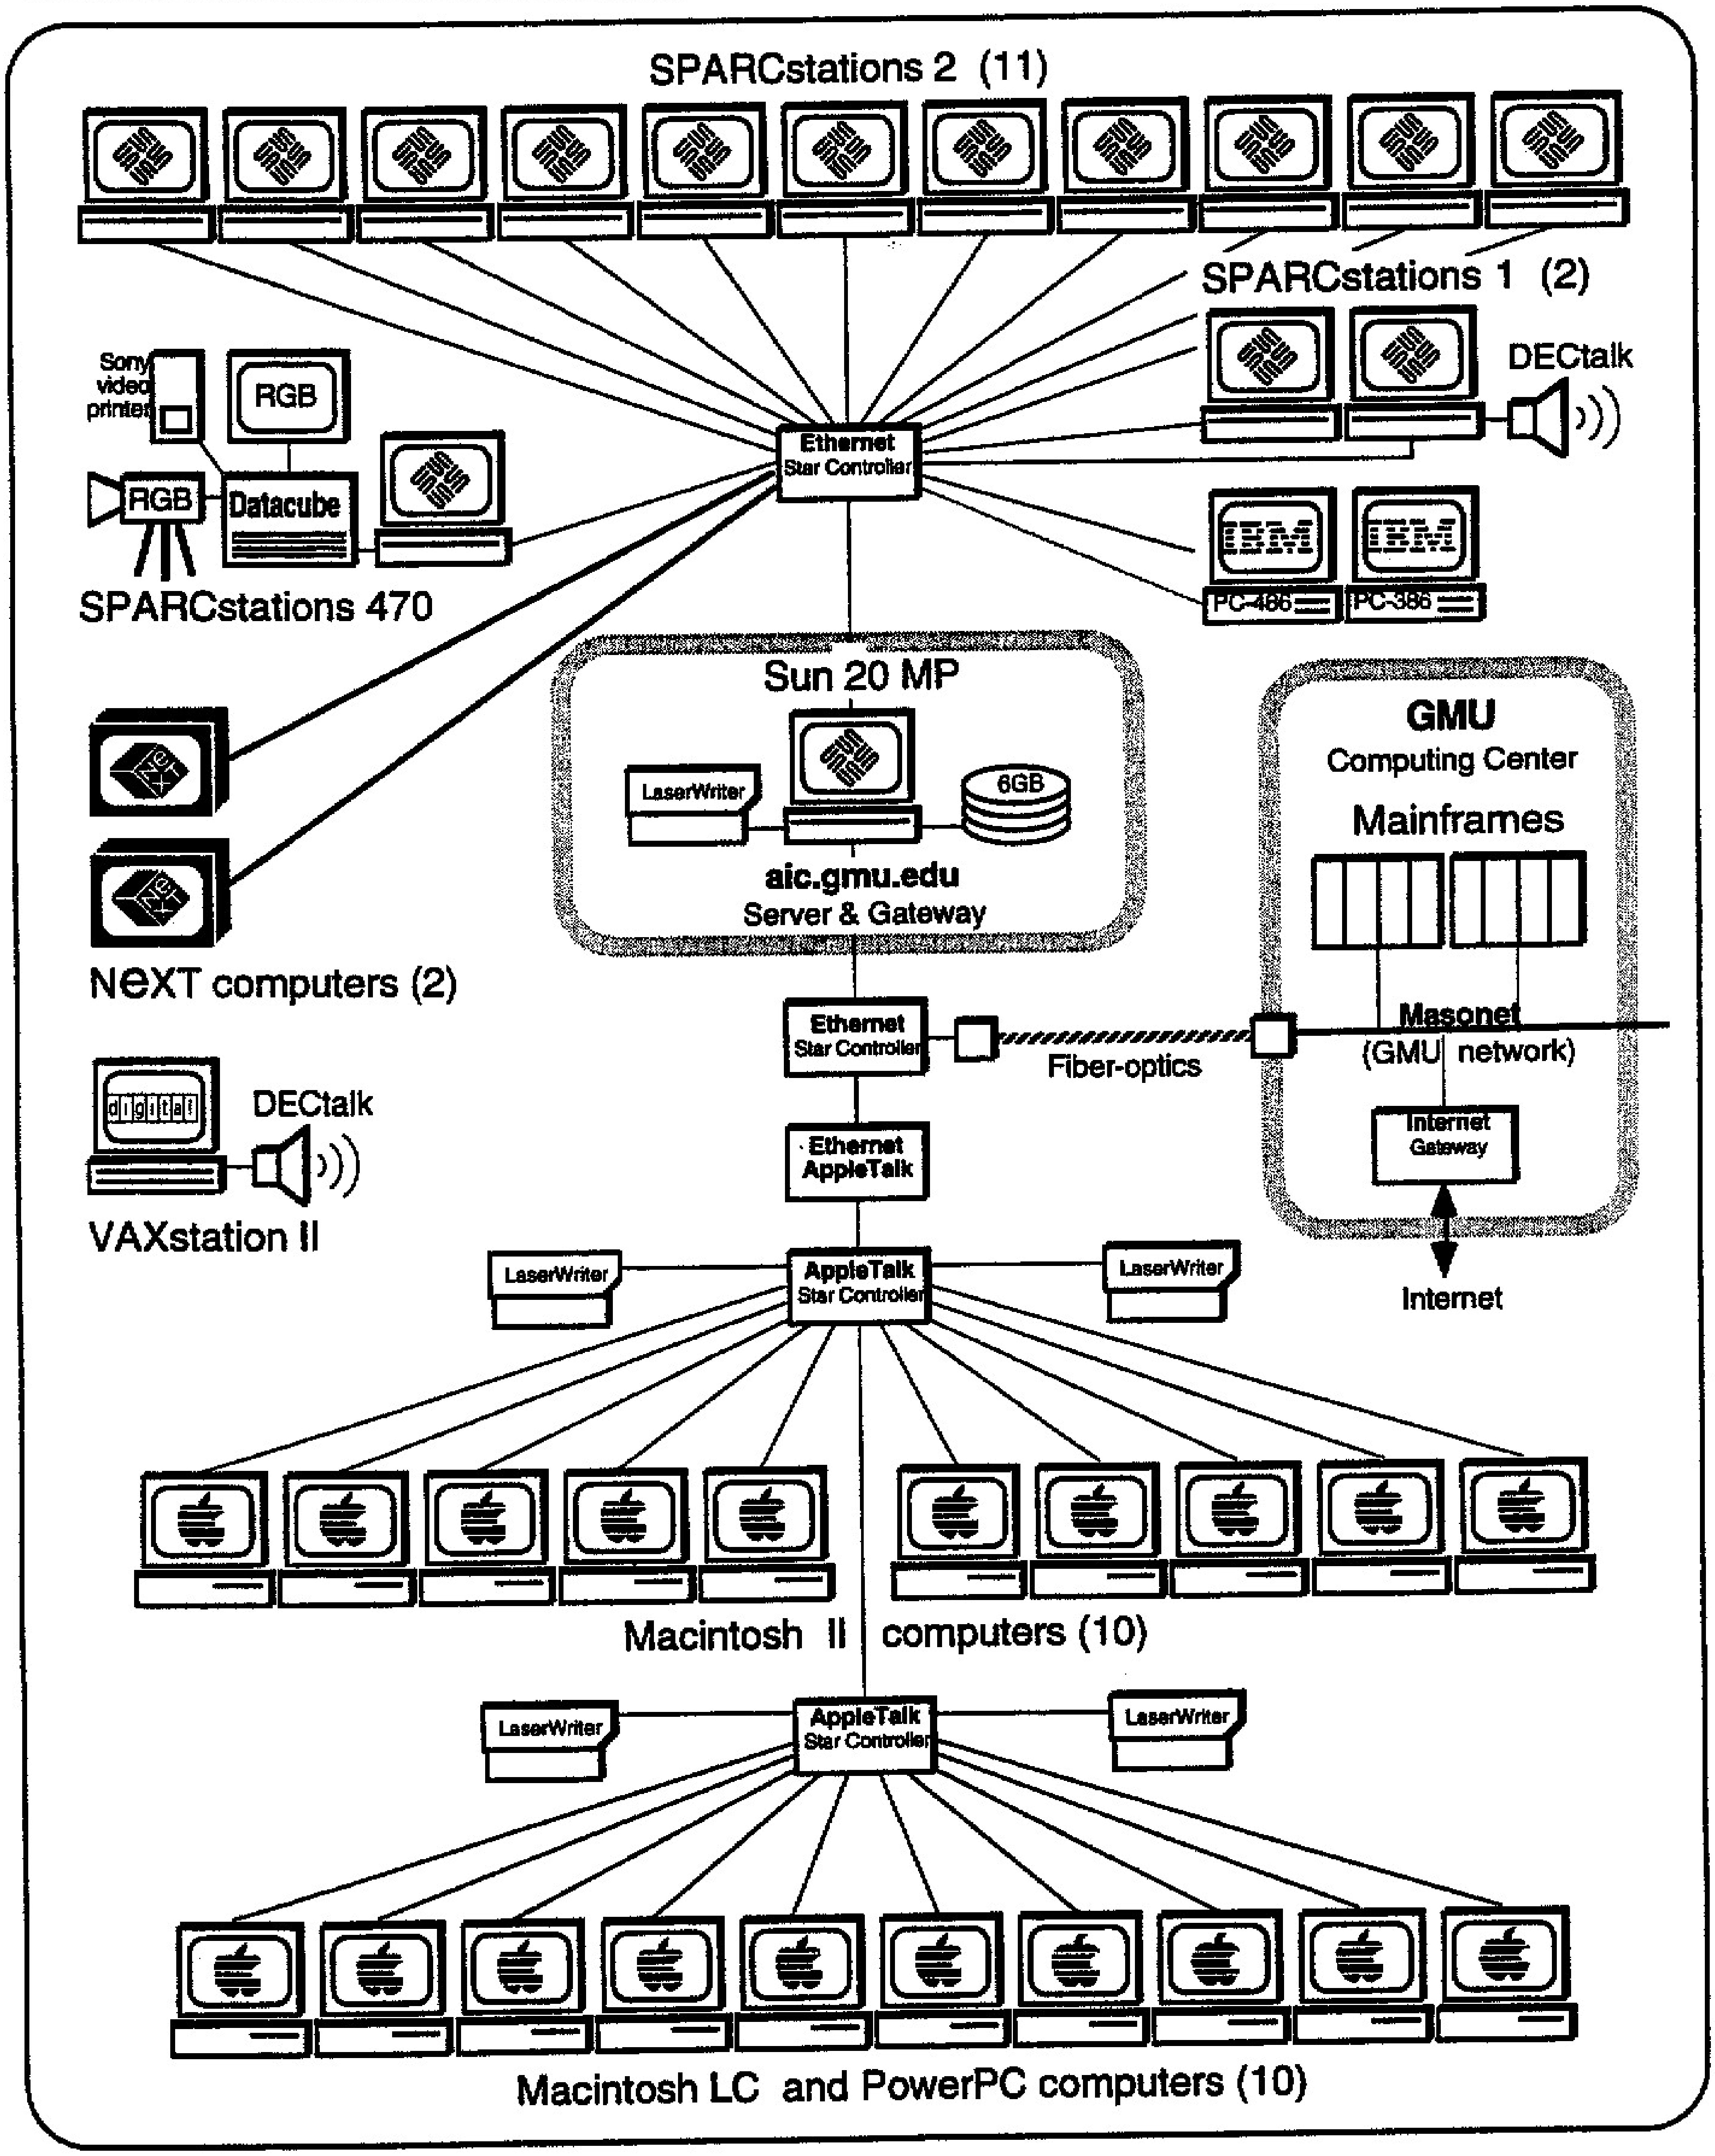

Machine Learning and Inference Colloquia

DATE	***SPEAKER***	***TOPIC***
2/8/94	Leslie Pack Kaelbling Computer Science Department Brown University	Hierarchical Reinforcement Learning in Stochastic Domains
2/9/94	Karamjit Gill University of Brighton, United Kingdom	European Perspective on Human-Centered Systems
2/25/94	Johanna D. Moore Learning Research & Development Ctr. Department of Computer Science University of Pittsburgh	Discourse Generation for Instructional Applications: Identifying and Exploiting Relevant Prior Explanations
3/8/94	Jerzy Bala and Mark Maloof MLI Laboratory, GMU	The KB Vision System and Its Use In Machine Learning
3/22/94	Vladimir Sazonov Computer Sciences Corporation	Evolving Neural Network Structure
3/29/94	David Wilkins Department of Computer Science University of Illinois	Designing Expert Shells for Learning in a Multi-agent Environment
3/29/94	Lawrence B. Wolff Computer Vision Laboratory Department of Computer Science Johns Hopkins University	Polarization Vision
4/5/94	Gary O'Brien Night Vision Laboratory United States Army	A Successful Application of Neural Nets
4/12/94	Bonnie J. Dorr Department of Computer Science University of Maryland	Machine Translation: Toward a Language-Independent Approach
7/11/94	Hugo de Garis ATR Human Information Processing Research Laboratories, Japan	The CAM-Brain Project, The Evolutionary Engineering of a Billion Neuron Artificial Brain which Grows/Evolves at Electronic Speeds inside Trillion Cell Cellular Automata Machines
7/21/94	Vladimir Sazonov Computer Sciences Corporation	Flexible Control Over Inductive Inference Based on Fuzzy Knowledge of an Expert and Assignments of a User
10/4/94	David W. Aha Navy Center for AI Naval Research Lab, Washington, DC	Learning Recursive Relations from Small Training Sets

10/11/94	Harold Morowitz Robinson Professor of Biology Director, Krasnow Institute, GMU	Databases, Expert Systems, and the Origin of Life
10/25/94	Jim Purtilo Department of Computer Science University of Maryland, College Park	Software Support for Prototyping in the Large
11/1/94	Vassilis Moustakis Institute of Computer Science, FORTH Hellas, Greece	An Overview of Machine Learning Activities at the Institute of Computer Science, Foundation of Research and Technology
11/29/94	Diana Gordon Navy Center for AI Naval Research Lab, Washington, DC	Evaluation of Inductive Biases
12/20/94	John Gero Key Centre of Design Computing University of Sydney, Australia	Recent Advances in Computational Models of Creative Design
2/7/95	James Hendler Department of Computer Science University of Maryland College Park, MD	High Performance Artificial Intelligence
2/21/95	George A. Hazelrigg National Science Foundation Washington, DC	Systems Engineering: An Approach to Information-Based Design
2/23/95	Reuven Karni Technion Israel	Industrial Applications of Knowledge-Bases
3/21/95	Wayne D. Gray Krasnow Institute George Mason University	Application, Science & Engineering Prediction & Validation
4/11/95	Menas Kafatos CSI George Mason University	The Earth Observing System and its Data Information System
9/5/95	Mary-Anne Williams Information Systems Group Department of Management The University of Newcastle Newcastle, Australia	Transmutations of Information Systems
10/10/95	Laveen Kanal Department of Computer Science University of Maryland College Park, MD	Sometimes the Magic Works: On ANN's, Statistics, and Pattern Recognition

10/24/95	Krzysztof Kuchcinski Department of CIS Linkoping University Linkoping, Sweden	High-level Synthesis of Digital Systems Specified in Behavioral VHDL
10/31/95	Mark Pullen Department of Computer Science George Mason University	Distributed Interactive Simulation Technology Challenges
11/13/95	Ed Fox Department of Computer Science Virginia Tech Blacksburg, Virginia	Digital Libraries and Computer Science Education
11/2195	Kathryn Blackmond Laskey Department of Systems Engineering Krasnow Institute, C3I Center George Mason University	Bayesian Decision Theory and Machine Learning
12/5/95	Ophir Frieder Department of Computer Science George Mason University	Integrating Database and Information Retrieval: A Parallel DBMS