

**MACHINE LEARNING and DYNAMIC RECOGNITION
for ROBOT VISION**

A Research Proposal Submitted to

The National Science Foundation

by the

**Center for Artificial Intelligence
George Mason University
Fairfax, VA 22030**

for consideration by

NSF-DARPA Initiative, IRIS/CISE

Principal Investigators:

R. S. Michalski, H. Wechsler and H. Ko

(703) 764-6259 or 764-6250 or 323-2713

MACHINE LEARNING and DYNAMIC RECOGNITION for ROBOT VISION

Summary

The proposed research represents an interdisciplinary effort to bridge two fields, machine learning and computer vision, and aims at investigating several novel ideas and topics related to these areas. Main research topics include:

- Combining symbolic and holistic methods for learning and recognizing 3D objects from 2D images.
- Exploring the method of *two-tiered representation* for learning imprecise and context-dependent visual concepts from examples.
- Applying the idea of *dynamic recognition* to automatically determine relevant visual classification clues and to control active perception.
- Determining a structure of operators for characterizing surface and texture by methods of inductive learning.

The proposed methods for learning and recognizing visual concepts are designed to handle noisy visual inputs, object occlusion/overlap, and class-preserving transformations, such as translation, perspective change, scale variations, and rotation. We plan to develop computational methods, implement them and experimentally test them on selected practical tasks. We will also explore the problem of integrating symbolic and holistic memory representations in the planned experimental system.

The proposed research builds upon the long-term past accomplishments of the investigators in areas such as machine learning, computer vision and expert systems with multiple knowledge representations and learning capabilities (see Appendices).

This research is highly relevant to the factory of the future, autonomous robotics and office automation, where visual feedback is crucial for execution and coordination of various tasks, such as object recognition, localization of objects in space, and collision avoidance. The research will be carried out at the Artificial Intelligence Center of George Mason University (GUM), in cooperation with the Computer Vision Laboratory (CVL) at the University of Maryland at College Park, with the Denver Aerospace division of the Martin Marietta Corporation and Machine Perception Research Department of the AT&T Bell Laboratories.

TABLE of CONTENTS

Summary

PART I: Research Program

1. Introduction

- 1.1. Research Goals
- 1.2. Relevant Research of Investigators and Project Justification

2. Learning and Recognition of 3D Visual Objects from 2D Images

- 2.1. Overview
- 2.2. Symbolic Approach
 - 2.2.1 Representing and learning imprecise and context—
dependent visual concepts using *two-tiered representations*
 - 2.2.2 Data and model-driven learning of structural descriptions of
3D objects
- 2.3. Holistic Approach
- 2.4. Integrating the Holistic and Symbolic Approaches

3. Dynamic Recognition

- 3.1. Basic principles
- 3.2. An outline of the theory of dynamic recognition
- 3.3. Summary of ideas and proposed research

4. Recognizing Objects Through Surface Properties and Texture

5. Conclusions

6. References

PART II: Research Team and Research Environment

- 1. Accomplishments of PIs on Previous NSF Grants
- 2. Biographical Sketches of Investigators
- 3. Facilities and Equipment

PART III: Budget

- 1. Proposed Budget
- 2. Current and Pending Support

Appendix 1. An Algorithm for Dynamic Recognition

Appendix 2. List of Publications of the AI Center for the Last 5 Years

Appendix 3. Selected Publications

PART I: Research Program

1. Introduction

1.1 Research Goals

Intelligent behavior is characterized by the ability to learn and to self-organize. Learning includes self-creation of knowledge through experience, self-improvement of control through practice, discovery of new concepts and relationships through observation and analogy, and reorganization of knowledge to perform given tasks better. When we consider human or animal vision, it is clear that they acquire the ability to recognize and handle visual objects to a large extent through learning, and only partially by pre-programming. Yet, current computer vision techniques are almost devoid of learning capabilities. We believe that further progress and the ultimate solution of computer vision depends on the implementation of machine learning capabilities in vision systems.

The proposed interdisciplinary research aims at applying modern machine learning techniques to computer vision problems, and involves a number of novel ideas and topics for investigation. These topics include:

- Combining holistic and symbolic methods for representing and learning 3D visual concepts from 2D images. A *holistic representation* is the result of some continuous transformations of the input images that preserve important topological and geometrical properties of whole objects, while a symbolic representation views objects as networks of named components and defined relationships among them. This project includes studying relative advantages and disadvantages of both representations, and developing methods for acquiring and integrating such representations in an experimental system for robot vision.
- Exploring the method of *two-tiered representation* for learning imprecise and context-dependent visual concepts from examples. The idea of two-tiered representation, introduced recently by one of the PIs (Michalski, 1987), provides a novel way for representing such concepts (they are called *flexible*). Experimental results on representing non-visual flexible concepts (e.g., human diseases or acceptable labor contracts) has shown that this representation may lead to a significant decrease in the size of description, and at the same time to an improvement in the performance of classifying new instances.
- Applying the idea of *dynamic recognition* to automatically determine relevant visual classification clues. The concept of dynamic recognition allows a system to recognize a given object from a great variety of different pieces of information about it, without storing explicitly the rules for every type of recognition.
- Determining rules for recognizing objects through properties of the surface and texture using methods of inductive learning. This is a novel approach to automatically acquiring surface and texture descriptions by determining the most significant spatial clues and unique texture characteristics. The method seems to offer a very powerful component in a system for recognizing visual objects.

In the subsequent sections, we will briefly describe our past research and the proposed research plan, and give more details about the projects mentioned above.

1.2 Relevant Research of Principal Investigators and Project Justification

The proposed research reflects the natural outgrowth of research conducted by the principal investigators and their collaborators over the years. R. S. Michalski has worked in the field of machine learning for many years, and is considered a co-founder of the field. He also worked in the areas of inference systems, cognitive modeling and expert systems. Prior to his work on machine learning and inference, he worked in the areas of computer vision and pattern recognition (e.g., he co-developed an early and very successful system for recognizing handwritten alphanumeric characters, developed a method for designing of what now would be called a connectionist system, and worked on various computer vision problems within the ILLIAC III project (Michalski, 1973).

Dr. Michalski has contributed a number of novel ideas and methods to the fields of artificial intelligence, pattern recognition and cognitive science, such as the concept of variable-valued logic, a theory and algorithm AQ for inductive learning, the concept of inductive generalization rules, conceptual clustering, constructive induction, qualitative prediction, variable-precision logic (with P. Winston), developed (with A. Collins) a theory of human plausible reasoning, and recently introduced the idea of two-tiered concept representation and dynamic recognition (for more details see Part II and Appendices 2 and 3).

Wechsler has been working on a number of computer vision problems, in particular, on problems of fault-tolerant object recognition. Several recent publications describe his current research (see attached list of MLI publications). The focus of his work is to test the usefulness of holistic, invariant image representations and PDP (parallel distributed processing) modeling for robust object recognition. Such processing includes implicit learning of object prototypes, and provides for adaptation with respect to stochastic variations. Most recently, he has been nominated to serve as one of the directors for the forthcoming NATO Advanced Study Institute on Active Perception and Robot Vision to be held in Maratea, Italy, July 16-29, 1989.

H. Ko has conducted research on several topics in machine learning, in particular, concept learning, the part-to-whole generalization and qualitative prediction. He also worked on intelligent computer-aided design, geometric reasoning and, most recently, on integrating planning and learning for automatic assembly and an intelligent autonomous robot (the IEX project).

The GMU Center of Artificial Intelligence is conducting research on many contemporary problems of machine learning, inference and computer vision, and on an integration of these capabilities in intelligent systems. One of our main application domains is intelligent autonomous robotics (the Intelligent Explorer Project - IEX). To apply our ideas to this domain, we need to have a computer vision system that could be integrated with learning and inference subsystems. Combining learning and vision is thus a very important research topic for our team. This proposal is to get support enabling such a research, and to facilitate cooperation between already supported projects and vision research.

The main thrust of this proposal is to investigate the applicability of several new ideas in machine learning, in particular, two-tiered representation, constructive learning and dynamic recognition, to computer vision tasks, such as object recognition, classification and detection of obstacles for safe navigation. We plan to study inferential symbolic machine learning and holistic techniques, and how to relate and integrate both approaches. We would also like to understand how the explicit visual knowledge in a rule-based form, obtained from high level symbolic processing, relates to the implicit knowledge acquired by a holistic method at a lower level of the visual architecture. Our goal is to develop methods that will work well with input images that may contain noise, object overlap or occlusion, changing viewpoints and image variations (due to the scale change, rotation, translation, and/or projection). As to the learning recognition rules (as opposed to their use for recognition), we will initially assume that visual objects are presented in relatively undistorted forms. In summary, we

are interested in developing a robust, fault-tolerant system for robot vision, characterized by graceful degradation of performance in the case of noisy inputs, and geometrically distorted and/or occluded objects.

Methods of structural inductive learning, such as implemented in the INDUCE series of programs (Larson and Michalski, 1976; Hoff, Stepp and Michalski, 1983), have been known to apply to computer vision, but this has not yet been done because of the lack of facilities and support. The new work on *constructive learning* (Michalski and Ko, 1988) that integrates empirical and analytical learning (e.g., constructive induction and explanation-based learning) is especially promising for learning visual objects. Also, the idea of dynamic recognition (see section 3) seem to ideally apply to computer vision, as it deals with the problem of automatically identifying relevant context-dependent recognition clues. But again, this idea has not yet been applied to computer vision problems. The proposed project will enable us to do this research, and have a collaboration between learning and computer vision projects at our AI Center, and other organizations, specifically, the Computer Vision Laboratory of the University of Maryland, the Denver Aerospace division of the Martin Marietta Corporation, and the Machine Perception Research Division of AT&T Bell Laboratories.

2. Learning and Recognition of 3D Visual Objects from 2D Images

2.1 Overview

We plan to develop a methodology for a vision system with learning capabilities. We will explore several new ideas on how to represent, learn and recognize visual objects, how to cope with the imperfection of the input, and the imprecision of visual concepts. We will design and implement an experimental system, and test it on selected practical problems. As to the testing, we will concentrate on objects occurring in typical office scenes, such as chairs, tables, pencils, cups, monitors, doors, letters, books and telephones.

The starting point of our approach is the realization that most human visual concepts lack precise definitions and their meaning depends on the context in which they are used. Consider, for example, generic concepts referring to the office objects mentioned above. It is very difficult to formulate general and precise definitions of such concepts in terms of visually detectable properties. Yet, humans usually have no difficulty in visually recognizing such objects in their environment. They do well even though the objects may occur in a very large number of different forms, can be partially occluded, and can be located in an environment full of other objects. It is clear that any recognition system for visual objects has to be able to construct and work with imprecise and context-dependent definitions of visual concepts, and to cope with the immense variety of appearances and forms in which objects occur. Most research in computer vision has been concerned with the second problem. There has been relatively little effort on the first problem (concept imprecision). In fact, most research assumes that concepts are precise entities, with well-defined boundaries, and are context-independent.


To deal with the problem of concept imprecision, we will pursue the idea of *two-tiered concept representation*, and to deal with the problem of variability of concepts appearances and forms we will pursue the idea of *dynamic recognition*. The concept of two-tiered representation was proposed by Michalski (1987) to explain how people learn and recognize imprecise and context-dependent concepts (non-visual or visual) and the variability of appearances and forms. According to this idea, the meaning of any concept is defined by two components: the Basic Concept Representation (BCR) and the Inferential Concept Interpretation (ICI). The BCR defines the most typical properties of the concept. The ICI makes the boundaries of the concept flexible by describing allowed modifications and transformations of the concept's properties, and their context-dependence. In general, the ICI may contain metarules and inference rules for deduction, analogy and induction, and definitions of allowed concept transformations. Concept anomalies, exceptions and context-dependencies are thus

handled by a reasoning process involving rules contained in the ICI and the context. Learning any concept requires constructing appropriate BCR and ICI. Experimental results with two-tiered representation of non-visual concepts have shown that by using such an approach one can significantly reduce the amount of memory for storing the representation, and simultaneously improve the performance of the representation in recognizing new instances of concepts (e.g., Michalski, 1987; Begadano, Matwin, Michalski and Zhang, 1988 a and b).

In the context of vision, the BCR is called the *base referent*, and the ICI is called the *focusing transform*. The base referent contains information characterizing typical and most relevant concept characteristics. In order to recognize an object, the given image of an object is compared to the information in the base referent. The focusing transform consists of rules and transformations that remove from the image any information that is irrelevant to the concept; ideally, the transform only preserves information relevant to the conceptual identity of the object represented by an image. There are two basic types of the base referent: holistic and symbolic. The holistic referent is a result of some continuous transformations of a typical image of an object, which preserve the basic topological and geometrical properties of the object. The simplest holistic referent is an object template. A symbolic referent represents a visual object (an instant or a concept) as a system of named components and named relationships among them (for example, as a semantic net or a logic-based description, such as the one used in the learning program INDUCE).

In our system, both holistic and symbolic representations will be employed. Such a *two-prong* approach seems to be consistent with findings in human vision. Studies done on patients suffering from disorders of visual perception due to brain damage led Warrington (1982) to suggest a modular organization for perceptual categorization and interpretation. It was suggested that the right hemisphere system (RHS) is concerned with perceptual (holistic) categorization, while the left hemisphere system (LHS) deals with semantic (symbolic) categorization. Holistic categorization is mostly concerned with matching objects by visual low level similarities, while symbolic categorization handles matching by functional or semantic similarity. We will assume that an input image undergoes some initial basic transformations (preprocessing and segmentation), and then there are two "prongs" that independently process information, one generating a holistic representation and one generating a symbolic representation. To generate symbolic representation, the image undergoes a variety of transformations (based on contour extraction, color measurement, surface and texture characterization, corner detection, etc.). At some higher level, the holistic and symbolic representations interact, and have access to a global knowledge base about visual objects, and other knowledge (Figure 1).

The first computational stage supplies sensing, preprocessing, and segmentation results to both the holistic and symbolic processing modules. Sensing and preprocessing provide the raw, analog data, like range, orientation and color information. The segmentation is done using standard methods based on edge detection, surface fitting, differential geometry (Besl and Jain, 1988), multiple cues, and surface and texture characterization. This yields planar or curved surfaces (needed for further holistic processing) and visual primitives needed for symbolic processing. The resulting surfaces and/or symbolic primitives correspond to 2D visual images, known as aspects or characteristic views (Koenderink, 1984), and they are linked by specific relationships. The holistic processing stage is concerned with achieving invariance to geometrical transformations, occlusion/overlap and noise, in the process of recognizing the identity of the surfaces resulting from segmentation. The symbolic processing builds symbolic logic-like descriptions, in which 2D and 3D visual concepts are modeled in terms of primitive units. When these representations are used for object recognition, any one of them may trigger the recognition. At the highest level there is an interaction between the holistic and symbolic representation. A partial and flexible matching can occur at the holistic level, and this allows for dealing with 'imperfect' images. Similarly, the symbolic processing can deal with imperfect symbolic representations. Holistic processing can suggest partial interpretations, those are instantiated as symbolic structures, and then sensing, preprocessing, segmentation, and holistic processing can be dynamically reset.



A Two-prong Representation of Visual Objects

Figure 1.

We propose to start our efforts by confining the visual tasks to the recognition of static objects. Humans recognize familiar objects with remarkable ease. For example, chairs come in a large number of forms. Some of them have four legs, other have only one or no legs at all; some of them have arms, while others do not; some have solid and straight backs but others do not, and so on. There is no simple way to define the general shape of a chair. As a result, currently no robot has the capability of recognizing 'simple' objects, such as chairs, tables, lamps, that a robot might encounter. We believe that the main reasons for the lack of success of present visual paradigms in recognizing objects are inappropriate representations of such flexible concepts and lack of learning capabilities. Learning processes should take care of automatic rule creation, the development of attentive processes (dynamic vision), and generalization from examples. Another example of learning is when a robot is viewing a partially occluded object or viewing an unfamiliar side of the object. Here the robot must learn which direction is more likely to provide a more informative view and hence decide to move in that direction. At a close distance, the proximity of two different objects does not pose any difficulty for the recognition of either of them by humans. This, however, must not be taken for granted when dealing with robots. For example, two chairs put side by side may be interpreted as only one object. We therefore need to develop capabilities for a robot to cope with all these situations and be prepared for unexpected contingencies. Again, such capabilities must be acquired through a learning process utilizing mostly depth discontinuities and surface orientation. For the above tasks, there is no need to compute depth accurately (which requires camera calibration and has proved to be a difficult problem; Kamgar-Parsi, 1987).

Different stages of visual representation are thus employed by the human visual system, according to the functionality of the task. In an analog way, low-level, holistic, invariant image representations might be appropriate for the very first stages of visual processing, while symbolic processing could be appropriate for the higher levels of visual cognition. Much of the fault tolerance exhibited by the visual system could then be the result of implicit invariant canonical image representations which are adaptive to stochastic changes due to their distributed and holistic representations. The symbolic stage of processing is responsible for much of the learning, which includes adaptation based on the functionality and the specificity of the visual task. There is much information exchange between the holistic and symbolic stages for solving ambiguity and speeding up the recognition task.

The vision system that we propose is based on the architecture discussed above. Thus, there are two structures involved in both recognition and learning; specifically, those for holistic and symbolic processing, and they correspond to perceptual and semantic interpretation.. While we have already outlined the basic conceptual architecture of such a system, it is quite likely that the implementation of specific modules will evolve as a result of actual experimentation. Our initial design assumes that the first computational stages, those of sensing and preprocessing, can feed both the holistic and symbolic modules. Sensing provides the raw, analog data, like range, intensity, and ultrasound. Appropriate understanding of the image formation process (physics, geometry, and optics) then yields intrinsic information like color (Klinker, Shafer, and Kanade, 1988) and orientation surface normals. The following sections discuss a perceptual and cognitive architecture of the proposed system, and the interface between the symbolic and holistic components.

2.2 Symbolic approach

2.2.1 Representing and learning visual concepts using two-tiered symbolic representations

In the symbolic approach, visual objects and classes of visual objects (concepts) are represented as systems of named components and named relationships binding these components. Such structures are organized into different levels of abstraction. Components that correspond to each other at different levels of abstraction are indicated. At the lowest level of processing, the symbolic representation will include information from the module that performs initial image transformations (e.g., local feature extraction, contour determination, color measurement, texture and surface characterization, and possibly also depth information (sec. 4, Part 1)). At the highest level, the symbolic representation will be fused with a holistic representation, by annotating image components by symbolic labels (an *annotated map*).

As mentioned before, visual concepts typically do not have precise definitions and can be manifested by objects appearing in a very large number of forms. To cope with these problems we will employ two tiered representations. To give a clearer characterization of such a representation, let us consider an example. Suppose we want to describe the concept of a *chair*. A two-tiered symbolic representation of this concept would consist of a base concept representation (BCR) that represents typical properties of chairs, and an inferential concept interpretation (ICI) that would define allowed chair modifications, context dependency and special cases:

BCR: A piece of furniture typically used for sitting by one person. Usually consists of a seat, four legs and a backrest. The seat is typically about 1.5 to 2.5 feet above the ground.

(A low-level holistic representation would consist of one or more typical and simplified images of a chair)

ICI: (no-of legs may vary from 1 to 4) (instead of legs there can be any object that provides a support for the seat) (the shape, the size, the color and the material of all components may vary as long as the function defined in the BCR is preserved) (if legs are replaced by wheels) --> wheelchair
(chair without the backrest) ---> (stool rather than chair)
(chair with arm-rests) ---> (an armchair)
(context = museum exhibit) --> (chair is not used for seating any more)
(context = toys) --> the dimensions and the height of a chair can be much smaller.

(A holistic counterpart would consist of various transforms that make the image invariant of translations, rotations, scale changes, small distortions and shape modifications)

This simple example illustrates several important features of the two-tiered representation. The BCR carries the typical properties which may not always occur. Any perceived object is first directly matched against the BCR. If this fails, then a flexible match is tried using information in the ICI. The ICI rules define allowable object transformations under different contexts. In general, contexts can have hierarchical organization. Also, the ICI rules may chain, though this is not shown in the example.

At the symbolic level, we will initially investigate the applicability of the already existing learning programs, specifically, INDUCE 3 and 4 (Michalski and Stepp, 1986; Bentrup, J., Mehler, G., and Riedesel, J., 1987). These programs construct general structural descriptions of classes of objects (a concept) from instances of the classes provided by a teacher.

Learning two-tiered representation will proceed in two stages: 1. determining the BCRs of visual concepts using the INDUCE constructive learning methodology, and 2. reducing the obtained representations by a description optimization method. Details on this procedure are in the enclosed papers on learning TT representations by Bergadano, Matwin, Michalski and Zhang (Appendix 3).

Thus, the basic idea is to use already developed programs for constructing the initial BCRs, and then optimize them in order to construct the final BCR. The ICI rules will initially be provided by direct programming. Later, we will develop separate programs for learning the ICI rules from examples. Subsequently we will work on the development of a general methodology, and new experimental programs for learning two-tiered representations. The input descriptions to the system will initially be encoded by an experimentator. When the low level feature extraction and preprocessing modules become available, the input descriptions will be generated by a program. We will also investigate problems of the interaction between the processes of constructing symbolic and holistic representations.

2.2.2. Data and Model-driven Learning of Structural Descriptions of 3-D Objects

We will investigate relative merits of data and model-driven approaches to learning 3D symbolic descriptions of objects from 2D projections. The underlying idea is to use projected 2D images of objects as input data for a learning system to synthesize descriptions of 3D objects.

The goals of this project are:

- 1) Incorporate domain-specific (vision related) knowledge into existing algorithms for learning structural descriptions from examples (e.g. INDUCE-4),
- 2) Develop methods for efficient mapping between structural descriptions of multiple images,
- 3) Develop heuristic methods for controlling image acquisition process, and
- 4) Develop methods for generating required 2D projections and estimating parameters for acquiring new images from the holistic module that disambiguate multiple candidate 3D models postulated by the learning system.

4) Develop methods for generating required 2D projections and estimating parameters for acquiring new images from the holistic module that disambiguate multiple candidate 3D models postulated by the learning system.

The proposed research will help to build a vision component of an implemented system NOMAD (NORMative Mechanical Assembly Device) that constructs an assembly sequence from the structural description of an assembly based on mating conditions [Ko & Lee 1987; Ko, 1988; Ko 1989]. Currently, all the 3D object models are given using a geometric modeling system. Because the kinematic connections in the assembly structure is described symbolically using mating conditions, it may be beneficial to integrate object recognition process that may start with symbolic description of the components of an assembly. With the object recognition capability, it can gracefully refine the symbolic description of the assembly structure for finer manipulations

A related project is concerned with an automatic construction and refinement of symbolic models of structured objects [Mozetic, 1987a, 1987b]. The system involves learning structural descriptions from examples [Mozetic & Lavrac, 1988] by locating incorrect and incomplete descriptions in the model and refining the model incrementally from the diagnosis of failures in predicting the behaviour of the model. The approach enables the representation and interpretation of a symbolic model on different levels of abstraction, thus increasing its diagnostic efficiency [Mozetic & Urbancic, 1987, Mozetic, Bratko & Urbancic, in press]. The initial system was successfully tested on reconstructing a model of the heart's electrical activity [Bratko, Mozetic & Lavrac, in press].

Based on the developed ideas for non-visual object modelling we will build a new system that given 3D generic primitives will generate a model of complex 3D objects. These 3D primitives may be at different levels of abstraction from geometric models to symbolic models with varying degrees of abstraction. For example, a cylinder may be a 3D generic primitive and modeled precisely at the geometric level; its corresponding symbolic description may describe it as a structural object with a top face, bottom face and a side face and their topological relationships such as the top face borders on the side face at one end and the bottom face borders on the other end. The shapes of these faces (circular for the top and bottom face and rectangular for the side face) are derived from the geometric model. Their exact boundaries are described at lower level of representation. For symbolic 3D primitives, their corresponding 2D projections will be derived by a symbolic processing using a set of inference rules: e.g., the projection of a top face of a cylinder may be a circle, an ellipse, or a line when the viewing angles are from the top, the slanted side, or the right side respectively.

When given a 2D image of an object, the system constructs first a 2D model of the image in terms of the 2D primitives derived from 3D primitives above. When 2D model of the image is constructed, candidate 3D models are hypothesized. This process requires an incremental modification of the candidate model descriptions and may require new images that would maximally disambiguate candidate models. For example, if an image of a cup is presented from the side, only the rectangular image can be explained as a projection of a cube or a cylinder each of which may be a part of the whole 3D object model. To disambiguate between these candidate 3D models, the system may ask an image of the cup from the top. The outline of the proposed system is given as follows:

GIVEN

1. 3D symbolic primitives
2. 2D projections of these 3D primitives

INPUT

1. 2D images that are transformed from 2D holistic to low-level 2D symbolic representations

LEARNER

1. Explain 2D symbolic images as projections of 3D models
2. Construct possible 3D models
3. Construct plausible 3D models

DEBUGGER

1. Plan a sequence of sensory actions to find the sources of mismatches.

CRITIC

1. Use 3D model to predict its projections
2. Debug the model if it is incomplete (some projections cannot be predictable)
3. Compare the predicted projections with real world images
4. Debug the model if it is incorrect (some predicted projections do not match with the real world images).

This research project will proceed as follows: 1) the development of symbolic languages at several levels of abstraction that qualitatively describes objects, 2) defining transformations between these levels of abstractions and then to the numerical representations, 3) improving existing learning methods to handle mappings between image models of the same object 4) the development of an efficient part-to-whole generalization methods (Michalski, Ko, and Chen, 1987) that postulate plausible 3D models from 2D image models. 5) the development of planning strategies for acquiring 2D images that would maximally disambiguate candidate images.

2.3 Holistic Approach

We will investigate the applicability of invariant and distributed image representations for holistic processing. The holistic processing is responsible for providing initial estimates regarding the type, properties and location of surfaces of objects being observed. Those estimates may be sufficient by themselves to trigger the recognition, or they can become units of information which are passed to the symbolic stage for further processing. Holistic processing draws its inputs from sensing and preprocessing and segmentation, or from suggestions provided by the symbolic module. Such suggestions could guide the holistic module in resetting specific parameters, as it might be needed for flexible matching and/or shifting the area of interest for further processing. We will explore how the robot vision system can incorporate geometric invariance and how it can withstand both stochastic changes (due to sensory noise or illumination changes) and incomplete (occluded/overlapped) images. The design and implementation of the holistic stage will be an extension of the current research on robust vision for robots (Wechsler and Zimmerman, 1989; see Appendix 3).

The basic paradigm of holistic processing is that of matching preprocessed and geometrically transformed input images against canonical memory prototypes (which represent classes of visual objects in different views). The input is transformed into a form invariant to geometrical transformations, like scaling, rotation, and perspective. These transformations are based on ideas developed by Wechsler (1988,1989). One of the experimental systems he developed achieved invariance to scale and rotation by transforming the original image representation from the cartesian domain to the polar domain, using the complex-log mapping. Another system, FOVEA (Jacobson and Wechsler, 1986), achieved perspective invariance using a limited range of reprojections. The required number of reprojections is usually small, and it is suggested by distortions of texture of the surfaces analyzed. A number of further issues need to be studied in order to achieve invariance to geometrical transformations. Specifically, one needs to determine what is the appropriate image resolution for implementing polar transformations, and reprojections. Integral components to the solution of this problem are the methods of two-tiered representation and active perception.

To recognize an object one needs to match one or more images of an object (that have been sensed from one of several possible viewpoints) against one or more canonical memory prototypes. Those prototypes represent specific 2D reprojected views of an object. Because there are many viewing angles possible, there may be many different images corresponding to them. The number of such viewpoints can be reduced by considering properties of the visible shape that stay approximately the same for a limited range of viewing angles. Such views are called aspects (or characteristic views or stable states). An aspect constitutes then a canonical object prototype corresponding to a given viewpoint of the object.

The stored aspects of an object might include information like depth, orientation, intensity, and color. Multiple streams of computation for sensory and / or intrinsic data are then integrated via data fusion methods. The data fusion approach is modeled in terms of Thin Plates (Terzopoulos et al, 1987), and it allows for flexible matching that may be guided by the result of previous memory searches. The flexible matching is with respect to what to be fused, over which area to do it, and how much confidence to attach to each of the inputs being fused.

Visual perception involves sensing of the environment in order to corroborate conjectures by redirecting the attention to those events and/or visual concepts which are most likely to increase the confidence in what is being perceived. A viewer is said to be active, if it changes the position of its sensors in order to improve the quality of perception. Although most research in robot vision has dealt with passively obtained images, humans view the environment through an active process, e.g., by rolling eyes and moving heads. In many situations passive perception has proved to be an ill-posed problem, giving rise to ambiguous interpretations. To compensate for the lack of sufficient information, various constraints have been imposed on the problem, such as smoothness assumptions and regularization, but this has met with a limited success. However, in active perception, the viewer, at least in principle, is able to acquire enough information about the scene, so that there is no need for the imposition of questionable constraints. The method of active perception is complementary to the derivation of invariants in the process of robust recognition. A major goal of this research is to demonstrate that active perception is one of the main ingredients for successful object recognition and safe navigation. The role of active perception can be:

- a) Exploratory, in which new sensed inputs are acquired to disambiguate or enhance recognition; or
- b) Functional, in which new rules and controls (e.g., reprojections) are initiated and/or parameters are set, in order to decide what and how to process the sensed image. As a result of the last interpretation attempted by the holistic stage, the symbolic processor decides what are the most likely intrinsic sources (color, and/or texture, and/or orientation) to be fused, what constraints have to be satisfied, what is an acceptable error for the recognition process, and what might be the most likely geometrical transformation the objects have been subject to.

Crucial to representing a 3D object in terms of aspects requires the ability to limit the number of needed aspects, in order to make computation efficient. Since multiple aspects represent multiple viewing angles, a link between a pair of aspects may be maintained. The link may be used to denote the egomotions that are necessary to retrieve additional aspects when some aspects are missing to validate the identity of the imaged object. Such information is passed to the holistic module by the symbolic processor, if the confidence in the hypotheses suggested so far (regarding object's identity) is low and additional information for identification is needed. Both, the *two-tiered representation* and *dynamic vision* will be explored and tested for reducing the number of aspects needed to represent an object and for improving efficiency of matching..

The stored aspects are distributed within the visual memory in order to make the overall vision system robust to occlusion/overlap situations, noisy inputs, and memory faults. The visual memory is of PDP type to facilitate searching in parallel, and to extract average prototypes for the characteristic views of an object. Our recent research on invariant object recognition strongly suggests that a distributed memory representation of the associative (DAM) type (Wechsler and Zimmerman, 1988, 1989) is capable of accomplishing robust and parallel recognition.

2.4 Integrating Holistic and Symbolic Approaches

One of the important goals for our research is to develop an interface between the holistic and symbolic processing. Past research in data and model-based vision assumed that complete 3D models of objects to be recognized are already known to the system, and that those models are precise. The robot vision system we plan to develop learns new 3D models and uses them for future

recognition. Furthermore, our research efforts will focus on developing a link between the holistic and symbolic processing that will enable the system to robustly recognize objects when facing noise and incomplete input images and to also reduce the complexity involved in dealing with the large repertoire of possible object forms, and their possible appearances.

An interface between the holistic and symbolic modules can be implemented by sharing a common database, called an *annotated map*. The map represents a current interpretation of the object being perceived and analyzed. The map includes the projected images that are initially constructed from low level sensing and preprocessing. The holistic module may annotate the type of surfaces and their locations on the map. Then, the symbolic module may annotate further what 3D object models in memory these surfaces may belong to. Because the 3D models in memory rarely match exactly with the image, we may explore the use of the novelty filter (Kohonen, 1988) to ascertain whether the map describes an object the system has learned before or not. If the map describes an object that is sufficiently similar to some object model in memory, the symbolic module will consider these models as candidates. When it is ambiguous (there are more than one candidates), the symbolic module may request the low level sensing module to acquire an additional image and the image elements are annotated by the holistic module again. The symbolic module checks which of the candidate 3D models are consistent with the newly annotated map. This process may iterate and incrementally refine the map until a satisfactory recognition results. Once a 3D model is recognized, the model can be used to predict new images of the map. Then, the holistic module may confirm those predictions by active perception as discussed in the previous section.

Since the holistic processing is basically data driven, the interface between the holistic and symbolic processing should involve geometrical reasoning as it relates to active perception. Holistic processing starts by conjecturing about the 2D visual images, and about the geometrical transformations undergone by such images when matched against some canonical, standard representations (obtained when the prototypes were originally learned). Such information is passed and used by the symbolic processing, in triggering parts of the semantic networks (visual primitives) and the geometrical attributes characterizing them and also suggesting the symbolic (both topological and geometric) relationships connecting the visual primitives. Once a partial identification at the symbolic level is achieved, the needed guidance for further sensing and possible resegmentation and/or matching is provided to the holistic module. The interface we plan to develop will enable the system to perform active sensing, preprocessing and recognition.

3. Dynamic recognition

3.1 Basic Principles

Recognizing objects in a scene is one of the basic acts of human vision. Humans are able to recognize objects from limited and noisy inputs, and use a variety of clues. Despite numerous efforts to understand human vision, the flexibility and ease with which the human mind can handle recognition problems remains unexplained and mysterious.

To recognize an object is to assign a concept name to a subset of observational data. The basic assumption made in current recognition techniques is that in order to assign a concept to a stream of data, one needs to find a *match* between this stream and stored concept description. Production rules, prototypes, statistical distributions, grammars and logical expressions are among the types of concept descriptions that have been used. The basic weakness of the developed methods is that they do not explain how people can recognize the same concept in so many different ways. For example, one may be able to recognize a given person by seeing just a part of his/her face, or back, by hearing the person's voice, seeing handwriting, or even by observing the person's silhouette in the dark.

Each of these pieces of information can appear in all kinds of forms and manifestations, and thus the number of different subsets of data items that can trigger recognition is practically infinite. Clearly, it is completely impractical to have one rule/description for each subset of data that is sufficient to recognize an object. But this is basically how our current recognition systems work.

This project proposes a computational theory and a method that explains how a recognition system could recognize a given concept from many different subsets of information about it, but using only one concept description. Instead of matching the description with the input, the system performs a form of inductive inference that determines the properties sufficient for object recognition in a given context.

3.2 An Outline of the Theory of Dynamic Recognition

The main idea behind dynamic recognition is that the system uses a single *characteristic concept description* of an object or a concept (Michalski, 1983). Such a description contains all the available visual and non-visual information about objects that system learns about. A given new object is recognized not by directly matching this description with the input image, but by conducting inductive inference involving the characteristic description, the observable object properties, and the context in which the recognition occurs. Thus, recognition is viewed not as a process of matching one description with another, but as an inference process that determines sufficient clues needed for object recognition in a given context. The proposed dynamic recognition method involves three major stages: 1. *REDUCE*, 2. *INDUCE* and 3. *INQUIRE*.

Given some initial visual properties/facts about an object, these properties are used to "reduce" the set of classes (concepts) that are expected in a given context (the *REDUCE* stage). If the set is not reduced to a singleton (the object is not yet recognized), an inductive procedure (in our approach, the *AQ* or *INDUCE* program) is applied to candidate descriptions to determine the simplest *discriminant* recognition rules (the *INDUCE* stage). The difference between the standard inductive process and the described process is that input data are not single object descriptions but rules (characteristic descriptions). After *discriminant* rules are generated, measurements/observations are made that provide information that is needed to satisfy the discriminant description. At each step, the algorithm *INQUIRE* determines what piece of information would foster maximal reduction of uncertainty. A more detailed description of the algorithm is in Appendix 1.

Here, to illustrate our ideas, we will give a simple example of dynamic recognition. In the example we will use concept descriptions in the form of attribute-based rules. It should be mentioned, however, that the described principles apply as well to complex structural descriptions. Suppose the following characteristic concept descriptions are given:

[C=3] ::= [X1 =5] & [X3 =7] & [X8 =2..8]
[C=1] ::= [X3 =3] & [X4 >7] & [X6 =7]
[C=3] ::= [X1 =7 or 8] & [X3 =6 or 8] & [X8 =1..7].

The operator ::= states that the concept specified in the LHS is *characterized* by assertions in the RHS. The variable C takes values that name various concepts.

Case 1. No initial facts

Because there are no initial facts, the *REDUCE* stage consists of determining the concept that can be considered in the given context (e.g., when entering an office in a building, we do not expect to see there a live elephant). One of the sets of discriminant concept recognition rules produced at the *INDUCE* stage this is done by the *AQ* inductive learning program (Michalski, Mozetic, Hong and, Lavrac 1986) is:

[X3 =7] ==> [C =3]
[X3 =3] ==> [C =1]
[X3 =6 or 8] ==> [C =2].

The operator ==> is interpreted: the statement in the LHS is a sufficient condition to recognize the concept specified in the RHS.

The INQUIRE stage generates a question: What value does X3 have? If the answer is, for example, 3, then we have
TRUE ==> [C1 =1], and the system returns the answer 3 as the recognized concept.

Case 2. Given are initial facts:
X8 =7 ; X4 =3

The REDUCE stage produces the following reduced characteristic descriptions:
[X1 =5] & [X3 =7] ==> [C3]
[X1 =7 or 8] & [X3 =6 or 8] ==> [C =2].

These descriptions are input to an inductive program that determines the minimal costs discriminant descriptions (concept recognition rules). One of the two alternative discriminant descriptions produced by the INDUCE step is:
[X3 =7] ==> [C =3]
[X3 =6 or 8] ==> [C =2]

Another one is:
[X1 =5] ==> [C =2]
[X1 =7 or 8] ==> [C =2].

To decide which ruleset to select, the system uses costs associated with determining values of individual variables. The ruleset that contains variables with the smallest total costs is selected. For example, suppose that cost(x1) =10 and cost(X3) =1. The system selects the first ruleset.

The value of variable X3 is can be determined by making a measurement, by conducting an experiment, or by inferring it from known information. Once the value is known, the rule that fires indicates the concept.

3.3 Summary of Ideas and Proposed Research

The presented method of dynamic concept recognition replaces the traditional rule matching by an inductive process of discrimination. The whole process of concept recognition consists of three phases: REDUCE, INDUCE and INQUIRE. In the REDUCE step the context, and an a priori supplied information about an object to be recognized is used to determine a candidate set of hypotheses, and reduce the characteristic concept descriptions for these hypotheses. The second step, INDUCE, constructs goal-oriented discriminant concept descriptions from the reduced characteristic descriptions. The last step involved determining the information needed for final recognition.

We propose to apply these ideas to the problem of recognizing visual objects on the basis of various descriptors that involve shape characteristics, color and also texture (as determined by the method described in the section 4). We also want to study how to deal with only partial satisfaction of characteristic descriptions in the REDUCE step, and only partial satisfaction of discriminant recognition rules in the INQUIRY step, so that the recognition process may generate a certainty

degree of recognition. Another problem to study is how to deal with different "significance" of observed properties for recognition in the process of dynamic recognition.

4. Recognizing Objects by Surface Properties and Texture

Among the most informative properties that play a role in recognizing visual objects are their color and texture. Many, if not the most objects in one's familiar environment can be recognized on the basis of just these two properties, i.e., without information about their shape, size or other characteristics. While measuring color is relatively easy, "measuring" texture is difficult.

Researchers have attacked this problem in a number of ways. Main approaches used statistical methods or joint space/spatial-frequency representations. Much of the work on the statistical approach uses Gibbs Random Fields (GRF) (Geman and Geman, 1984; Cohen and Cooper, 1987; Goutsias, 1987). The GRF models are usually applied to images made up of known texture geometries and can be computationally quite expensive. Furthermore, they do not provide a straightforward derivation of perceptual primitives for describing textures, and do not reflect the present knowledge on how low-level vision is represented in terms of oriented banks of spatial/spatial-frequency filters. Generally, the accuracy of statistical methods is limited by the degree to which the assumed statistical model agrees with the actual distribution of data. Moreover, if a texture consists of a repeated pattern of structured objects, and the structural properties rather than statistical are significant, then such methods are not adequate.

We propose to work on an approach to surface and texture description that employs inductive learning techniques. These techniques are used for automatically extracting the most significant spacial properties of a surface and texture. This approach was originally proposed by one of the PIs (Michalski, 1973), and tested using ILLIAC III computer facilities. Early experiments produced very good results in discriminating between even very similar structural textures. Subsequently, this approach was applied to determine faults in laminates for aircraft wings using ultrasound images (Channic, 1988). Recently, Michalski's students Channic and Bala have implemented a simple version of a system, called TEXTPERT, that utilizes this approach. TEXTPERT is designed for recognizing objects by the properties of their surface and texture on the basis of two-dimensional digital images. Given a series of images classified by a human tutor into named surface and textural regions, TEXTPERT generates a procedure for classifying pixels into these regions. Such a procedure consists of a sequence of operators that transform any given texture into a uniform set of labels characterizing the individual texture type. The major step in the procedure is the formulation of rules characterizing spacial properties of a texture. This is done by the inductive learning program AQ15. To generate the whole procedure, the system first searches for the most discriminatory image quantization levels (the number of pixels and the number of grey levels), and then extracts a set of spacial texture samples (*events*) from different texture regions. These events are supplied to the learning program that formulates the rules. In the testing phase, the system applies the procedure to unclassified images and partitions them into different texture regions. The performance of the rules is measured by the degree of agreement between the system's and expert's classifications. If the procedure applied to the testing area produces a high degree of agreement with a classification done by an expert, then it is assumed that future images can be processed without the assistance of an expert. The system is using AQ15 machine learning module to generate the set of rules.

We will extend techniques used in TEXTPERT for general application in computer vision, and as a module in our vision system. We will investigate the usefulness of these techniques in recognizing visual objects. We will test them on problems of detecting important properties (e.g., faults) of materials. By detecting textures of materials we can also extend the range of other senses such as touch. A system capable of identifying materials at a distance might possess the capacity to infer the tactile sensations that it would experience while touching the distance objects. There are several important questions that arise when considering the use of TEXTPERT like techniques in

computer vision: what types of event operators (features) are the best for identifying different materials, how large a region must material occupy in the image for it to be correctly identified in the learning phase of TEXPert; and how can the system automatically choose the best operator for a given type of surface and texture. Main research topics include

1. acquiring surface and texture descriptors
2. applying inductive learning for constructing object recognition rules based on surface and texture properties, and
3. applying conceptual clustering for segmentation.

5. Conclusion

The research proposal addresses problems related to the application of modern machine learning and dynamic vision techniques to robot vision tasks. We propose to investigate a number of novel approaches for representing and learning of visual concepts, such as two-tiered representation, dynamic recognition, and a two-prong visual architecture that includes both holistic and symbolic representations. Another novelty of the proposed research is in the specific hierarchical architecture of the system, which allows for different abstraction levels and facilitates symbolic and holistic processing.

The proposed research will be closely connected to the Independent Explorer Project (IEX). The primary application of IEX project is the field of autonomous intelligent robotics, and the main idea is to build a system which will be able to acquire general and domain specific knowledge and will use this knowledge by deciding on its own which particular learning, planning, and inference strategies should be applied. We plan to design robotic system with the capability of safe navigation, the gathering, delivering, and sorting of materials and objects located in different places, and performing some simple assembly tasks. Building this system requires the integration of different modules that will be able to process dynamic knowledge acquired from changing environment within time constraints. One of the most important module of such a system is the intelligence computer vision module. We believe that domain of autonomous intelligent robotics will be the best testing ground for our research in applying machine learning to computer vision.

The ultimate goal of our research is to develop a robot vision system that is robust and capable of representing and learning imprecise and context-dependent concepts. The system should be robust in the sense that it has the ability to deal with an imperfect world, where the images can be noisy, the observed objects partly occluded, and/or have been geometrically transformed as compared to the training phase. Furthermore, the robot can learn to handle a world where objects can assume a large number of appearances and forms. It is one of the major goals of this proposal to explore the possibility of using distributed processing towards robust perceptual categorization and to interface between the PDP and symbolic processing. Achieving such a goal will be an important step towards better understanding of how to build practical computer vision systems, and will advance our present knowledge on intelligent and adaptive behavior.

REFERENCES

- Bajcsy, R., Active Perception, Proc.IEEE, Vol.76, 8, 996-1005,1988
- Ballard, D., Eye Movements and Spatial Cognition, TR 218,CS, University of Rochester, 1987.
- Baskin, A.B., and R. S. Michalski, R. S. , An Integrated Approach to Knowledge-based Systems: Experience with ADVISE and Related Programs, chapter in Expert system Design, 1988.
- Bentrup, J., Mehler, G., and Riedesel, INDUCE-4: a Program for Incremental Learning Structural Descriptions from Examples, ISG report 87-2, Department of Computer Science, University of Illinois, 1987.

- Bergadano, F., Matwin, S., Michalski, R. S., and Zhang, J., Representing and Acquiring Imprecise and Context-dependent Concepts in Knowledge-based Systems, Invited paper, Proceedings of the International Symposium on Methodologies for Intelligent Systems, Torino, 1988a.
- Bergadano, F., Matwin, S., Michalski, R. S., and Zhang, J., Measuring Quality of Concept Descriptions, Accepted for presentation to the Third European Working Session on Learning, 1988b.
- Besl, P.J., and Jain, R.C., Segmentation through variable-order surface fitting, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol.10, No.2, 167-192, 1988.
- Bratko, I., Mozetic, I., Lavrac, N. KARDIO: a study in deep and qualitative knowledge for expert systems. The MIT Press, Cambridge, MA (in press).
- Burt, P., Smart sensing with a pyramid vision machine, Proc. IEEE, Vol.76, 8, 1006-1015, 1988.
- Gallant, S.I., Connectionist expert systems, Comm of the ACM, Vol.31, No.2, 152-169, 1988.
- Hoff, W., Michalski, R. S., and Stepp INDUCE 2: A Program for Learning Structural Descriptions from Examples, ISG 83-4, UIUCDCS-F-83-903, Department of Computer Science, University of Illinois, Urbana, Jan. 1983.
- Jacobson, L., and Wechsler, H., FOVEA-A system for invariant visual form recognition, 2nd Int. Symposium on Optical and Electro-optical Applied Science and Engineering, Cannes, France, 1986.
- Kamgar-Parsi, B., Minimization of the quantization error in camera calibration, Proc. DARPA Image Understanding Workshop, 671-680, 1987.
- Klinker, G.D., Shafer, S.A., and Kanade, T., The measurement of highlights in color images, Int. Journal of Computer Vision, 2, 7-32, 1988.
- Ko, H. and Lee, K., Automatic Assembling Procedure Generation from Mating Conditions, Computer Aided Design, Vol 19, Num. 1, pp. 3-10, 1987
- Ko, H., Reasoning about Kinematics from Mating Conditions, ISG-16, UIUC-ENG-88-1708, Department of Computer Science, University of Illinois, Urbana, 1988.
- Ko, H., Integration of Planning and Learning for Robot Assembly, PhD thesis forthcoming, University of Illinois, Urbana, January 1989.
- Kohonen, T., Self-Organization and Associative Memories, Springer-Verlag, 1988.
- Koenderink, J. J., and van Dorn, A.J., Dynamic Shape, Biol. Cybernetics 53, 383-396, 1986.
- Michalski, R. S., AQVAL/1--Computer Implementation of a Variable-Valued Logic System VL₁ and Examples of its Application to Pattern Recognition, Proc. of the 1st International Joint Conf. on Pattern Recognition, Washington, D.C., pp. 3-17, 1973.
- Lippman, R. P., An introduction to computing with neural nets, IEEE ASSP, 4-22, 1987.
- Michalski, R. S., Carbonell, J. G., Mitchell, T. M., "Machine Learning: An Artificial Intelligence Approach", Tioga Publishing Co, Palo Alto, CA, 1983.
- Michalski, R. S., "A Theory and Methodology of Inductive Learning", Chapter in the book "Machine Learning, an Artificial Intelligence Approach", Michalski, R. S., Carbonell, J. G., Mitchell, T. M. (Eds.), Tioga Pub. Co., Palo Alto, CA, 1983.
- Michalski, R. S., Mozetic, Hong, J., I., Lavrac, The Multi-purpose Incremental Learning System AQ15 and its Testing Application to Three Medical Domains, Proc. 5th AAAI, pp. 1041-1045 1986.

- Michalski, R. S., and Stepp, R., INDUCE-3: A Program for Learning Structural Descriptions from Examples, ISG 86-9, UIUCDCS-F-86-960, Department of Computer Science, University of Illinois, Urbana, 1986.
- Michalski, R. S., Ko, H., and Chen, K. Qualitative Prediction: The SPARC/G Methodology for Describing and Predicting Discrete Processes, Current Issues in Expert Systems, A. Van Lamsweerde and P. Dufour (eds), Academic Press Incorporated, pp. 125-158, 1987.
- Michalski, R. S. Two-Tiered Concept Meaning, Inferential Matching and Conceptual Cohesiveness, in *Similarity and Analogical Reasoning*, S. Vosniadou and A. Ortony, (eds), Cambridge University Press 1988.
- Michalski, R.S. and Ko, H. On the Nature of Explanation or Why did the Wine Bottle Shatter, Proceedings of the Symposium on Explanation-based Learning, Stanford University, March 1988.
- Mozetic, I., Learning of qualitative models. in *Progress in Machine Learning* (I.Bratko, N.Lavrac, Eds.), pp. 201-217, Sigma Press, Wilmslow, UK, 1987a.
- Mozetic, I., The Role of Abstractions in Learning Qualitative Models. Proc. 4th Intl. Workshop on Machine Learning, pp.242-255, Irvine, CA, July 22-25, Morgan Kaufmann, 1987b.
- Mozetic, I., Bratko, I., Urbancic, T. Varying Level of Abstraction in Qualitative Modelling. In Machine Intelligence 12 (J.E.Hayes, D.Michie, E.Tyugu, Eds.), Oxford University Press (in press).
- Mozetic, I., Lavrac, N. Incremental Learning from Examples in a Logic-based Formalism. Proc.. Intl. Workshop on Machine Learning, Meta-Reasoning and Logics, pp. 109-127, Sesimbra, Portugal, February 14-17, 1988.
- Mozetic, I., Urbancic, T. Hierarchical qualitative modelling: a deep knowledge representation for efficient diagnostics. Proc. IASTED Intl. Symposium Expert Systems, pp. 162-166, Geneve, Switzerland, June 16-18, ACTA Press, 1987.
- Terzopoulos, D., Witkin, A., and Kass,M., Symmetry-Seeking Models for 3D Object Reconstruction,Proc.1st Int. Conference on Computer Vision, 1987.
- Warrington, E. K., Neuropsychological Studies of Object Recognition, Phil. Trans. of the Royal Soc. Lond. B 298, 15-33, 1982
- Wechsler, H., and Zimmerman,L, 2D Invariant Object Recognition Using Distributed Associative Memories,IEEE Trans. on Pattern Analysis and Machine Intelligence, Nov.1988(to appear).
- Wechsler,H., and Zimmerman,L.,Distributed Associative Memories for Bin-Picking, IEEE Trans. on Pattern Analysis and Machine Intelligence,1989 (to appear).
- Wechsler,H., and Zimmerman, L., Redundant vision for robotics,in Robots with Redundancy, A.K.Bajczy, Springer-Verlag,1989 (to appear).

PART II: Research Team and Research Enviroment

1. Accomplishments of PIs on Previous NSF ~~and DARPA Grants~~ During the Last Five Years

R. S. Michalski

NSF Grant: *Studies in Inductive Learning and Plausible Reasoning*
MCS-83-06614 (1982-1986)

Research on this grant was concerned with understanding the principles of inductive learning and plausible reasoning, and developing experimental programs implementing the developed methods. We conducted many specific projects on these topics and published over 100 papers reporting our results during the period of the grant. Beacuse we had also other support for research on closely related topics, the publications list joint support by NSF and other agencies. These publications are listed in Appendix 2 (years 82-86).

Major topics of our research included conceptual clustering, learning structural descriptions from examples, determining most relevant attributes for describing classes of objects, developing a large scale inference system with multiple knowledge representation and learning capabilities (ADVISE), the part-to-whole generalization, qualitative prediction, variable-precision logic (with P. Winston) and developing a cognitive theory of plausible reasoning (with A. Collins).

ACCOMPLISHMENTS FROM PREVIOUS NSF GRANTS

- (a) NSF/ECS-8105168 (University of Minnesota) originally awarded as ECS79-15216
(University of Wisconsin)
\$57,622
1980-1982
Texture Analysis

Summary

One of our major concerns in texture analysis was that of segmentation of textured imagery [1]. We also looked at the formal requirements of the task and found what is the computational complexity of the task. In particular, we were able to show that (i) there exists no algorithm that decides for any arbitrarily given predicate P whether or not P is a uniformity predicate and (ii) that there exists no algorithm that decides for any arbitrarily given uniformity predicate P whether or not the predicate induces more than one segment in any image [2].

Next, we look at the problems caused by looking at texture images in only one of the spatial or spatial-frequency domains. If the analysis is performed in any of those two domains taken by itself, there is poor separability of information and data items cluster together. Based on a synergetic approach, motivated by signal analysis, neurophysiology and psychophysics research, we were the first to suggest the use of cojoint image representations for image processing [3]. Specifically, we introduced the Wigner Distribution (WD), a joint/simultaneous representation of both space and spatial frequency. This was a major step and led to much further research and interest in the potential usefulness of such representations. One such application was directed toward object recognition [4].

References

- [1] H. Wechsler (1980), "Taxonomy and Segmentation of Textured Images", Proceedings of the 5th International Joint Conference on Pattern Recognition, Miami, Florida.
- [2] E. Gurari and H. Wechsler (1982), "On the Difficulties Involved in the Segmentation of Pictures", IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. PAMI-4, No. 3, 304-306.
- [3] L. Jacobson and H. Wechsler (1982), "The Wigner Distribution and its Usefulness for 2-D Image Processing", Proceedings of the 6th International Joint Conference on Pattern Recognition, Munich, West Germany.
- [4] L. Jacobson and H. Wechsler (1982), "A Paradigm for Object Recognition of Brightness, Optical Flow and Binocular Disparity Images", Pattern Recognition Letters, Vol. 1, No. 1, 61-68.

(b) NSF/ECS-8310057

\$118,624

1984-1986

Invariant Image Recognition of Brightness and Optical Flow

Summary

Within the framework provided by the Wigner Distribution (WD) we considered the issue of invariance. Specifically, we were interested to find out what invariants were available for pattern recognition and how to implement them. Toward that end we considered distortions caused by both linear transformations (LT) and perspective. The joint spatial/spatial-frequency representation provided by the WD was made invariant to LT and perspective by using conformal mappings and deprojections, respectively, [1,2]. A formal analysis of the intrinsic capabilities of the WD in terms of computation and resolution [3] was undertaken and the results showed that the WD can achieve the highest resolution among available joint representations [4,11].

Following our development and analysis of invariant joint representations we considered applications to real problems like texture analysis and segmentation [6] and invariant object recognition [12]. A system called FOVEA-for invariant visual form recognition was developed [8] and preliminary results show that it can indeed cope with both distorted and occluded objects.

Finally, we extended the concept of simultaneous spatial/spatial-frequency representations to include the temporal dimension as well [9]. In particular, we were interested to derive optical flow which characterizes temporal changes, i.e., how object change position. Toward that end, we were able to make our method consistent with the coherence phenomena.

The work performed under this grant led us to better understand the capabilities of invariant joint representations and their potential applications toward real tasks. We were also in the position to be invited to survey the problem of invariant pattern recognition, its state-of-the-art status and to suggest directions for future research [10], and to show the use of distributed associative memories (DAM) [12] for invariant pattern recognition tasks.

References

[1] L. Jacobson and H. Wechsler (1984), "A New Theory for Invariant Object Recognition in the Frontoparallel Plane", IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. PAMI-6, No. 3, 325-331.

[2] L. Jacobson and H. Wechsler (1984), "Invariant Analogical Image Representations and Pattern Recognition", Pattern Recognition Letters, Vol. 2, 289-299.

[3] L. Jacobson, U. Thomanschefsky, and H. Wechsler (1984), Resolution and Invariance Criteria for Object Pattern Recognition, SPIE/Cambridge Symposium on Intelligent Robots and Computer Vision, Boston, MA.

[4] L. Jacobson and H. Wechsler (1985), Joint Spatial/Spatial-Frequency Representations for Image Processing, SPIE/Cambridge Int. Conference on Intelligent Robots and Computer Vision, Boston, MA.

[5] L. Jacobson and H. Wechsler (1984), "Multipurpose Low-Level Image Processing", Proceedings of the 7th Int. Conference on Pattern Recognition, Montreal, Canada.

[6] H. Wechsler (1984), "Invariant Local Spectral Signatures for Image Processing", Proceedings of the 7th Int. Conference on Digital Signal Processing, Florence, Italy.

[7] L. Jacobson and H. Wechsler (1984), "A Path Toward Solving the Bin-Picking Problem", Proceedings of the Int. Conference on Robotics, Atlanta, Georgia.

[8] L. Jacobson and H. Wechsler (1985), FOVEA-A System for Invariant Visual Form Recognition, 2nd Int. Technical Symposium on Optical and Electro-Optical Applied Science and Engineering, Cannes, France.

[9] L. Jacobson and H. Wechsler (1987), "Derivation of Optical Flow Using Spatiotemporal Frequency Analysis", Computer Vision, Graphics, and Image Processing, 38, 29-65.

[10] H. Wechsler (1987), INVARIANCE IN PATTERN RECOGNITION in Advances in Electronics and Electron Physics, P. W. Hawkes (Ed.), Academic Press.

[11] L. Jacobson and H. Wechsler (1988), Joint spatial/spatial-frequency representations, Signal Processing, Vol. 14, No. 1, 37-68.

[12] H. Wechsler and L. Zimmerman (1988), 2-D Invariant Object Recognition Using Distributed Associative Memory, IEEE Trans. on Pattern Analysis and Machine Intelligence, (in print).

(c) NSF/EET-8713563
\$29,983
1987-1988
Fault-Tolerant Recognition

We considered the problem of invariant object recognition using parallel distributed processings (PDP) models of computation. Our research efforts led to the development of an experimental system for the bin-picking problem [1,2]. The bin-picking problem, highly relevant to industrial automation, aims at recognizing noisy, partially distorted and occluded objects. The thrust of our work is towards redundant and robust performance for robotic vision [3]. There is much analogy between pattern recognition and informational retrieval. We have applied similar PDP concepts to relational data base retrieval and inference [4].

References

- [1] H. Wechsler and L. Zimmerman (1989), Distributed Associative Memories for Bin-Picking, IEEE Trans. on Pattern Analysis and Machine Intelligence (to appear).
- [2] H. Wechsler and L. Zimmerman (1988), Distributed Associative Memories and Data Fusion, 2nd Int. Conference on Neural Networks.
- [3] H. Wechsler and L. Zimmerman (1989), Redundant Vision for Robotics, in Robotics with Redundancy, A. K. Bejczy (Ed.), Springer-Verlag.
- [4] J. M. Char, V. Cherkassky, H. Wechsler, and L. Zimmerman (1988), Distributed and Fault-Tolerant Computation for Retrieval Tasks Using DAMs, IEEE Trans. on Computers, Vol. 37, No. 4, 484-490.

2. Biographical Sketches of Investigators

RESUME

Ryszard Spencer Michalski

Planning Research Corporation Professor
of Information Technology and Computer Science
Director of the Center for Artificial Intelligence
George Mason University
Fairfax, VA

Address:

George Mason University
Artificial Intelligence Center
School of Information Technology
4400 University Drive
Fairfax, VA 22030

Phone:

Office: 703/323-2713
Home: 703/764-9142

Citizenship: USA

Areas of Interest:

Machine Learning
Cognitive Models of Inference
Integrated AI Systems
AI Applications

Plausible Reasoning
Expert Systems
Logic Programming
Autonomous Robots

Biographical Sketch:

Dr. Michalski was born and educated in Poland. Before moving to the United States in 1970, he was a research scientist at the Polish Academy of Sciences in Warsaw. He conducted there research in the areas of computer logic design, switching and automata theory, computer vision, and threshold logic systems.

From 1970 until 1987 Dr. Michalski was a Professor of Computer Science and Medical Information Science at the University of Illinois at Urbana-Champaign. He was engaged there in a variety research projects in areas of artificial intelligence and pattern recognition, concentrating on methods of machine learning and inference, and their applications to medicine, agriculture, and recently also to autonomous intelligent robotics. In 1988 he joined the faculty of Computer Science at George Mason University, and became director of the Center for Artificial Intelligence.

Dr. Michalski is one the founders of the field of machine learning. His Laboratory is one of the leading centers of research in this field. He has developed a theoretical basis for inductive learning, and with his collaborators has implemented a series of inductive learning programs. These programs include AQ and INDUCE for concept learning, CLUSTER for conceptual clustering, SPARC for qualitative prediction, and ABACUS for quantitative discovery. These programs have been applied to a wide spectrum of domains by various organizations in the

United States and abroad. Simplified versions of these programs are demonstrated at the exhibition "ROBOTS and BEYOND: Age of Intelligent Machines," which is currently touring major Science Museums in USA.

In 1971 Dr. Michalski introduced the concept of variable-valued logic, a generalization of multiple-valued logic, which was the basis for the development of numerous machine learning programs.

In 1976 he started research on the application of artificial intelligence to agriculture. This research produced the first agricultural expert system, PLANT/ds, for the diagnosis of soybean diseases. It was developed in collaboration with Professor Sinclair, a world authority on soybean diseases. It was the first expert system that learned its diagnostic rules from examples. The inductively learned rules outperformed rules acquired by the conventional method of interviewing an expert. PLANT/ds is now in regular use by plant pathologists.

In 1980 Dr. Michalski originated the research on *conceptual clustering*, a novel approach to clustering and building classifications, which has now become an active area of AI research. Since 1982 he has collaborated with cognitive psychologist Allan Collins from Bolt, Beranek and Newman, Inc., on the development of a theory of human plausible inference. A major paper on this theory will appear in *Cognitive Science Journal*.

In 1985, Michalski jointly with Patrick Winston from MIT, introduced the concept of *variable precision logic*, which combines ideas of nonmonotonic logic with multiple-valued logic, and offers a novel approach to inference under time and resource limitations.

Dr. Michalski has lectured extensively worldwide, and held Visiting Professorships at various Universities in the United States, Great Britain and France.

Dr. Michalski co-organized the 1st International Machine Learning Workshop in Pittsburgh in 1980, and co-organized subsequent workshops/conferences held in 1983, 1985, 1986 and 1988.

Dr. Michalski edited (with Drs. Carbonell and Mitchell) three books in the field of machine learning, *Machine Learning: An Artificial Intelligence Approach*, Vols. I and II, and *Machine Learning: Guide to Current Research*, which have received wide acclaim in the field of machine learning. He authored or co-authored over 130 research and technical publications in his areas of interest. He is an editor of *Machine Learning Journal*, the basic journal in his field.

Education:

Institute and Location	Years	Degree
Polytechnical University of Cracow, Poland	1954-56	Certificate
Polytechnical University of Warsaw, Poland	1956-59	B.S.
Leningrad Polytechnical Institute, USSR	1959-61	M.S.
Advanced Course in Computer Science, Leningrad Polytechnical Institute, USSR	1961	Certificate
Polytechnical University of Silesia, Poland	1967-1969	Ph.D.

Professional Experience:

Planning Research Corporation Professor of Information Technology and Computer Science Director of Artificial Intelligence Center George Mason University	1988 - present
Director of Artificial Intelligence Laboratory Department of Computer Science, University of Illinois, Urbana	1983 - 1987
Professor of Computer Science, Department of Computer Science, University of Illinois,	1982 - 1987
Professor of Medical Information Science, College of Medicine, University of Illinois,	1982 - 1987
Visiting Scientist, Artificial Intelligence Laboratory, MIT, Cambridge, Massachusetts	Fall 84 - Spring 85
Visiting Scientist and Guest Lecturer Boeing Computer Services, Seattle, Washington	Summer 83, 84 & 85
Visiting Scientist, Artificial Intelligence Laboratory, MIT, Cambridge, Massachusetts	May - June 1983
Associate at the Center for Advanced Study, University of Illinois, Urbana	1982 - 1983
Associate Professor of Medical Information Science, University of Illinois, Urbana	1979 - 1982
Associate Professor of Computer Science, Department of Computer Science, University of Illinois, Urbana	1972 - 1982
Visiting Professor, Universite de Paris IX (Dauphine) Paris, and INRIA Research Institute Rocquencourt, Le Chesnay, France	Summer 1980
Visiting Research Scientist, Bolt, Beranek and Newman, Inc., Cambridge, Massachusetts	Summer 1978
Visiting Professor at the University of Essex and the University of London, Queen Mary College U.K.	Fall 1978
Visiting Professor at the Department of Computer Science at Carnegie-Mellon University, Pittsburgh,	Spring 78
Visiting Professor at the Department of Computer Science at the University of Wisconsin, Madison	Summer 77

Visiting Professor at the University of Illinois at Chicago Summers 75, 76

Visiting Research Assistant Professor of Computer Science, Department of Computer Science, University of Illinois, Urbana 1970 - 72

Research Scientist, and then the Head of the Pattern Recognition Group, Institute of Automatic Control, Polish Academy of Sciences, Warsaw 1962 - 70

Lecturer in Computer Science and Electronics, State Technical College 1964 - 68
Technical Scientific Organization, Warsaw

Research Scientist, Computer Design Department 1961 - 62
Institute of Mathematical Machines,
Polish Academy of Sciences, Warsaw

Honors and Past Grants:

Member of the USA Delegation to the first US-Japan Symposium on Artificial Intelligence, Tokio, November 30-December 3, 1987.

Elected to Executive Council of the American Association of Artificial Intelligence, 1985.

International Chair in Informatics, University of Mons, Belgium, Spring 1985.

Member of the USA Delegation to the First US-China Joint Seminar on Automation and Intelligent Systems, Beijing, China, May 28-June 1, 1984.

Associate at the Center for Advanced Study, University of Illinois, 1982-1983 (released from teaching for a year of advanced research).

Distinguished Speaker of the IEEE Computer Society, 1976-1981.

Listed in *Who's Who in Frontiers of Science and Technology*.

Included in *American Men and Women of Science*.

Included in *Who's Who in Technology Today*.

Selected for the *International Book of Honor*.

Co-Principal Investigator of an ARO/DoD Equipment Grant *Research in Artificial Intelligence: Qualitative Reasoning, Machine Learning, and Mechanized Induction* (with K. Forbus and L. Rendall), 1986-87 (\$298,000).

Principal Investigator of the project *Machine Learning and Discovery*, DARPA grant, 1985-1988.

Principal Investigator of the project *Studies in Computer Inductive Learning and Plausible Inference*, NSF Grant No. DCR 84-06801, 1984-1987.

Co-Principal Investigator (with Prof. T. Fermanian) of a grant to develop a "Turf Advisor" expert system from the Research Board and Agricultural Experiment Station, University of Illinois, 1985-1986 (\$20,000).

An Investigator of the *Center of Excellence* equipment grant from the National Science Foundation to the Department of Computer Science, University of Illinois at Urbana-Champaign, 1981 - 1986.

Principal Investigator of an NSF Grant for Artificial Intelligence Equipment, MCS 83-04913, 1983-1985.

Principal Investigator of the project *Studies in Machine Inductive Learning and Plausible Reasoning*, (in collaboration with Allan Collins and David Kuck,) NSF Grants No. MCS-82-05166, 1982-1984.

Principal Investigator of the project *Development of Expert Systems with Multiple Knowledge Representation and Learning Capabilities*, (with collaborators A. B. Baskin and R. Stepp), ONR Grant No. N00014-82-K-0186, 1984-1987.

Principal Investigator of NSF grants in the area of machine learning and inductive inference, MCS-74-03514 (1974-76), MCS-76-22940 (1977-79), MCS-79-06614 (1979-82) and MCS-83-06614 ((1982-1986).

Co-Principal Investigator (with Prof. A. B. Baskin, School of Medical Sciences and Professor Robert Stepp, Electrical and Computer Engineering) of the project *Integrating Multiple Knowledge Representations and Learning Capabilities in an Expert System*, ONR Grant No. N00014-81-K-0186, 1981-1984.

Co-Principal Investigator of the project *Knowledge-Based Programming Assistant*, IBM Grant, 1981-1983.

Co-Principal Investigator (with Prof. J. Sinclair and Prof. W. Ruesink) of the project on the *Applications of Expert Systems to Agriculture*, USDA, Entomology and Pest Losses, 1979-1982.

On the list of Teachers Ranked as Excellent by Their Students, Fall 1976.

Senior Visiting Fellowship from the British Science Research Council, Fall 1977.

Fullbright Fellowship, 1970 (not used).

Current Grants:

Principal Investigator of the project *Machine Learning: A Multistrategy Approach* (with Ken DeJong of George Mason University as Co-principal Investigator), ONR Grant No. N00014-88-K-0186, 1988-1991.

Principal Investigator of the project *Integrated Machine Learning, Inference and Discovery*, (faculty collaborators: Professors A. B. Baskin, S. Lu, G. Collins and R. Stepp), DARPA 1987-1990.

Principal Investigator of the project *Cognitive Models of Learning and Inference* (with collaborators Doug Medin and Maria Zemankova), Office of Naval Research, 1988-1990.

Editorships:

Editor (with Jaime Carbonell, Tom Mitchell and Patrick Langley) of the Journal *Machine Learning*

Editor of *Reports of Machine Learning and Inference Laboratory*, George Mason University (until 1988, *Reports of Intelligent Systems Group*, University of Illinois at Champaign-Urbana)

Member, Editorial Board of *Research Notes in Artificial Intelligence*

Member, Editorial Board of the *International Journal of Policy Analysis and Information Systems*

Member, Editorial Board of the *International Journal of Intelligent Systems*

Member, Editorial Board of the *Future Generations Computer Systems*

Major Organizational Functions in the Last Several Years:

Member of the George A. Miller Lecture Committee at the University of Illinois, 1982 - 1984 (an interdisciplinary committee to invite distinguished scholars).

Co-Organizer of the First International Workshop on Machine Learning, Carnegie Mellon University, June 16-18, 1980, Pittsburgh, PA.

Organizer, General Chairman and Speaker, Second International Workshop on Machine Learning, Allerton Park, University of Illinois, Urbana, IL, June 22-24, 1983.

Organizer of a Presentation on Artificial Intelligence Research at the University of Illinois "Thinking and Learning" (with participating scientists: D. Medin, B. Ross, A. B. Baskin, D. Plaisted, M. Harandi and C. Sammut), October 10, 1984.

Co-Organizer and Speaker, Third International Machine Learning Workshop, Skytop Lodge, Rutgers University, NJ, June 24-26, 1985.

Co-Organizer and Speaker of the International Meeting on Advances in Learning, (IMAL), Les Arcs, France, July 28-August 1, 1986.

Co-Organizer of the Fourth International Machine Learning Workshop, Irvine, June 22-25, 1987.

Member of the Program Committee for the International Symposium on Methodologies of Intelligent Systems, University of North Carolina, October 14-17, 1987.

Member of the Program Committee for the Conference on Applications of Artificial Intelligence, November 1987.

Member of the Organizing Committee and the Program Committee of the 5th International Conference on Machine Learning, University of Michigan, Ann Arbor, June 12-14, 1988.

Other Professional or Public Services:

Staff Search and Initial Screening Committee
Fellowships, Assistantships and Admissions
George Miller Committee, 1982-1984
Piano Recitals for Senior Citizens and Other Groups

Foreign Languages:

Polish and Russian (fluent)

Other Interests:

Cognitive psychology, skiing (former ski instructor), hiking, tennis, piano.

Selected Professional Activities for the Last Three Years:

Invited guest and speaker at the Warsaw University, and Jagiellonian University (Cracow), July 24 - August 8, 1988.

Lecturer, Summer School on Machine Learning, Les Arces, France, July 18-23, 1988.

Invited speaker, Conference on the Management of Uncertainty, Urbino, Italy, July 4-7, 1988.

Invited speaker, Communications Network Seminar Series, Department of Computer Science, University of North Carolina at Charlotte and Microelectronics Center of North Carolina, Charlotte, N.C. April 28, 1988

Invited speaker, Sigma Xi Society Meeting, University of North Carolina at Charlotte, Charlotte, N.C., April 27, 1988.

Invited speaker, Distinguished Lecture Series, The University of Toledo, College of Engineering, Toledo, April 7, 1988 (two lectures).

Invited speaker at the International Conference on Information Systems, Italian Department of Communication, Rome, March 10-12, 1988.

Member of the American delegation to the first USA-Japan Symposium on Artificial Intelligence, November 30-December 3, 1987, Tokio, Japan (co-sponsored by NSF and ICOT, MITI).

Invited Lecture at the 2nd International Conference on Methodologies for Intelligent Systems, University of North Carolina at Charlotte, October 14-18, 1987.

Invited Lecture at the Workshop on Knowledge Representation and Organization in Machine Learning, Technical University of Berlin, Schloss Eringerfeld, Geseke, West Germany, October 5-9, 1987.

Tutorial Presentation on Machine Learning, 10th International Joint Conference on Artificial Intelligence, Hotel Milano Fiori, University of Milano, Milano, August 23-28, 1987.

Invited Plenary Lecture at the Second International Conference on Artificial Intelligence, Institut International de Robotique et d'Intelligence Artificielle, Marseilles, December 3-5, 1986.

Invited Speaker, Knowledge-Based Systems Workshop, Banff, Canada, Nov. 3-7, 1986

Invited Speaker, International Symposium on Methodologies for Intelligent Systems, Knoxville, TN, October 21-25, 1986

Invited Speaker, International Meeting on Advances in Learning, Les Arcs, France, July 28-August 1, 1986

Session Chairman at AAAI-86 Conference, Philadelphia, PA, August 10-15, 1986.

Invited Speaker, Allerton Conference on Analogy and Similarity, June 1986

Consultant for the National Security Agency and other governmental agencies on the state-of-the-art in machine learning, June 1986.

Invited Speaker, Two lectures (public and technical), Northern Illinois University, DeKalb, IL, April 1986.

Colloquium Speaker, Department of Computer Science, Duke University, Durham, NC, March 12, 1986

Invited Speaker, Department of Computer Science, Yale University, New Haven, CT, February 10, 1986.

Colloquium Speaker, Department of Psychology, University of Illinois at Urbana-Champaign, January 30, 1986.

Invited Speaker, Conference on Machine Learning, sponsored by the Knowledge-Based Systems Centre at the Polytechnic of the South Bank, Bloomsbury Crest Hotel, London, November 11-13, 1985.

Presentation at the Navy-wide Technical Review of AI Research and Development, National Academy of Sciences, Washington DC, September 17-18, 1985.

Invited Speaker, Knowledge Systems and Logic Programming, IBM Europe Institute, Oberlech, Austria, July 29 - August 2, 1985.

Organizer and Lecturer, Advanced Two-Week Course on Knowledge Representation, Boeing Computer Services, Seattle, WA, July 8-19, 1985.

Guest Lecturer, Aiken Science Center, Harvard University, Cambridge, MA, June 26, 1985.

Chairman of a Panel Discussion and Speaker at the Third International Machine Learning Workshop, Skytop Lodge, Rutgers University, NJ, June 24-26, 1985.

Two presentations at the Artificial Intelligence Laboratory, MIT, Cambridge, MA, April 26, 1985 and May 3, 1985.

Two invited lectures at McGill University, Montreal, Canada, April 9 and 10, 1985.

Invited Speaker at the Machine Intelligence Workshop 11, Ross Priory, University of Strathclyde, March 26-29, 1985.

Seminar Speaker at the Imperial College, London, March 19, 1985.

Invited Lecturer, IBM Arthur K. Watson International Education Centre, European Systems Research Institute, La Hulpe, Belgium, March 15, 1985.

International Professorship in Computer Science Program, Universite' de l'E'tat MONS, Faculte's Universitaires NAMUR, March 3-16, 1985 (a series of lectures).

Presentation on Machine Learning and Artificial Intelligence at DARPA, Washington, DC, December 11, 1984.

Guest Lecturer, General Electric Research and Development Center, Schenectady, NY, December 6, 1984.

Guest Lecturer, Electrical and Computer Engineering Department, MIT, Cambridge, MA, November 19, 1984.

Invited Speaker at the Artificial Intelligence Revolving Seminar, Artificial Intelligence Laboratory MIT, Cambridge, MA, November 14, 1984.

Invited Speaker at the International Conference on Artificial Intelligence, Dubrovnik, Yugoslavia, August 27 - September 3, 1984.

Invited Speaker at the European Workshop on Expert Knowledge, Bled, Yugoslavia, August 22-24, 1984.

Member of the Panel Discussion on Machine Learning, National Conference on Artificial Intelligence (AAAI-84), Austin, TX, August 6-10, 1984.

Seminar Lecturer, Department of Psychology, University of Washington, Seattle, WA, July 26, 1984.

Organizer and Lecturer, Advanced Two-Week Course on Machine Learning, Boeing Computer Sciences, Seattle, WA, July 6-20, 1984.

Member of the American Delegation for the First US-China Joint Seminar on Automation and Intelligent Systems, Beijing, China, May 28 - June 1, 1984.

Guest Lecturer, Electrotechnical Institute, Science City, Tokyo, Japan, May 25, 1984.

Guest Lecturer, NIPPON Electric Company (NEC), Tokyo, Japan, May 24, 1984.

Visitor to ICOT (Institute for the New Generation of Computer Technology), Tokyo, Japan, May 23, 1984.

Invited Guest at the *Unification Workshop* organized by IBM Los Angeles Research Center, February 27-29, 1984.

Visited Stanford University, Stanford Research Institute, and IBM San Jose Research Center, February 15-19, 1984.

Selected Courses Taught:

Computer Inference and Knowledge Acquisition
 Pattern Recognition and Machine Learning
 Information Systems
 Data Structures
 Logic Programming
 Introduction to Artificial Intelligence
 Computer Programming

PUBLICATIONS

Books Authored, Edited or Co-Edited:

R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach*, TIOGA Publishing Company, Palo Alto, CA, 1983.

R. S. Michalski (Ed.), *Proceedings of the International Machine Learning Workshop*, University of Illinois Allerton House, Urbana, IL, June 22-24, 1983.

R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach*, Vol II, Morgan Kaufman Publishers, Inc., Los Altos, CA, 1986.

T. Mitchell, J. Carbonell, and R. S. Michalski (Eds.), *Machine Learning: Guide to Current Research*, Kluwer Publishing Co., 1986.

Y. Kodratoff and R. S. Michalski (Eds.), *Machine Learning: An Artificial Intelligence Approach*, Vol. III (in preparation).

R. S. Michalski, *An Introduction to Machine Inference and Learning*, (in preparation).

Chapters in Books:

R. S. Michalski, Problems of Computer Simulation of Pattern Recognition Systems and a Description of System Recognizing Alphanumeric Characters, *Learning Automata and Biological Processes of Perception*, Polish Academy of Sciences, pp. 152-182, Warsaw, October 1966 (in Polish).

R. S. Michalski, A Variable-Valued Logic System as Applied to Picture Description and Recognition, *Graphic Languages*, F. Nake and A. Rosenfeld (Eds.), North-Holland Publishing Co., New York, pp. 20-47, 1972.

R. S. Michalski, Learning by Inductive Inference, *Computer Oriented Learning Processes*, J. C. Simon (Ed.), NATO Advanced Study Institutes Series, Series E, No. 14, pp. 321-337, 1976.

R. S. Michalski and P. Negri, An Experiment on Inductive Learning in Chess End Games, *Machine Representation of Knowledge*, Machine Intelligence 8, E. W. Elcock and D. Michie (Eds.), Ellis Horwood Ltd., New York, pp. 175-192, 1977.

R. S. Michalski, R. Stepp and E. Diday, A Recent Advance In Data Analysis: Conceptual Clustering, *Progress in Pattern Recognition*, Vol. 1, L. Kanal and A. Rosenfeld (Eds.), pp. 33-55, 1981.

R. S. Michalski and R. L. Chilausky, Knowledge Acquisition by Encoding Expert Rules versus Computer Induction From Examples: A Case Study Involving Soybean Pathology, *Fuzzy Reasoning and its Applications*, 1981.

R. S. Michalski and R. Stepp, *Revealing Conceptual Structure in Data by Inductive Inference*, *Machine Intelligence 10*, D. Michie, J. E. Hayes and H-H Pao (Eds.), Ellis Horwood Ltd., New York, pp. 173-196, 1982.

R. S. Michalski and R. Stepp, Learning From Observation: Conceptual Clustering, *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, CA, pp. 331-363, 1983.

R. S. Michalski, A Theory and Methodology of Inductive Learning, *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, CA, pp. 83-134, 1983.

T. Dietterich and R. S. Michalski, A Comparative Review of Selected Methods for Learning from Examples, *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, CA, pp. 41-81, 1983.

J. Carbonell, R. S. Michalski and T. Mitchell, An Overview of Machine Learning, *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, CA, pp. 3-23, 1983.

R. S. Michalski, Inductive Learning as Rule-Guided Generalization of Symbolic Descriptions: A Theory and Implementation, *Automatic Program Techniques*, Macmillan Publishing Company, New York, pp. 517-552, 1984.

R. S. Michalski, Generalization of Symbolic Descriptions: A Theory and Implementation, chapter in *Automatic Program Techniques*, Macmillan Publishing Company, New York, pp. 517-552, 1984.

R. S. Michalski, Understanding the Nature of Learning: Issues and Research Directions, *Machine Learning: An Artificial Intelligence Approach Vol. II*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), Morgan Kaufman Publishers, Los Altos, CA, pp. 3-25, 1986.

R. S. Michalski and R. Stepp, Conceptual Clustering: Inventing Goal-Oriented Classifications of Structured Objects, *Machine Learning: An Artificial Intelligence Approach Vol. II*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), Morgan Kaufman Publishers, Los Altos, CA, pp. 471-498, 1986.

T. Dietterich and R. S. Michalski, Learning to Predict Sequences, *Machine Learning: An Artificial Intelligence Approach Vol. II*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), Morgan Kaufman Publishers, Los Altos, CA, pp. 63-106, 1986.

R. S. Michalski, S. Amarel, D. Lenat, D. Michie, and P. Winston, *Machine Learning: Challenges of the 80's*, (Edited transcripts of a panel discussion), *Machine Learning: An Artificial Intelligence Approach Vol. II*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), Morgan Kaufman Publishers, Los Altos, CA, pp. 27-41, 1986.

R. E. Reinke and R. S. Michalski, Incremental Learning of Decision Rules: A Method of Experimental Results, in *Machine Intelligence 11*, J. E. Hayes, D. Michie and J. Richards (Eds.), Oxford University Press, 1987, in press.

R. S. Michalski, Learning Strategies and Automated Knowledge Acquisition: An Overview, *Computational Methods of Learning*, L. Bolc (Ed.), Springer-Verlag, 1987, in press.

R. E. Reinke and R. S. Michalski, Incremental Learning of Decision Rules, to appear in *Machine Intelligence 11*, J. E. Hayes, D. Michie and J. Richards (Eds.), Oxford University Press, In press, 1987.

R. S. Michalski, H. Ko and K. Chen, Qualitative Process Prediction: A Method and Program SPARC/G, *Expert Systems*, C. Guetler, (Ed.), Academic Press Inc., London, 1987.

R. S. Michalski, Two-tiered Concept Meaning, Inferential Matching and Conceptual Cohesiveness, in *Analogy and Similarity*, S. Vosniadou and A. Ortony (eds.), Cambridge University Press, 1988 (to appear).

Monographs:

J. Karpinski and R. S. Michalski, A Digital System Recognizing Alphanumeric Characters: General Mathematical Concepts and Experimental Data from an Operational Computer Simulator, *Reports of the Institute of Automatic Control*, No. 35, Polish Academy of Sciences, Warsaw, 89 pp., 1966 (in Polish).

R. S. Michalski, Variable-Valued Logic and Its Applications to Pattern Recognition and Machine Learning, Chapter in the monograph, *Computer Science and Multiple-Valued Logic Theory and Applications*, D. C. Rine (Ed.), North-Holland Publishing Co., New York, pp. 506-534, 1975.

Articles in Refereed Journals and Proceedings:

R. S. Michalski, A Systematic Procedure for Obtaining The Minimal Form of a Switching Function by the Local Extraction of Prime Implicants, *Papers of the IV National (Polish) Conference on Automatic Control*, Vol. II, pp. 29-42, Krakow, June 20-24, 1967 (in Polish).

R. S. Michalski, Some Problems in Design and Application Possibilities of Reading Machines, *Proceedings of the II National (Bulgarian) Conference on Automatic Control*, Vol. III, pp. 81-90, Bulgaria, Varna, September 21-25, 1967 (in Russian).

R. S. Michalski, On the Problems of Construction of Reading Machines, *Mathematical Machines*, No. 4, pp. 9-13, April 1968 (in Polish).

R. S. Michalski, Recognition of Total or Partial Symmetry in a Completely or Incompletely Specified Switching Function, *Proceedings of the IV Congress of the International Federation on Automatic Control (IFAC)*, Vol. 27, pp. 109-129, (Finite Automata and Switching Systems), Warsaw, June 16-21, 1969.

R. S. Michalski, Algorithm A^Q for the Quasi-Minimal Solution of the General Covering Problem, *Archiwum Automatyki i Telemekhaniki*, No. 4, Polish Academy of Sciences, 1969 (in Polish).

R. S. Michalski, On the Quasi-Minimal Solution of the General Covering Problem, *Proceedings of the V International Symposium on Information Processing (FCIP 69)*, Vol. A3 pp.125-128, (Switching Circuits), Bled, Yugoslavia, October 8-11, 1969.

R. S. Michalski, Covering Problems in the Structural Synthesis of Finite Automata, *Journal of the Silesia Technical University 'Automatyka'*, No. 14, pp. 229, Gliwice, 1969 (in Polish).

R. S. Michalski, Automatic Synthesis of the Quasi-Minimal Multiple-Output Switching Circuits, *VI International Symposium on Information Processing (FCIP 70)*, pp. 1-7, Bled, Yugoslavia, September 23-26, 1970.

R. S. Michalski and B. H. McCormick, Interval Generalization of Switching Theory, *Proceedings of the Third Annual Houston Conference on Computer and System Science*, Houston, Texas, April 26-27, 1971.

R. S. Michalski, A Method for Quasi-Minimal Solution of the General Covering Problem and Its Application to the Synthesis of Automata, *Avtomatika i Telemekhanika*, No. 1, pp. 118-130, 1972 (in Russian).

R. S. Michalski and Z. Kulpa, A System of Programs for the Synthesis of Switching Circuits Using the Method of Disjoint Stars, *Information Processing 71*, pp. 61-65, North-Holland Publishing Co., New York, 1972.

R. S. Michalski, Discovering Classification Rules Using Variable-Valued Logic System, *Proceedings of the Third International Joint Conference on Artificial Intelligence*, pp. 162-172, Stanford, CA, August 20-23, 1973.

R. S. Michalski, AQVAL/1--Computer Implementation of a Variable-Valued Logic System VL₁ and Examples of its Application to Pattern Recognition, *Proceedings of the First International Joint Conference on Pattern Recognition*, pp. 3-17, Washington, DC, October 30 - November 1, 1973.

R. S. Michalski, Variable-Valued Logic: System VL₁, *Proceedings of the 1974 International Symposium on Multiple-Valued Logic*, pp. 323-346, West Virginia University, Morgantown, May 29-31, 1974.

R. S. Michalski, A Variable Decision Space Approach for Implementing a Classification System, *Second International Joint Conference on Pattern Recognition*, pp. 71-75, Copenhagen, Denmark, August 13-15, 1974.

R. S. Michalski, Problems of Designing an Inferential Medical Consulting System, *First Illinois Conference on Medical Information Systems*, pp. 151-157, Urbana, IL, October 17-18, 1974.

R. S. Michalski, Synthesis of Optimal and Quasi-Optimal Variable-Valued Logic Formulas, *Proceedings of the 1975 International Symposium on Multiple-Valued Logic*, pp. 76-87, Bloomington, IN, May 13-16, 1975.

R. Chilausky, B. Jacobsen and R. S. Michalski, An Application of Variable-Valued Logic to Inductive Learning of Plant Disease Diagnostic Rules, *Proceedings of the 1976 International Symposium on Multiple-Valued Logic*, Utah State University, Logan, Utah, May 25-28, 1976.

J. B. Larson and R. S. Michalski, Inductive Inference of VL Decision Rules, Invited paper for the Workshop in Pattern-Directed Inference Systems, Hawaii, May 23-27, 1977 and published in *SIGART Newsletter*, ACM No. 63, June 1977.

R. S. Michalski, A System of Programs for Computer-Aided Induction: A Summary, *Proceedings of the Fifth International Joint Conference on Artificial Intelligence (ICJAI)*, pp. 319-320, MIT, Cambridge, MA, August 22-25, 1977.

R. S. Michalski, Conceptual Clustering: A Theoretical Foundation and a Method for Partitioning Data into Conjunctive Concepts, *Seminaires IRIA, Classification Automatique et Perception par Ordinateur*, pp. 253-295, IRIA (now INRIA), Versailles, France, 1979.

R. S. Michalski, Detection of Conceptual Patterns Through Inductive Inference, *Seminaires IRIA, Classification Automatique et Perception par Ordinateur*, pp. 297-339, Inria, France, 1979.

T. Dietterich and R. S. Michalski, Learning and Generalization of Characteristic Descriptions: Evaluation Criteria and Comparative Review of Selected Methods, *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pp. 223-231, Tokyo, Japan, August 20-23, 1979.

R. S. Michalski and R. Chilausky, Knowledge Acquisition by Encoding Expert Rules versus Computer Induction From Examples: A Case Study Involving Soybean Pathology, *International Journal for Man-Machine Studies*, No. 12, pp. 63-87, 1980.

R. S. Michalski and R. L. Chilausky, Learning by Being Told and Learning from Examples: An Experimental Comparison of the Two Methods of Knowledge Acquisition in the Context of Developing an Expert System for Soybean Disease Diagnosis, *International Journal of Policy Analysis and Information Systems*, Vol. 4, No. 2, pp. 125-161, 1980.

R. S. Michalski, Pattern Recognition as Rule-Guided Inductive Inference, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-2, No. 4, pp. 349-361, July 1980.

R. S. Michalski, Inductive Learning As Rule-Guided Generalization and Conceptual Simplification of Symbolic Descriptions: Unifying Principles and a Methodology, *Workshop on Current Developments in Machine Learning*, Carnegie-Mellon University, Pittsburgh, PA, July 16-18, 1980.

R. S. Michalski, Adaptation through Generalization, *Invited paper at the Fifth Annual Ann Arbor Adaptive Systems Workshop*, University of Michigan, Ann Arbor, MI, July 21-23, 1980.

R. S. Michalski, Inductive Inference As Rule-Guided Transformation of Symbolic Descriptions, *Invited paper for International Workshop on Program Construction*, Chateau de Bonas, France, September 8-12, 1980.

R. S. Michalski, Knowledge Acquisition Through Conceptual Clustering: A Theoretical Framework and an Algorithm for Partitioning Data into Conjunctive Concepts, *Journal of Policy Analysis and Information Systems*, Vol. 4, No. 3, pp. 219-244, September 1980.

R. S. Michalski and R. Stepp, A Method of Organizing Data into Conceptual Hierarchies, *Proceedings of the Classification Society Meeting*, Toronto, Canada, May 31 - June 2, 1981.

T. G. Dietterich and R. S. Michalski, Inductive Learning of Structural Descriptions: Evaluation Criteria and Comparative Review of Selected Methods, *Artificial Intelligence Journal*, Vol. 16, No. 3, pp. 257-294, July 1981.

A. Collins and R. S. Michalski, Toward a Formal Theory of Human Plausible Reasoning, Third Annual Conference of the Cognitive Science Society, Berkeley, CA, August 19-21, 1981.

R. S. Michalski and R. Stepp, An Application of AI Techniques to Structuring Objects into an Optimal Conceptual Hierarchy, *Proceedings of the Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, Vancouver, Canada, August 24-28, 1981.

R. S. Michalski, J. H. Davis, V. S. Bisht and J. B. Sinclair, PLANT/ds: An Expert Consulting System for the Diagnosis of Soybean Diseases, *Proceedings of the 1982 European Conference on Artificial Intelligence*, pp. 133-138, Orsay, France, July 1982 (subsequently published in *Plant Diseases*, 1982).

R. S. Michalski, A. B. Baskin and K. A. Spackman, A Logic-based Approach to Conceptual Database Analysis, *Proceedings of the Sixth Annual Symposium on Computer Applications in Medical Care (SCAMC-6)*, pp. 792-796, George Washington University Medical Center, Washington, DC, November 1-2, 1982.

R. S. Michalski, J. H. Davis, V. S. Bisht and J. B. Sinclair, A Computer-Based Advisory System for Diagnosing Soybean Diseases in Illinois, *Plant Diseases*, pp. 459-463, April 1983.

R. S. Michalski, A Theory and Methodology of Inductive Learning, *Artificial Intelligence*, Vol.20, No.2, pp. 111-161, 1983.

R. S. Michalski, Inductive Learning: A Review of Some Recent Work, Invited paper at the Third Yale Workshop on Applications of Adaptive Systems Theory, June 15-17, 1983.

R. S. Michalski and R. Stepp, How to Structure Structured Objects, *Proceedings of the International Machine Learning Workshop*, pp. 156-160, University of Illinois, Allerton House, Urbana, IL, June 22-24, 1983.

T. Dietterich and R. S. Michalski, Discovering Patterns in Sequences of Objects, *Proceedings of the International Machine Learning Workshop*, pp. 41-57, University of Illinois, Allerton House, Urbana, IL, June 22-24, 1983.

R. S. Michalski and R. Stepp, *Automated Construction of Classifications: Conceptual Clustering versus Numerical Taxonomy*, IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-5, No. 4, pp. 396-410, July 1983.

R. S. Michalski and A. B. Baskin, Integrating Multiple Knowledge Representation and Learning Capabilities in an Expert System: The ADVISE System, *Proceedings of the Eighth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 256-258, Karlsruhe, West Germany, August 8-12, 1983.

J. G. Carbonell, R. S. Michalski and T. M. Mitchell, Machine Learning: A Historical and Methodological Analysis, *The AI Magazine*, pp. 69-79, 1983.

R. S. Michalski, A. B. Baskin and K. Spackman, A Logic-Based Approach to Conceptual Data Base Analysis (Revised version), *Med. Inform.*, Vol. 8, No. 3, pp. 187-195, 1983.

T. Dietterich and R. S. Michalski, Discovering Patterns in Sequences of Events, *Artificial Intelligence Journal*, Vol. 25, No. 2, pp. 187-232, February 1985.

R. E. Reinke and R. S. Michalski, Incremental Learning of Decision Rules, Invited paper at the Machine Intelligence Workshop 11, Ross Priory, University of Strathclyde, March 26-29, 1985.

R. S. Michalski, Knowledge Repair Mechanisms: Evolution vs Revolution, *Proceedings of the Third International Machine Learning Workshop*, Skytop Lodge, Rutgers University, NJ, June 24-26, 1985.

T. Fermanian, R. S. Michalski and B. Katz, An Expert System to Assist Turfgrass Managers in Weed Identification, presented at the Proceedings of the 1985 Summer Computer Simulation Conference, Chicago, IL, July 1985.

R. S. Michalski, Knowledge Growth Mechanisms, Invited paper at 1985 Conference on Machine Learning, sponsored by the Knowledge-Based Systems Centre at the Polytechnic of The South Bank, Bloomsbury Crest Hotel, London, November 11-13, 1985.

R. Stepp and R. S. Michalski, Conceptual Clustering of Structured Objects: A Goal-Oriented Approach, *AI Journal*, 1986.

R. S. Michalski, Machine Learning Research in the Artificial Intelligence Laboratory at Illinois, *Machine Learning: A Guide to Current Research*, T. Mitchell, J. Carbonell, and R. S. Michalski (eds.), Kluwer Publishing Company, 1986.

J. Hong, I. Mozetic and R. S. Michalski, An Extension Matrix Approach to General Covering Problem, Proceedings of the SPIE Applications of Artificial Intelligence III, Sheraton Town Towers, Orlando, FL, April 1-3, 1986.

R. S. Michalski, I. Mozetic, J. Hong and N. Lavrac, The AQ15 Inductive Learning System: An Overview and Experiments, *Proceedings of the American Association for Artificial Intelligence Conference (AAAI)*, 1986.

R. S. Michalski and P. Winston, Variable Precision Logic, *Artificial Intelligence Journal*, No. 29, pp. 121-146, 1986.

B. Falkenhainer and R. S. Michalski, Integrating Quantitative and Qualitative Discovery: The ABACUS System, *Machine Learning Journal*, Vol. 4, 1987.

R. S. Michalski and R. Stepp, Clustering, *Artificial Intelligence Encyclopedia*, John Wiley and Sons Publishers, New York, NY 1987.

R. S. Michalski, *Concept Learning*, Artificial Intelligence Encyclopedia, John Wiley and Sons Publishers, New York, NY 1987.

P. Langley and R.S. Michalski, Machine Learning and Discovery, *Machine Learning Journal*, Vol.1, No.3, 1986

Nayel El-Shafei and R. S. Michalski, Quantitative Discovery and Reasoning about Failure Mechanisms in Pavement, to appear in *International Journal of Engineering Application of AI*, 1987.

R. S. Michalski, How to Learn Imprecise Concepts: A Method for Employing a Two-Tiered Knowledge Representation in Learning, Invited Paper, Proceedings of the 4th International Workshop on Machine Learning, University of California, Irvine, June 22-25, 1987.

D. Medin, W. D. Wattenmaker and R. S. Michalski, Constraints and Preferences in Inductive Learning: An Experimental Study Comparing Human and Machine Performance, accepted for publication in *Cognitive Science*, 1988.

A. Collins and R. S. Michalski, Logic of Plausible Reasoning, accepted for publication in *Cognitive Science*, 1988.

R. S. Michalski, Dynamic Recognition: An Outline of a Theory, in preparation.

Technical Reports:

R. S. Michalski, Graphical Minimization of Switching Functions in the Class of Disjunctive Normal Forms, *Reports of the Institute of Automatic Control*, No. 52, Polish Academy of Sciences, Warsaw, 1967 (in Polish).

R. S. Michalski, A Survey of Progress, Present Problems and New Ideas in the Field of Automatic Reading, *Technical Report*, Institute of Automatic Control, Polish Academy of Sciences, September 1968 (in Polish).

R. S. Michalski and B. H. McCormick, Interval Generalization of Switching Theory, report No. 442, Department of Computer Science, University of Illinois, Urbana, May 1971.

R. S. Michalski, A Geometrical Model for the Synthesis of Interval Covers, Report No. 461, Department of Computer Science, University of Illinois, Urbana, June 1971.

J. Larson and R. S. Michalski, AQVAL/1 (AQ7) User's Guide and Program Description, Report No. 731, Department of Computer Science, University of Illinois, Urbana, June 1975.

R. S. Michalski, On the Selection of Representative Samples from Large Relational Tables for Inductive Inference, Report No. M.D.C. 1.1.9, Department of Engineering, University of Illinois at Chicago Circle, Chicago, July 1975.

R. S. Michalski, Toward Computer-Aided Induction: A Brief Review of Currently Implemented AQVAL Programs, Report No. 874, Department of Computer Science, University of Illinois, Urbana, May 1977.

R. S. Michalski, A Planar Geometrical Model for Representing Multi-Dimensional Discrete Spaces and Multiple-Valued Logic Functions, Report No. 897, Department of Computer Science, University of Illinois, Urbana, January 1978.

R. S. Michalski, Designing Extended Entry Decision Tables and Optimal Decision Trees Using Decision Diagrams, Report No. 898, Department of Computer Science, University of Illinois, Urbana, March 1978.

R. S. Michalski and J. B. Larson, Selection of Most Representative Training Examples and Incremental Generation of VL_1 Hypotheses: The Underlying Methodology and the Description of Programs ESEL and AQ11, *Intelligent Systems Group Reports*, No. 867, Department of Computer Science, University of Illinois, Urbana, May 1978.

R. S. Michalski, Pattern Recognition as Knowledge-Guided Computer Induction, *Report No. 927*, Department of Computer Science, University of Illinois, Urbana, June 1978.

T. G. Dietterich and R. S. Michalski, Learning and Generalization of Structural Descriptions: Evaluation Criteria and Comparative Review of Selected Methods, *Report No. 1007*, Department of Computer Science, University of Illinois, Urbana, February 1980.

R. S. Michalski and R. Stepp, Concept-based Clustering versus Numerical Taxonomy, *Report No. 1073*, Department of Computer Science, University of Illinois, Urbana, October 1981.

B. Hoff, R. S. Michalski and R. Stepp, INDUCE 2: A Program for Learning Structural Descriptions from Examples, *Reports of the Intelligent Systems Group*, UIUCDCS-F-83-904, Department of Computer Science, University of Illinois, Urbana, January 1983.

R. S. Michalski and J. B. Larson, Incremental Generation of VL₁ Hypotheses: The Underlying Methodology and the Description of Program AQ11, *Reports of the Intelligent Systems Group*, Report 83-5, UIUCDCS-F-83-905, Department of Computer Science, University of Illinois, Urbana, January 1983.

R. S. Michalski, A Theory and Methodology of Inductive Learning, *Report UIUCDCS-R-83-1122*, Department of Computer Science, University of Illinois, Urbana, January 1983.

M. Stauffer and R. S. Michalski, Determining Computer Architectures and Compiler Structures through Inductive Inference: A Preliminary Investigation, *Intelligent Systems Group Reports*, No. 83-6, UIUCDCS-F-83-906, Department of Computer Science, University of Illinois, Urbana, January 1983.

R. S. Michalski, Learning Strategies and Automated Knowledge Acquisition: An Overview, *Reports of the Intelligent Systems Group*, ISG 84-6, UIUCDCS-F-84-926, Department of Computer Science, University of Illinois, Urbana, November 1984.

R. S. Michalski, Understanding the Nature of Learning: Issues and Research Directions, *Reports of the Intelligent Systems Group*, ISG 85-8, UIUCDCS-F-85-938, Department of Computer Science, University of Illinois, Urbana, February 1985 (an early version of the chapter in *Machined Learning II*, 1986).

T. G. Dietterich and R. S. Michalski, Learning to Predict Sequences, *Reports of the Intelligent Systems Group*, ISG 85-9, UIUCDCS-F-85-939, Department of Computer Science, University of Illinois, Urbana, February 1985.

R. Stepp and R. S. Michalski, Conceptual Clustering: Inventing Goal-Oriented Classifications of Structured Objects, *Reports of the Intelligent Systems Group*, ISG 85-10, UIUCDCS-F-85-940, Department of Computer Science, University of Illinois, Urbana, February 1985.

R. S. Michalski, H. Ko and K. Chen, SPARC/E(V.2), An Eleusis Rule Generator and Game Player, *Reports of the Intelligent Systems Group*, ISG 85-11, UIUCDCS-F-85-941, Department of Computer Science, University of Illinois, Urbana, February 1985.

R. S. Michalski, Knowledge Repair Mechanisms: Evolution vs Revolution, *Reports of the Intelligent Systems Group*, ISG 85-15, UIUCDCS-F-85-946, Department of Computer Science, University of Illinois, Urbana, July 1985.

R. S. Michalski and Patrick Winston, Variable Precision Logic, *Artificial Intelligence Memo*, #857, Massachusetts Institute of Technology, Cambridge, MA.

A. Collins and R. S. Michalski, The Logic of Plausible Reasoning: An Advance Report, *Reports of the Intelligent Systems Group*, ISG 85-17, UIUCDCS-F-85-951, Department of Computer Science, University of Illinois, Urbana, November 1985.

R. S. Michalski, Knowledge Repair Mechanisms: Evolution vs Revolution, Abstracts of the Third International Machine Learning Workshop, Skytop Lodge, Rutgers University, NJ, June 24-26, 1985.

D. L. Medin, Wm D. Wattenmaker, and R. S. Michalski, Constraints and Preferences in Inductive Learning: an Experimental Study Comparing Human and Machine Performance, *Reports of the Intelligent Systems Group*, ISG 86-1, UIUCDCS-F-86-952, Department of Computer Science, University of Illinois, Urbana, January 1986.

P. Haddawy, R. S. Michalski, A Variable Precision Logic with Dempster-Shafer Evaluation Scheme: Reasoning under Uncertainty and Time Constraints, *Reports of the Intelligent Systems Group*, 1987.

J. Hong, I. Mozetic and R. S. Michalski, AQ15: Learning Attribute-Based Descriptions from Examples Algorithm and User's Guide, *Reports of the Intelligent Systems*, ISG 86-17, UIUCDCS-F-85-949, Department of Computer Science, University of Illinois, Urbana, January 1986.

P. Haddawy, R. S. Michalski, A Variable Precision Logic with Dempster-Shafer Evaluation Scheme: Reasoning under Uncertainty and Time Constraints, *Reports of the Intelligent Systems Group*, October 1987.

R. S. Michalski and R. Stepp, INDUCE 3: A Program for Learning Structural Descriptions from Examples, Department of Computer Science, University of Illinois, Urbana, 1987.

R. S. Michalski and H. Ko, On the Nature of Explanation, Proceedings of the Symposium on the Explanation-based Learning, Stanford University, March 21-23, 1988

R. S. Michalski and L. Watanabe, Constructive Learning: Basic Ideas, Reports of the Machine Learning and Inference Laboratory, No. 88-1, Artificial Intelligence Center, George Mason University, 1988.

H. Ko and R.S. Michalski, Types of Explanation and Their Role in Constructive Learning, Reports of the Machine Learning and Inference Laboratory, No. 88-2, Artificial Intelligence Center, George Mason University, 1988.

F. Bergadano, S. Matwin, R.S. Michalski, Z. Zhang, A Measure of Quality of Descriptions, Reports of the Machine Learning and Inference Laboratory, No. 88-3, Artificial Intelligence Center, George Mason University, 1988.

F. Bergadano, S. Matwin, R.S. Michalski, Z. Zhang, Learning Flexible Concept Descriptions Using a Two-tiered Knowledge Representation: Ideas and a Method, Reports of the Machine Learning and Inference Laboratory, No. 88-4, Artificial Intelligence Center, George Mason University, 1988.

F. Bergadano, S. Matwin, R.S. Michalski, Z. Zhang, Learning Flexible Concept Descriptions Using a Two-tiered Knowledge Representation: Implementation and Experiments, Reports of the Machine Learning and Inference Laboratory, No. 88-35, Artificial Intelligence Center, George Mason University, 1988.

HARRY WECHSLER

CURRICULUM VITA

(October 18, 1988)

PERSONAL

Address: 4600 Briar Patch Ct.
Fairfax, Va. 22032

Telephone: (H) 703-425-1810
(O) 703-323-2650
(FAX) 703-323-2630

Born: January 23, 1948

EDUCATION:

Ph.D. Computer Science, Univ. of California, Irvine, Oct. 1975.
M.S. Computer Science, Weizmann Institute of Science, Israel,
June 1973.
B.S. Mathematics and Statistics, Tel-Aviv University, Israel,
July 1971.

WORK EXPERIENCE:

1988 - present	Professor of Computer Science, George Mason University (with tenure).
9/80 - 12/87	Associate Professor of Electrical Engineering, University of Minnesota (with tenure).
9/78 - 8/80	Assistant Professor of Computer Science and Electrical Engineering, University of Wisconsin.
8/76 - 8/78	Visiting Assistant Professor of Electrical Engineering and Systems Manager for the Advanced Automation Research Laboratory (AARL), Purdue University

PROFESSIONAL EXPERIENCE:

Teaching

I have taught undergraduate courses on Introduction to Programming, Linear Systems, Digital Logic Design, Probability, Statistics, and Computer Organization and Assembly Language, and graduate courses on Data Base Managements Systems, Computer Architecture, Computer Vision, Image Processing, Pattern Recognition, Artificial Intelligence, Design of Intelligent Systems (Expert Systems), Computer Networks and a Computer Science Workshop on PASCAL. I have also been invited by academic and business organizations abroad to teach short courses on computer vision & AI (AMK-Berlin, Selenia-Italy).

Consulting

I have been a consultant with Eaton, Honeywell, the National Institutes of Health (NIH), Perkin-Elmer, Proctor & Gamble, MIT Lincoln Lab, and 3M. The consulting work was connected with industrial automation, quality inspection, target identification, neural networks, videotracking, artificial intelligence, and biomedical image processing.

Scientific Visits

- (a) During the summer of 1980 (June-August) I held the position of Visiting Professor at INRIA (Institut National de Recherche en Informatique et en Automatique), France. I also attended and lectured at the NATO Advanced Study Institute on Digital Image Processing and Analysis, Bonas, France, June 23 - July 4, 1980.
- (b) Visiting Professor with the Department of Electrical Engineering at Kyoto University, Japan, July - August, 1982. Visited and lectured at the main research facilities of ETL, FUJITSU, HITACHI, NIPPON and TOSHIBA.
- (c) Attended and lectured at the NATO Advanced Study Institute on Robotics and Artificial Intelligence, Pisa, Italy, June 27 - July 8, 1983.
- (d) Visiting Professor with the Department of Electrical Engineering, Technion, Haifa, Israel, November-December 1984.
- (e) Invited lecturer at the NATO Advanced Study Institute on Pattern Recognition, Spa-Balmoral, Belgium, June 9-20, 1986.
- (f) Visiting Professor with the Center for Automation Research, University of Maryland, December 1986 - February 1987.
- (g) Invited lecturer at the NATO Advanced Research Workshop on Robots with Redundancy, Salo, Italy, June 27-July 1, 1988.

Invited Talks

I lectured in the United States (Carnegie-Mellon, JPL, SRI), Bell Labs), Austria, Canada, Denmark, France, Hungary, Israel, Italy, Japan, Mexico, The Netherlands, Sweden, Switzerland (ETH), UK (Cambridge, the Turing Institute), and West Germany.

RESEARCH INTERESTS:

1. Summary of Past and Present Research

Some of my first achievements were related to Medical Image Processing and Texture Analysis. My work on the ANALYSIS of CHEST RADIOGRAPHS was published as a book by Birkhauser Verlag and it also received the best paper award of the PATTERN RECOGNITION SOCIETY in 1978.

My previous work on medical image processing led to the establishment of joint research efforts with MAYO CLINIC. A Ph.D. thesis on RECONSTRUCTION ALGORITHMS for DYNAMIC CT IMAGERY was one result of such efforts. The thesis deals with the induced motion blur caused by body organs motion and shows how by using a Kalman filter, it can be significantly reduced. The experiments were done on the latest DSR (Dynamic Spatial Reconstructor) developed at Mayo and the results could greatly benefit the next CT generation of scanners.

Image interpretation and understanding is an emerging field, and there is great interest for its potential applications. Most of computer vision systems require a substantial amount of task-specific fine tuning. Such systems can be defined as "surface systems" which directly associate input states with interpretations (like the so-called Production Systems - PS). Progress in the field and robust performance can come only from developing "deep systems" which draw their deductions from a collection of fundamental principles. Toward that goal we developed the WIGNER (DISTRIBUTION) TRANSFORM (WD) as a means for providing basic low-level distributed image representations for tasks like invariant image recognition, optical flow derivation, and clustering. Our work was the first to introduce the WD as a computer vision tool and several papers have been published.

My main areas of research are those of Computer Vision and Image Processing, and AI (Artificial Intelligence). Both areas are vital for the ongoing revolution in automation. The field of Computer Vision and Image Processing is concerned with issues related to both image enhancement and image understanding. My emphasis has been on image understanding. Tasks such as developing vision as a sensor feedback for intelligent/programmable robotics, quality inspection techniques, target identification and tracking for real-time varying imagery are examples of my work. This work was done both in academia and as a consultant for private industry (3M and Honeywell).

Another interest is in how the memory should be organized and accessed and on the relationships between low-level and high-level processes. The 'recognition' process, both for knowledge-base engineering and computer vision is a 'loop' where the raw data and low-level representations should help in 'indexing' the memory, while the models should 'drive' the process of acquiring such low-level representations by providing the appropriate parameters. Characteristics to this research were the cooperative efforts undertaken with HONEYWELL on real-time target acquisition and interpretation, and my recent work on Distributed Associative Memories (DAM).

2. Present and Future Research Goals

My research goals are within the Computer Vision, Neural Networks, and AI areas. I am interested in Fault-Tolerant 3-D Object Recognition and massively parallel and distributed computation (PDP) through the use of Neural Networks. The redundant aspect provided by PDP computation is crucial to both robotics and data retrieval applications and our recent research showed the feasibility of our approach. One particular model of Neural Networks under investigation is that of Distributed Associative Memories (DAMs). In addition to the intrinsic scientific merit of such research in advancing the state-of-the-art one would also benefit by finding improved solutions for the bin-picking and the data fusion problems, respectively. The bin-picking problem is concerned with picking-up partially-occluded objects from a bin of scrambled mechanical parts. The problem is unsolved

and its solution will have a tremendous impact on industrial automation and automatic assembly. (Similar issues have also to be dealt with within the context of target recognition and tracking.) The data fusion problem attempts to enhance the decision-making process by integrating different signal modalities. Within the AI area my interest is to investigate machine learning aimed at developing intelligent computer aided instruction (ICAI) systems, conceptual development, concept definition, and self organization (adaptiveness).

The neural network research is characteristic to attempts of using parallel and distributed models of computation. The approach is synergetic and draws from optimization, psychology, neurophysiology and cognitive sciences. Our effort is directed towards developing fault-tolerant models endowed with advanced learning capabilities. Specifically, modeling should be robust to noise in both the system input and/or the memory storage and be endowed with regeneration (error-correcting) capabilities. Learning is related to the ability to categorize by providing taxonomy skills like generalization and abstraction. It is only through context dependent learning/discovery that computers, like the human beings, can become truly intelligent.

An ongoing goal is that of writing a book on Computational Vision. Academic Press decided to assume editorial responsibility for publishing my forthcoming book. My book will describe a computational vision theory, functionally appropriate for tasks like recognition and navigation, and whose goal is to capture image invariants. Appropriate image representations and the corresponding computational algorithms operating on them are the main topics explored throughout the book. Emphasis is put on novel approaches based on neural networks (distributed and parallel computation), motion, expert systems for computer vision and data fusion.

RESEARCH GRANT AWARDS:

1. NATO (International Scientific Exchange Programmes) to organize and co-direct (with A. Sood) an Advanced Study Institute on Active Perception and Robot Vision in Maratea, Italy, July 16-29, 1989 (\$62,450).
2. National Science Foundation (NSF-EET8713563) Emerging Engineering Technology - for research on 'Fault-Tolerant 3-D Object Recognition'. Project supported 1987-1988 (\$29,983).
3. Department of the Army (DASD-H-CRS/87-08-067) for research on 'Target Recognition in Cluttered Environments.' The original award was for 3 years (\$237,170) and the project was in effect for 1987/8.
4. Defense Advanced Research Project Agency (DARPA-F19628-85-C-0002) for research on 'Invariant Object Recognition Using Cojoint Image Representations.' Project supported 1986-1988 (\$231,016) and monitored by MIT Lincoln Lab (CX-7401/CS-7916/CX-9333).
5. National Science Foundation (NSF-ECS8310057) Automation, Bioengineering and Sensing System Programs - for research on 'Invariant Image Recognition'. Project supported 1984-1986 (\$118,624).

6. Microelectronics and Information Science (MEIS) Center - for research on 'Intelligent Systems'. Project supported 1984-1986 (\$800,000) (the other three Co-PIs from CS, Psychology, and Management Sciences).
7. National Science Foundation Grant (NSF-ECS8105168) Automation, Bioengineering, and Sensing Programs for research on 'Texture Analysis'. Project supported 1981-1983 (\$57,622).

AWARDS:

I was a co-winner of the Fourth Annual Pattern Recognition Society Award (1979). This award is made each year for the best paper published in the Pattern Recognition Journal. Manuscripts published in both 1977 and 1978 have been judged for this award.

PROFESSIONAL SOCIETIES:

I am an Associate Editor and Member of the Editorial Board for the Pattern Recognition Journal. I am a Senior Member of IEEE and a Member of AAAI. I am a member of the Technical Committee on Pattern Analysis and Machine Intelligence of the IEEE Computer Society. I am also a regular reviewer for the IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI), Chair Paper Review Process for Computer Science Conference (CSC '90).

ADVISING:

Four Ph.D. students and several Master students completed their research under my supervision. Presently I supervise 3 Ph.D. students (N. Vassilas, L. Zimmerman and B. Gipe) in Computer Vision. The Ph.D. students who graduated are:

James Botic (now @ EATON Corp.)
Artificial Intelligence (5/83)
Thesis: Triggering and Control of Inferences Processes

Raj Acharya (now @ SUNY-Buffalo)
Image Processing (12/83)
Thesis: Time Domain Signal Processing Approach to Obtain Improved Reconstruction of Dynamic Events Using High Speed CT

Lowell Jacobson (now @ University of Massachusetts)
Computer Vision (1/87)
Thesis: Spatial/Spatial-Frequency Representations for Object Recognition

Todd Reed(now @ University of Minnesota)
Computer Vision (4/88)
Thesis: Segmentation of Textured Images and Gestalt Organization Using the Wigner Distribution

PUBLICATIONS

(a) Books and Book Chapters

- 1) H. Wechsler (1977). AUTOMATIC DETECTION OF RIB COHTOURS IN CHEST RADIOGRAPHS. AN APPLICATION OF PATTERN RECOGNITION TECYNIQUES IN MEDICAL DIAGNOSIS, Birkhauser, Basel, Switzerland.
- 2) H. Wechsler (1987), INVARIANCE IN PATTERN RECOGNITION in Advances in Electronics and Electron Physics, P. W. Hawkes (Ed.), Academic Press, 262-322.
- 3) H. Wechsler (1987), NETWORK REPRESENTATIONS AS MATCH FILTERS for INVARIANT OBJECT RECOGNITION, in Pattern Recognition Theory and Applications, P. Devijver and J. Kitler (Eds.), Springer-Verlang, 269-276.
- 4) H. Wechsler and L. Zimmerman (1989), REDUNDANT VISION FOR ROBOTICS, in Robots with Redundancy, A.K. Bejczy (Ed.), Springer Verlag.
- 5) A. Sood and H. Wechsler (1989), FAULT-TOLERANT MEDICAL IMAGE INTERPRETATION, in The Formation, Handling, and Evaluation of Medical Images, A. Todd-Pokropek (Ed.), Springer-Verlag.
- 6) H. Wechsler (1989), COMPUTATIONAL VISION, Academic Press (Forthcoming).
- 7) H. Wechsler (Ed.) (1989), NEURAL NETWORKS for VISUAL and MACHINE PERCEPTION, Oxford University Press. (Forthcoming)

(b) Journal Papers

- 1) H. Wechsler (1976). A Fuzzy Approach to Medical Diagnosis, Int. Journal of Biomedical Computing, Vol. 7, No. 3, 191-203.
- 2) H. Wechsler and J. Sklansky (1977), Finding the Rib Cage in Chest Radiographs, Pattern Recognition, Vol. 7, No. 3, 191-203.
- 3) H. Wechsler and M. Kidode (1977), A New Edge Detection Technique and its Implementation, IEEE Trans. on Systems, Man and Cybernetics, Vol. SMC-7, No. 12, 827-836.
- 4) H. Wechsler (1978), A New Low-Level Procedure for Image Segmentation, Computer Graphics and Image Processing, Vol. 7, No. 1, 120-129.
- 5) H. Wechsler and K. S. Fu (1978), Image Processing Algorithms Applied to Rib Boundary Detection in Chest Radiographs, Computer Graphics and Image Processing, Vol. 7, No. 3, 375-390.

- 6) H. Wechsler (1979), A Structural Approach for Shape Analysis Using Mirroring Axes, Computer Graphics and Image Processing, Vol. 9, No. 3, 246-266.
- 7) H. Wechsler and M. Kidode (1979), A Random Walk Procedure for Texture Discrimination, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. PAMI-1, No. 3, 272-280.
- 8) H. Wechsler (1980), Texture Analysis-A Survey, Signal Processing, Vol. 2, No. 3, 272-282.
- 9) H. Wechsler and T. Citron (1980), Feature Extraction for Texture Classification, Pattern Recognition, Vol. 12, No. 5, 301-311.
- 10) H. Wechsler (1981). A New and Fast Algorithm for Estimating the Perimeter of Objects for Industrial Vision Tasks, Computer Graphics and Image Processing, Vol. 17, No. 4, 375-386.
- 11) E. Gurari and H. Wechsler (1982), On the Difficulties Involved in the Segmentation of Pictures, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. PAMI-4, No. 3, 304-306.
- 12) L. Jacobson and H. Wechsler (1982), A Paradigm for Invariant Object Recognition of Brightness, Optical Flow and Binocular Disparity Images, Pattern Recognition Letters, Vol. 1, No. 1, 61-68.
- 13) L. Jacobson and H. Wechsler (1984), A New Theory for Invariant Object Recognition in the Frontoparallel Plane, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. PAMI-6, No. 3, 325-331.
- 14) L. Jacobson and H. Wechsler (1984), Invariant Analogical Image Representations and Pattern Recognition, Pattern Recognition Letters, Vol. 2, No. 5, 289-299.
- 15) L. Jacobson and H. Wechsler (1987), Derivation of Optical Flow Using Spatiotemporal Frequency Analysis, Computer Vision, Graphics, and Image Processing, 38, 29-65.
- 16) R. Acharya, P. B. Heffernan, R. Robb and H. Wechsler (1987), High Speed 3-D Imaging of the Beating Heart Using Temporal Estimation, Computer Vision, Graphics and Image Processing, 39, 279-290.
- 17) L. Jacobson and H. Wechsler (1988), Joint Spatial/Spatial-Frequency Representations, Signal Processing, Vol. 14, No. 1, 37-68.
- 18) T. Reed and H. Wechsler (1988), Tracking of Non-Stationarities for Texture Fields, Signal Processing, Vol. 14, No. 1, 95-102.
- 19) J. M. Char, V. Cherkassky, H. Wechsler, and L. Zimmerman (1988), Distributed and Fault-Tolerant Computation for Retrieval Tasks Using Distributed Associative Memories (DAMs), IEEE Trans. on Computers Vol. 37, No. 4, 484-490.

- 20) H. Wechsler and L. Zimmerman (1988), 2-D Invariant Object Recognition Using Distributed Associative Memories, IEEE Trans. on Pattern Analysis and Machine Intelligence (to appear).
- 21) H. Wechsler and L. Zimmerman (1989), Distributed Associative Memories for Bin-Picking, IEEE Trans. on Pattern Analysis and Machine Intelligence, (to appear).
- 22) J. M. Char, V. Cherkassky, and H. Wechsler (1989), Fault-Tolerant Database Using Distributed Associative Memories, Information Sciences (to appear).
- 23) P. Meer, S. Wang, and H. Wechsler (1989), Edge Detection by Associative Mapping, Pattern Recognition (to appear).
- 24) T. Reed and H. Wechsler (1988), Segmentation of Textured Images and Gestalt Organization Using Spatial/Spatial-Frequency Representations, IEEE Trans. on Pattern Analysis and Machine Intelligence (under review).
- 25) B. Kangar-Parsi, B. Kangar-Parsi, and H. Wechsler (1988), Simultaneous Fitting of Several Planes to Point Sets Using Neural Networks, Computer Vision, Graphics, and Image Processing (under review)
- 26) T. Reed, H. Wechsler, and M. Werman (1989), Texture Segmentation Using a Diffusion Region Growing Technique, Signal Processing (under review).

(c) (Reviewed) Proceedings Papers

- 1) H. Wechsler (1975), Applications of Fuzzy Logic to Medical Diagnosis, 5th International Symposium in Multiple-Valued Logic, Indiana University, Bloomington, Indiana.
- 2) H. Wechsler and J. Sklansky (1975), Automatic Detection of Rib Contours in Chest Radiographs, 4th International Joint Conference on Artificial Intelligence, Tbilissi, USSR.
- 3) H. Wechsler (1976), Interpretation of Scenes Consisting of Planar Geometrical Shapes, 3rd International Joint Conference on Pattern Recognition, San Diego, California.
- 4) H. Wechsler and K. S. Fu (1977), Automatic Analysis of Chest Radiographs, MEDINFO, Toronto, Canada.

- 5) H. Wechsler and K. S. Fu (1977), Texture Analysis in Chest Radiographs, IEEE Symposium on Systems, Man and Cybernetics, Washington, DC.
- 6) H. Wechsler (Invited Panelist) (1977), Software for Image Processing Recognition, 1st International Computer Software and Applications Conference (COMSAC 1977), Chicago, Illinois.
- 7) H. Wechsler (1978), Color Edge Detection, IEEE Workshop on Pattern Recognition and Artificial Intelligence, Princeton, New Jersey.
- 8) H. Wechsler, J. Mui, and K. S. Fu (1978), Biomedical Image Processing and Understanding, Society of Photo-Optical Instrumentation Engineers (SPIE) Symposium, San Diego, California.
- 9) H. Wechsler and M. Kidode (1978), Mathematical Models for Scene Analysis, 4th International Joint Conference on Pattern Recognition, Kyoto, Japan.
- 10) H. Wechsler (1979), Feature Extraction for Texture Discrimination, Pattern Recognition and Image Processing Conference, Chicago, Illinois.
- 11) H. Wechsler (1980), Taxonomy and Segmentation of Textured Images, 5th International Joint Conference on Pattern Recognition, Miami, Florida.
- 12) H. Wechsler (1981), How to Improve on the Measurement Process Despite Digitization Effects, 2nd Scandinavian Conference on Image Analysis, Helsinki, Finland.
- 13) L. Jacobson and H. Wechsler (1982), the Wigner Distribution as a Tool for Deriving an Invariant Representation of 2-D Images, Pattern Recognition and Image Processing Conference, Las Vegas, Nevada.
- 14) L. Jacobson and H. Wechsler (1982), The Wigner Distribution and its Usefulness for 2-D Image Processing, 6th International Joint Conference on Pattern Recognition, Munich, West Germany.
- 15) L. Jacobson and H. Wechsler (1983), The Composite Pseudo Wigner Distribution: a Computable and Versatile Approximation to the Wigner Distribution, IEEE Int. Conference on Acoustics, Speech and Signal Processing, Boston, MA.
- 16) L. Jacobson and H. Wechsler (1983), Spatiotemporal frequency analysis of time-varying imagery, Computer Vision and Pattern Recognition Conference, Washington, DC.
- 17) R. S. Acharya, and H. Wechsler (1983), Time Domain Signal Processing Approaches to Dynamic Medical Imaging, Computer Vision and Pattern Recognition Conference, Washington, DC.

- 18) L. Jacobson and H. Wechsler (1984), A Path Toward Solving the Bin-Picking Problem, Int. Conference on Robotics, Atlanta, Georgia.
- 19) R. Acharya, R. Robb, and H. Wechsler (1984), Time Varying Image Sequence Analysis and the DSR, a High Temporal Resolution 3-D Synchronous Volume Scanner, ICAASP '84, San Diego, CA.
- 20) L. Jacobson, and H. Wechsler (1984), Multipurpose Low-Level Image Processing, 7th Int. Conference on Pattern Recognition, Montreal, Canada.
- 21) H. Wechsler (1984), Invariant Local Spectral Signatures for Image Processing, 7th Int. Conference on Digital Signal Processing, Florence, Italy.
- 22) L. Jacobson, U. Thomanschefskey, and H. Wechsler (1984), Resolution and Invariance Criteria for Object Pattern Recognition, SPIE/Cambridge Symposium on Intelligent Robots and Computer Vision, Boston, MA.
- 23) L. Jacobson and H. Wechsler (1985), Joint Spatial/Spatial-Frequency Representations for Image Processing, SPIE/Cambridge Int. Conference on Intelligent Robots and Computer Vision, Boston, MA.
- 24) L. Jacobson and H. Wechsler (1985), FOVEA-A System for Invariant Visual Form Recognition, 2nd Int. Technical Symposium on Optical and Electro-Optical Applied Science and Engineering, Cannes, France.
- 25) R. S. Acharya, C. Benchimol, and H. Wechsler (1986), Analysis of 3-D Biomedical Images, 8th Int. Conference on Pattern Recognition, Paris, France.
- 26) H. Wechsler and L. Zimmerman (1987), Fault-Tolerant Recognition Using DAMs, 1st Int. Conference on Neural Networks, San Diego, CA.
- 27) T. Reed and H. Wechsler (1987), Texture Segmentation via Non-Stationarity Detection Using Joint Spatial/Spatial-Frequency Representations, 10th Int. Conference on Digital Signal Processing, Florence, Italy.
- 28) H. Wechsler and L. Zimmerman (1987), Invariant Object Recognition Using Distributed Associative Memories, IEEE Conference on Neural Information Processing Systems, Denver, Colorado.
- 29) P. Meer, S. Wang, and H. Wechsler (1987), Edge Detection in Feature Space, IEEE Workshop on Computer Vision, Miami, Florida.
- 30) H. Wechsler and L. Zimmerman (1988), Rotation and Scale Invariant Object Recognition Using a Distributed Associative Memory, Symposium on Innovative Science and Technology, Los Angeles, California.

- 31) H. Wechsler (1988), Parallel and Distributed Computation for Fault-Tolerant Object Recognition, 1988 Goddard Conference on Space Applications of Artificial Intelligence, NASA Goddard Space Flight Center, Greenbelt, MD.
- 32) T. Reed and H. Wechsler (1988), Texture Segmentation and Organization Using the Wigner Distribution, EUSIPCO, Grenoble, France.
- 33) H. Wechsler and L. Zimmerman (1988), Distributed Associative Memories and Data Fusion, 2nd Int. Conference on Neural Networks, Boston, MA.
- 34) J. M. Char, V. Cherkassky, and H. Wechsler (1988), Distributed Processing for Database Design, 2nd Int. Conference on Neural Networks, Boston, MA.
- 35) H. Wechsler and D. Rine (1988), Object Oriented Programming and its Relevance to Designing Intelligent Software Systems, Int. Conference on Computer Languages (ICCL '88), Miami Beach, Florida.
- 36) T. Reed and H. Wechsler (1988), Texture Analysis and Clustering Using the Wigner Distribution, 9th Int. Conference on Pattern Recognition, Rome, Italy.
- 37) V. Cherkassky and H. Wechsler (1989), Fault Tolerant User Interfaces for Data Bases, 11th Int. Conference on Software Engineering, Pittsburgh, PA, May 15-18 (under review).
- 38) H. Wechsler and A. Sood (1989), Active Sensing and Multistage Processing for Dexterous Robot Manipulation, NASA Conference on Space Telerobotics, Pasadena, CA (under review)
- 39) V. Cherkassky and H. Wechsler (1989), Fault-Tolerant Information Retrieval, Int. Conference on R&D in Information Retrieval, Cambridge, MA.

Heedong Ko

Department of Computer Science
George Mason University
4400 University Drive
Fairfax, VA 22030-4444

PRESENT ADDRESS

5938 Cove Landing Rd #201
Burke, VA 22015
Home Ph: (703)250-2467
Office Ph: (703)323-3884

PERMANENT ADDRESS

5938 Cove Landing Rd #201
Burke, VA 22015

PERSONAL DATA

Citizen of USA, June 1988.

VITA

Heedong Ko was born on September 8, 1959 in Seoul, Korea. He grew up there until he moved to USA in 1977. He attended State University of New York at Albany in 1979 and received his B.S. in Computer Science and Applied Mathematics in 1982. After spending two months at the Korean Advanced Institute of Science and Technology under summer scholarship program by the Korean Scientists and Engineers Association in America inc., he arrived at the University of Illinois at Urbana-Champaign in 1982 to study for his Ph.D. degree in Computer Science. While at the University of Illinois, he was a teaching assistant in the Department of Computer Science in the fall of 1982, a university fellow for 1982-83 academic year and a research assistant under Prof. R. S. Michalski from the spring of 1983 to the spring of 1988. Now, he is a visiting assistant professor in the department of computer science at George Mason University.

EDUCATIONAL RECORD

State University of New York at Albany
1/79-5/82: BS degree in Computer Science and Applied Mathematics, May 82.
University of Illinois, College of Engineering, Urbana, IL
9/82-11/88: PhD degree in Computer Science, expected in 1988.

HONORS

Tau Beta Pi: Illinois Alpha 1984.
University Fellow of the Department of Computer Science at University of Illinois at Urbana-Champaign 1982-1983.
Selected Participant of Summer Scholarship Program by Korean Scientists and Engineers Association in America, 1982.
Summa Cum Laude: State University of New York at Albany 1982.
Scholarship for Physics/Mathematics Majors at the State University of New York at Albany, 1979.

MEMBERSHIP

AAAI Special Interest Group in Manufacturing
American Association of Artificial Intelligence
Tau Beta Pi
Korean Scientists and Engineers Association in America

PROFESSIONAL EXPERIENCE

- Visiting assistant professor at the department of computer science at George Mason University 9/1988-present.
- Research assistant at the department of computer science at University of Illinois 1983-8/1988.
- Teaching assistant system at the department of computer science at University of Illinois 9/82-1/83.
- Invited Participant of the Fourth International Workshop on Machine Learning at Irvine 1987.
- Invited Participant of Intelligent Manufacturing Workshop at Philadelphia (AAAI) 1986.
- Invited Participant of Third International Workshop on Machine Learning at Skytop 1985.
- Reviewer of American Association of Artificial Intelligence Conference at Philadelphia 1986.
- Codeveloper of Machine Learning Portion of National Exhibition, "ROBOTS and BEYOND: The Age of Intelligent Machines.", 1987.
- System programmer at Korean Advanced Institute of Science and Engineering, 6/82-7/82.

PUBLICATIONS

JOURNALS

- H. Ko and K. Lee, **Automatic Assembling Procedure Generation from Mating Conditions**, Computer Aided Design, Vol. 19, Num 1, pp. 3-10, Jan/Feb 1987.

BOOKS

- H. Ko, **Symbolic Process Prediction: Another Learning Problem**: in **Machine Learning: A Guide to Current Research**, Mitchell, T. M., Carbonell, J. G. and Michalski, R. S. (eds), Kluwer Academic Publishers, pp. 141-144, 1986.
- R. S. Michalski, H. Ko and K. Chen, **Qualitative Prediction: The SPARC/G Methodology for Describing and Predicting Discrete Processes**, in **Current Issues in Expert Systems**, A. Van Lamsweerde and P. Dufour (eds), Academic Press Incorporated, pp 125-158, 1987.
- H. Ko and K. Lee, **Automatic Assembling Procedure Generation from Mating Conditions**, published in a text for the correspondence course PT616 Computer-aided Engineering by OPEN university, England, 1987.

SYMPOSIUM

- R. Michalski and H. Ko, **On the Nature of Explanation: Why Wine Bottle Shattered?**, AAAI 1988 Spring Symposium Series, Mini-Symposium: Explanation-Based Learning, March 1988.

ABSTRACTS

- H. Ko and K. Lee, **Automatic Assembling Procedure Generation from Mating Conditions**, Cambridge Scientific Abstracts, vol 35 (5), 1987.

TECHNICAL REPORTS

- R. S. Michalski, H. Ko and K. Chen, **SPARC/E(V.2), An Eleusis Rule Generator and Game Player**, Intelligent System Group, ISG 85-11, UIUCDCS-F-85-941, Department of Computer Science, University of Illinois, Urbana, February 1985.
- H. Ko and K. Lee, **Toward A Practical Planning System for Assembly Tasks**, Intelligent System Group, ISG 86-7, UIUCDCS-F-86-957, Department of Computer Science, University of Illinois, Urbana, September 1986.

H. Ko, **Reasoning About Kinematics from Mating Conditions**, Intelligent System Group, ISG-16, UIUC-ENG-88-1708 Department of Computer Science, University of Illinois, Urbana, January 1988.

R. Michalski and H. Ko, **On the Nature of Explanation: Why Wine Bottle Shattered?**, Intelligent System Group, ISG-15, Department of Computer Science, University of Illinois, Urbana, 1988.

Behrooz Kamgar-Parsi

Computer Science Department
George Mason University
Fairfax Virginia, 22030
Phone: 703-323-2713

Citizenship: USA

RESEARCH INTERESTS

Computer vision, image processing, artificial intelligence
neural networks and learning.

EDUCATION

Ph.D. in physics from the Catholic University of America, Washington D.C.,
1978. (Thesis topic: Triangular Ising Lattice.)

Graduate study in computer science at the University of Maryland, College Park,
Maryland, January–December 1982.

PROFESSIONAL EXPERIENCE

Assistant Professor at the Computer Science Department, George Mason
University, teaching and carrying out research in computer vision, image processing
and neural networks, August 1988 to present.

Research Scientist at the Center for Automation Research, University of Maryland;
carrying out research in computer vision and related areas. I worked on motion,
stereo, 3D recovery, matching of intensity and range images, and applications of
neural networks to a number of problems in image analysis and A.I. June
1985 to August 1988.

Visiting Scientist at the Division of Computer Research and Technology, National Institutes
of Health; developing numerical algorithms, January 1983 to June 1985.

Research in Disturbances of Human Circadian Rhythms in collaboration with physicians
at the National Institutes of Health, 1978 to December 1981. (This research led to the first
successful treatment of human non-24-hour sleep-wake syndrome.)

PENDING GRANT

Army Research Office, 1988–1990.

PUBLICATIONS

1. "Critical Temperature of Triangular Ising Lattice with Small Second-Neighbor Interaction", Behrooz Kamgar-Parsi and P. Meijer, *Physics Lett.* 66A, 69 (1978).
2. "Successful Treatment of Human Non-24-Hour Sleep-wake Syndrome", Behrooz Kamgar-Parsi, T. Wehr and C. Gillin, *SLEEP*. Vol. 6, No. 3 (1983).
3. "An Improved Branch and Bound Algorithm for Computing k-Nearest Neighbors", Behrooz Kamgar-Parsi and L. Kanal, *Pattern Recognition Lett.*, Vol.3, pp. 7-12 (1985).
4. "Closed-Form Solutions to Image Flow Equations", A. Waxman, Behrooz Kamgar-Parsi and M. Subbarao, *International Journal of Computer Vision*, Vol. 1, pp. 239-258, (1987); also the Proceedings of the First International Conference on Computer Vision, p. 12-24 (1987).
5. "Minimization of the Quantization Error in Camera Calibration", Behrooz Kamgar-Parsi, Proceedings of the DARPA Image Understanding Workshop, p. 671-680 (1987).
6. "Correspondence from Correspondence", J. Aloimonos and Behrooz Kamgar-Parsi, *Topical Meeting on Machine Vision*, pp. 46-51 (1987).
7. "An Efficient Model of Neural Networks for Optimization", Behzad Kamgar-Parsi and Behrooz Kamgar-Parsi, Proceedings of the IEEE First International Conference on Neural Networks, p. 785-790 (1987).
8. "A New Interpolation Technique for Line Search in Optimization of Multivariate Functions", Behrooz Kamgar-Parsi. Kamgar-Parsi, submitted to *Mathematics of Computation*.
9. "Calibration of a Stereo System with Small Relative Angles, Behrooz Kamgar-Parsi and R. Eastman", Proceedings of the IEEE Computer Vision and Pattern Recognition Conference, pp. 44-51 (1988); also *Computer Vision, Graphics, and Image Processing (CVGIP)*, in press.
10. "A Nonparametric Method for Fitting a Line to Noisy Images", Behzad Kamgar-Parsi and Behrooz Kamgar-Parsi, CAR-TR-315, Center for Automation Research, University of Maryland (1987); submitted to *IEEE Trans, Pattern Analysis and Machine Intelligence*.
11. "Evaluation of Quantization Error in Computer Vision", Behrooz Kamgar-Parsi and Behzad Kamgar-Parsi, Proceedings of the IEEE Computer Vision and Pattern Recognition Conference, pp. 52-62 (1988), and the Proceedings of the DARPA Image Understanding Workshop, pp. 720-730 (1988); also *IEEE Trans, Pattern Analysis and Machine Intelligence* (in press).

12. "Simultaneous Fitting of Several Planes to Point Sets Using Neural Networks", Behrooz Kamgar-Parsi, Behzad Kamgar-Parsi and Harry Wechsler, submitted to Computer Vision, Graphics, and Image Processing (CVGIP).
13. "On the Error Analysis in Stereo Recovery of 3D Point Position", Behrooz Kamgar-Parsi, Behzad Kamgar-Parsi and Azriel Rosenfeld, submitted to IEEE Trans, Pattern Analysis and Machine Intelligence, (1988).
14. "Quantization Error In Hexagonal Sensory Configuration", Behzad Kamgar-Parsi, Behrooz Kamgar-Parsi and William A. Sander, III, submitted to IEEE Trans, Pattern Analysis and Machine Intelligence, (1988).
15. "How to Cluster in Parallel with Neural Networks", Behzad Kamgar-Parsi, J.A. Gualtieri, J.E. Devaney and Behrooz Kamgar-Parsi, to appear at the Proceedings of the Second Symposium on Massively Parallel Computation (George Mason University, October, 1988).
16. "Mapping of the Ocean Floor", Behrooz Kamgar-Parsi and Jeffrey L. Jones, CAR-TR-377, Center for Automation Research, University of Maryland; also the final report to the Naval Research Laboratory, Washington, D.C., 1988.
17. "The Importance of Being Synchronous in Neural Networks", Behzad Kamgar-Parsi and Behrooz Kamgar-Parsi, abstract in the Proceedings of the International Neural Networks Society Conference, p. 104 (Boston, September, 1988).
18. "Exact Results for the Sum of Uniform Random Variables", Behrooz Kamgar-Parsi, Behzad Kamgar-Parsi and Menashe Brosh, submitted to SIAM Journal of Statistics and Scientific Computing, (1988).
19. "Registration of Multiple Overlapping Range Images: Scenes without Distinctive Features", Behrooz Kamgar-Parsi, Jeffrey L. Jones and Azriel Rosenfeld, submitted to IEEE Trans, Pattern Analysis and Machine Intelligence, (1988).
20. "Applicability Analysis of Structure from Motion Using Image Flow", Menashe Brosh, Behrooz Kamgar-Parsi and Behzad Kamgar-Parsi, (in preparation).

Igor Mozetic
Department of Computer Science
George Mason University
4400 University Drive
Fairfax, VA 22030-4444

RESUME

- Personal data:* Born on February 3, 1957 in Ljubljana, Yugoslavia; married, two children.
- Education:* 1979 B.Sc. in Computer Science (Implementation and Use of Advice Language AL1.5)
1984 M.Sc. in Computer Science (Qualitative Model of the Heart)
1988 Ph.D. in Computer Science (Machine Learning of Qualitative Models)
All from University of Ljubljana, Yugoslavia
- Employment:* Feb. 1980 - July 1988
Positions: Research Assistant, Research Associate
Department of Computer Science and Informatics
Jozef Stefan Institute, Ljubljana, Yugoslavia
Sept. 1984 - Sept. 1986
Positions: Fulbright Scholar, Visiting Research Associate
Department of Computer Science
University of Illinois, Urbana-Champaign, IL
Sept. 1988 - now
Position: Visiting Assistant Professor
Department of Computer Science
George Mason University, Fairfax, VA
- Awards:* 1979 Slovenian Research Council award for excellent students
1982 Award for a medical expert system on the RAST-Rijeka inventions exhibit
1984 Fulbright scholarship for 9 months research in the USA
1986 Boris Kidric award (with collaborators) for automatic derivation of knowledge from a deep model
- Research interests:* artificial intelligence, knowledge based systems, knowledge representation and acquisition, logic programming, machine learning, qualitative modeling

Research experience:

I was involved in the development of the KARDIO expert system for diagnosis and treatment of cardiac arrhythmias. As a main contribution to the project I developed a qualitative model of the heart which simulates its electrical activity. The model was used to automatically generate the ECG knowledge base for all physiologically possible multiple arrhythmias (over 2400 rules).

During my stay at the University of Illinois I was working mainly in machine learning. In one project I applied the learning from examples techniques to compress a large medical knowledge base to 3% of its original size. Another project involved taking medical records of patients with known diagnoses, using learning from examples to construct diagnostic rules, and measuring their diagnostic accuracy. Results show that the rules reached the same level of diagnostic accuracy as medical experts.

My current research, based on my Ph.D. thesis, involves development of the system for automatic construction and refinement of qualitative models. The approach enables the representation and interpretation of a qualitative model on different levels of abstraction. The system involves learning from examples, location of incorrect and incomplete rules in the model, their incremental refinement, and diagnosis of failures in the model. The system was successfully tested on a non-trivial problem of reconstructing a model of the heart's electrical activity.

Igor Mozetic
Department of Computer Science
George Mason University
4400 University Drive
Fairfax, VA 22030-4444

SELECTED PUBLICATIONS

Lavrac, N., Bratko, I., Mozetic, I., Cercek, B., Grad, A., Horvat, M. KARDIO-E - an expert system for electrocardiographic diagnosis of cardiac arrhythmias. *Expert Systems 2 (1)*, pp. 46-49, 1985.

Michalski, R.S., Mozetic, I., Hong, J., Lavrac, N. The multi-purpose incremental learning system AQ15 and its testing application to three medical domains. *Proc. AAAI-88*, pp. 1041-1045, Philadelphia, August 11-15, 1986, Morgan Kaufmann.

Mozetic, I. Knowledge extraction through learning from examples. In *Machine Learning: A Guide to Current Research* (T.M.Mitchell, J.G.Carbonell, R.S.Michalski, Eds.), pp. 227-231, Kluwer, Boston, 1986.

Mozetic, I. Learning of qualitative models. In *Progress in Machine Learning* (I.Bratko, N.Lavrac, Eds.), pp. 201-217, Sigma Press, Wilmslow, UK, 1987.

Mozetic, I. The role of abstractions in learning qualitative models. *Proc. 4th Intl. Workshop on Machine Learning*, pp. 242-255, Irvine, CA, July 22-25, 1987, Morgan Kaufmann.

Mozetic, I., Lavrac, N. Incremental learning from examples in a logic-based formalism. *Proc. Intl. Workshop on Machine Learning, Meta-Reasoning and Logics*, pp. 109-127, Sesimbra, Portugal, February 14-17, 1988.

Bratko, I., Mozetic, I., Lavrac, N. Automatic synthesis and compression of cardiological knowledge. In *Machine Intelligence 11* (J.E.Hayes, D.Michie, J.Richards, Eds.), pp. 435-454, Clarendon Press, Oxford, 1988.

Mozetic, I., Bratko, I., Urbancic, T. Varying level of abstraction in qualitative modelling. In *Machine Intelligence 12* (J.E.Hayes, D.Michie, E.Tyugu, Eds.), Oxford University Press (in press).

Bratko, I., Mozetic, I., Lavrac, N. *KARDIO: A Study in Deep and Qualitative Knowledge for Expert Systems*. The MIT Press, Cambridge, MA (in press).

INVITED LECTURES

Acquisition and refinement of deep knowledge bases. Technische Universitat Munchen, West Germany, December 4, 1986.

Automatic construction of qualitative models. The Turing Institute, Glasgow, Scotland, December 8, 1986.

The KARDIO project. Leiden University, The Netherlands, March 3, 1987.

Qualitative modelling: an overview. Graduate course on Knowledge Representation, University of Stockholm, Sweden, March 7, 1988.

Learning and interpretation of qualitative models. Graduate course on Design of Knowledge Based Systems, University of Stockholm, Sweden, March 8, 1988.

Artificial intelligence techniques and expert systems. Graduate seminar, Institute for applied mechanical engineering, University of Novi Sad, Yugoslavia, April 14, 15 and 16, 1988.

PART . III

1. Proposed Budget

ORGANIZATION		PROPOSAL NO.		DURATION (MONTHS)	
George Mason University				Proposed	Grant
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR		AWARD NO.			
Ryzsard Michalski					
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)		NSF FUNDED PERSON-MOS		FUNDS REQUESTED BY PROPOSER	FUNDS GRANTED BY (IF DIFFERENT)
		CAL.	ACADSUMR		
1. Ryzsard Michalski, P.I.			6/9	\$ 65,725	\$
2. H. Wechsler, P.I.			6/9	47,866	
3. H. Ko, P.I.			6/9	28,577	
4. P. Pachowich, Investigator			6/9	28,577	
5. (2) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)			36wks	39,570	
6. (6) TOTAL SENIOR PERSONNEL (1-5)			144wks	210,315	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. () POST DOCTORAL ASSOCIATES					
2. (2) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.)			27mo	99,625	
3. (3,5) GRADUATE STUDENTS				158,620	
4. () UNDERGRADUATE STUDENTS					
5. () SECRETARIAL-CLERICAL					
6. () OTHER					
TOTAL SALARIES AND WAGES (A+B)				468,560	
7. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				33,044	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				501,604	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING \$1,000:) Please see attached equipment list.					
TOTAL PERMANENT EQUIPMENT				77,499	
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				19,500	
2. FOREIGN					
F. PARTICIPANT SUPPORT COSTS					
1. STIPENDS \$					
2. TRAVEL					
3. SUBSISTENCE					
4. OTHER					
TOTAL PARTICIPANT COSTS					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				6,300	
2. PUBLICATION COSTS/PAGE CHARGES				4,500	
3. CONSULTANT SERVICES				3,000	
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Equipment Maintenance, Phone, Photocopy				22,800	
TOTAL OTHER DIRECT COSTS				36,600	
H. TOTAL DIRECT COSTS (A THROUGH G)				635,203	
I. INDIRECT COSTS (SPECIFY)					
TOTAL INDIRECT COSTS				262,394	
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)				897,597	
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS SEE GPM 252 AND 253)					
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				\$897,597	\$
PI/PD TYPED NAME & SIGNATURE*		DATE		FOR NSF USE ONLY	
NST. REP. TYPED NAME & SIGNATURE*		DATE		INDIRECT COST RATE VERIFICATION	
				Date Checked Date of Rate Sheet Initials - DGC	

COMPLETING) 7/1/89 - 6/30/90		SUMMARY		APPENDIX 1	
ORGANIZATION		FOR NSF USE ONLY			
George Mason University		PROPOSAL NO.		DURATION (MONTHS)	
		AWARD NO.		Proposed	Granted
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR					
Ryzsard Michalski					
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)		NSF FUNDED PERSON-MOS.		FUNDS REQUESTED BY PROPOSER	FUNDS GRANTED BY NSI (IF DIFFERENT)
		CAL.	ACAD	SUMR	
Ryzsard Michalski, P.I.			2/9	\$ 20,444	\$
H. Wechsler, P.I.			2/9	14,889	
3H. Ko, P.I.			2/9	8,889	
P. Pachowich, Investigator			2/9	8,889	
(2) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)			12wks	12,308	
6. (6) TOTAL SENIOR PERSONNEL (1-5)			48wks	65,419	
OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
() POST DOCTORAL ASSOCIATES					
2. (2) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.)			9mo	32,240	
3. (3, 5) GRADUATE STUDENTS				48,510	
() UNDERGRADUATE STUDENTS					
5. () SECRETARIAL-CLERICAL					
6. () OTHER					
TOTAL SALARIES AND WAGES (A+B)				146,169	
FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				10,361	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				156,530	
PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING \$1,000:)					
Please see attached equipment list.					
TOTAL PERMANENT EQUIPMENT				77,499	
L. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				6,500	
2. FOREIGN					
PARTICIPANT SUPPORT COSTS					
1. STIPENDS \$					
2. TRAVEL					
3. SUBSISTENCE					
4. OTHER					
TOTAL PARTICIPANT COSTS					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				2,100	
2. PUBLICATION COSTS/PAGE CHARGES				1,500	
3. CONSULTANT SERVICES				1,000	
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Equipment Maintenance, Phone, Photocopy				7,600	
TOTAL OTHER DIRECT COSTS				12,200	
TOTAL DIRECT COSTS (A THROUGH G)				252,729	
I. INDIRECT COSTS (SPECIFY) 56% of salaries and wages					
TOTAL INDIRECT COSTS				81,855	
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)				334,584	
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS SEE GPM 252 AND 253)					
AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				\$334,584	\$
I/PD TYPED NAME & SIGNATURE*		DATE	FOR NSF USE ONLY		
NST. REP. TYPED NAME & SIGNATURE*		DATE	INDIRECT COST RATE VERIFICATION		
			Date Checked	Date of Rate Sheet	Initials - DGC
					Program

FOR NSF USE ONLY			
PROPOSAL NO.		DURATION (MONTHS)	
		Proposed	Granted
ORGANIZATION George Mason University			
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Ryzsard Michalski		AWARD NO.	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)		NSF FUNDED PERSON-MOS.	FUNDS REQUESTED BY PROPOSER
		CAL. ACADSUMR	FUNDS GRANTED BY NSF (IF DIFFERENT)
Ryzsard Michalski, P.I.		2/9	\$ 21,875
4.H. Wechsler, P.I.		2/9	15,931
3.H. Ko, P.I.		2/9	9,511
P. Pachowich, Investigator		2/9	9,511
(2) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)		12wks	13,170
6. (6) TOTAL SENIOR PERSONNEL (1-5)		48wks	69,998
OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)			
() POST DOCTORAL ASSOCIATES			
2. (2) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.)		9mo	33,186
(3.5 GRADUATE STUDENTS			55,055
() UNDERGRADUATE STUDENTS			
5. () SECRETARIAL-CLERICAL			
6. () OTHER			
TOTAL SALARIES AND WAGES (A+8)			158,239
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)			11,031
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)			169,270
PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING \$1,000:)			
TOTAL PERMANENT EQUIPMENT			
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)			6,500
2. FOREIGN			
PARTICIPANT SUPPORT COSTS			
1. STIPENDS \$			
2. TRAVEL			
3. SUBSISTENCE			
4. OTHER			
TOTAL PARTICIPANT COSTS			
OTHER DIRECT COSTS			
1. MATERIALS AND SUPPLIES			2,100
2. PUBLICATION COSTS/PAGE CHARGES			1,500
3. CONSULTANT SERVICES			1,000
4. COMPUTER (ADPE) SERVICES			
5. SUBCONTRACTS			
6. OTHER Equipment Maintenance, Phone, Photocopy			7,600
TOTAL OTHER DIRECT COSTS			12,200
TOTAL DIRECT COSTS (A THROUGH G)			187,970
I. INDIRECT COSTS (SPECIFY) 56% of salaries and wages			
TOTAL INDIRECT COSTS			88,614
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)			276,584
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS SEE GPM 252 AND 253)			
AMOUNT OF THIS REQUEST (J) OR (J MINUS K)			\$276,584
PI/PD TYPED NAME & SIGNATURE*		DATE	FOR NSF USE ONLY
IST. REP. TYPED NAME & SIGNATURE*		DATE	INDIRECT COST RATE VERIFICATION
		Date Checked	Date of Rate Sheet
			Initials - DGC

COMPLETING) 7/1/91 - 6/30/92		PROPOSAL BUDGET		APPENDIX 1					
ORGANIZATION George Mason University				FOR NSF USE ONLY					
				PROPOSAL NO.	DURATION (MONTHS) Proposed Granted				
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Ryzsard Michalski				AWARD NO.					
1. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)				NSF FUNDED PERSON-MOS		FUNDS REQUESTED BY PROPOSER		FUNDS GRANTED BY NSF (IF DIFFERENT)	
				CAL.	ACAD	SUMR			
Ryzsard Michalski, P.I.						2/9	\$ 23,406	\$	
H. Wechsler, P.I.						2/9	17,046		
3.H. Ko, P.I.						2/9	10,177		
P. Pachowich, Investigator						2/9	10,177		
(2) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)						12wks	14,092		
6. (6) TOTAL SENIOR PERSONNEL (1-5)						48wks	74,898		
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)									
() POST DOCTORAL ASSOCIATES									
2. (2) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.)					9mo		34,199		
3. (3,5) GRADUATE STUDENTS							55,055		
() UNDERGRADUATE STUDENTS									
() SECRETARIAL-CLERICAL									
6. () OTHER									
TOTAL SALARIES AND WAGES (A+B)							164,152		
FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)							11,652		
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)							175,804		
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING \$1,000:)									
TOTAL PERMANENT EQUIPMENT									
TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)							6,500		
2. FOREIGN									
E. PARTICIPANT SUPPORT COSTS									
1. STIPENDS \$									
2. TRAVEL									
3. SUBSISTENCE									
4. OTHER									
TOTAL PARTICIPANT COSTS									
G. OTHER DIRECT COSTS									
1. MATERIALS AND SUPPLIES							2,100		
2. PUBLICATION COSTS/PAGE CHARGES							1,500		
3. CONSULTANT SERVICES							1,000		
4. COMPUTER (ADPE) SERVICES									
5. SUBCONTRACTS									
6. OTHER Equipment Maintenance, Phone, Photocopy							7,600		
TOTAL OTHER DIRECT COSTS							12,200		
TOTAL DIRECT COSTS (A THROUGH G)							194,504		
I. INDIRECT COSTS (SPECIFY) 56% of Salaries and Wages									
TOTAL INDIRECT COSTS							91,925		
TOTAL DIRECT AND INDIRECT COSTS (H + I)							286,429		
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS SEE GPM 252 AND 253)									
AMOUNT OF THIS REQUEST (J) OR (J MINUS K)							\$286,429	\$	
/PD TYPED NAME & SIGNATURE*				DATE	FOR NSF USE ONLY				
INDIRECT COST RATE VERIFICATION									
INIT. REP. TYPED NAME & SIGNATURE*				DATE	Date Checked	Date of Rate Sheet	Initials - DGC		
				Program					

BUDGET EXPLANATION

1. The two other senior personnel are Drs. Kamgar-Parsi and Mozetic.
2. Fifteen percent effort during the academic year will be funded by George Mason University so that Professors Michalski and Wechsler can work on this research project.

3. Facilities and Equipment

Current equipment of the Center for Artificial Intelligence includes:

- 11 SUN-3 machines, 1 SUN 3/260 file server (12MB memory), 9 SUNs 3/60 (8MB memory each) and 1 SUN 3/50 (4MB memory). The file server has two disk drives with 280 MB each, and one Eagle SCSI disk drive with 140 MB. The SUN 3/260 also has a 1/4" cartridge tape drive.
- Two printers available on the net : PRISM dot-matrix printer, and Apple Laser Writer
- Two Symbolics 3640, running Genera 7.2
- Two MAC-II computers, with LaserView displays and 1MB of memory each.
- Two microVAX II/GPX workstations, connected to the same network as SUN-3s. Each of them has a speech generator (DECTalk). They are used for demonstrations of the Emerald system: an integrated system of learning programs.
- Hero 2000 Robot system connected to a SUN3 workstation.

Requested Equipment:.

The proposed research, which requires a development of software for combining computer vision with machine learning, will involve very extensive computations. Therefore, we need to enhance our facilities with more powerful workstations than we currently have. We are requesting two SUN-4/110ME-8-P5 workstations, and would like to equip them with additional specialized hardware and software for image processing.

The specialized hardware hardware is needed for the acquisition, processing, storage and display of images. The MaxVideo family of image processing boards seems to be appropriate for our work.

The DIGIMAX and FRAMESTORE modules will provide real-time black and white aquisition, display and storage of RS-170 video.

DIGIMAX will also provide digitalization of RGB color video at 1/3 real-time (i.e.: 3 frame times).

The MAX-SP module will allow real-time averaging and mask substraction of the video signal.

The VFIR module will perform either 3*3 or 10*1 convolutions, for edge-detection, at the video rate.

The MAX-box is a 20 slot 32 bit VME card cage which will power, cool and house the MAX-VIDEO cards with room for future expansion. The VME-VME bus converter will provide the necessary connection path into the SUN computer. The 6U to 9U adapter is needed to secure the 6U bus converter in the 9U SUN computer.

The requested equipment thus includes:

- | | |
|--|-------------|
| • 2 SUN-4/110xx workstations with
153A (Floating point unit) and 514A
(141 MByte disk with 60 MByte
1/4-inch tape drive) each | \$55,800.00 |
| • CCD Color Video Camera | \$ 1,500.00 |
| • DIGIMAX 10MHz | \$ 2,463.00 |
| • FRAMESTORE 10MHz | \$ 3,678.00 |
| • VFIR 10MHz | \$ 3,675.00 |
| • MAX-SP 10MHz | \$ 1,608.00 |
| • VME to VME Bus Converter | \$ 3,000.00 |
| • MAX-BOX P2 | \$ 4,500.00 |
| • 9 U to 6U Adapter | \$ 600.00 |
| • 10 SUN 1/4" Cartridges | \$ 675.00 |

Total equipment cost	\$77,499.00
----------------------	-------------

2, Current and Pending Support

Current and Pending Support for Research and Education in Science and Engineering

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of the proposal.

I. Name of Principal Investigator	Source of Support	Project Title	Award Amount (or Annual Rate)	Period Covered by Award	Person-Months or % of Effort Committed to the Project			Location of Research
					ACAD.	SUMM.	CAL. YR.	
Ryszard Michalski								
A. Current Support List—if none, report none	See attached							
B. Proposals Pending 1. List this proposal	NSF/DARPA	Machine Learning and Dynamic Recognition for Robot Vision	\$897,597	7/1/89–6/30/92	15%	2 mo		George Mason University
2. Other pending proposals, including renewal applications. If none, report none.								
3. Proposals planned to be submitted in near future. If none, report none.	NSF	INRIA Research Institute France US Cooperation	\$30,000					
II. Name of co-principal investigator and/or faculty associate A. _____ B. _____								
III. Transfer of Support If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.								
IV. Other agencies to which this proposal has been/will be submitted								

USE ADDITIONAL SHEETS AS NECESSARY

Attachment:

Current support (R.S. Michalski)

Source: Office of Naval Research

No. N00014-88-K-0226

Title: Multistrategy Learning

PIs: R.S. Michalski, K. DeJong

Annual rate: \$150 000/year

Period covered: Jan 1, 88 - Dec. 31, 90

% of effort: Academic and Summer - 10%

Location of Research: George Mason University

Source: Office of Naval Research

No. N00014-88-K-0397

Title: Human Plausible Reasoning and Learning:

Development of a Computational Model and Unified Theory

PIs: R.S. Michalski, D. Medin

Annual rate: \$120 000/year

Period covered: May15, 88 - May 14, 90 .

% of effort: Academic and Summer - 10%

Location of Research: George Mason University

Source: DARPA (a subcontract from the University of Illinois at UI)

No. NO001487-K-0874

Title: Integrating Machine Learning, Inference and Planning

PIs: R.S. Michalski, K. DeJong, J. Zytkow

Annual rate: \$300 000/year

Period covered: January 1, 88 - August 31, 89

% of effort: Academic and Summer - 13%

Location of Research: George Mason University

Current and Pending Support for Research and Education in Science and Engineering

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of the proposal.

I. Name of Principal Investigator	Source of Support	Project Title	Award Amount (or Annual Rate)	Period Covered by Award	Person-Months or % of Effort Committed to the Project			Location of Research
					ACAD.	SUMM.	CAL. YR.	
Harry Wechsler								
A. Current Support List—if none, report none								
B. Proposals Pending 1. List this proposal	NSF/DARPA	Machine Learning and Dynamic Recognition for Robot Vision	\$897,597	7/1/89-6/30/92	15%	2 mo.		George Mason University
2. Other pending proposals, including renewal applications. If none, report none.	NIH	Artificial Neural Systems (ANS) models of Visual Learning	\$275,243	6/1/89-5/31/92	10%	1 mo.		George Mason University
3. Proposals planned to be submitted in near future. If none, report none.								
II. Name of co-principal investigator and/or faculty associate								
A. _____								
B. _____								
III. Transfer of Support If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.								
IV. Other agencies to which this proposal has been/will be submitted								

USE ADDITIONAL SHEETS AS NECESSARY

Current and Pending Support for Research and Education in Science and Engineering

The following information should be provided for each investigator and other senior personnel. Failure to provide this information may delay consideration of the proposal.

I. Name of Principal Investigator Heedong Ko	Source of Support	Project Title	Award Amount (or Annual Rate)	Period Covered by Award	Person-Months or % of Effort Committed to the Project			Location of Research
					ACAD.	SUMM.	CAL. YR.	
A. Current Support List—if none, report none								
B. Proposals Pending 1. List this proposal	NSF/DARPA	Machine Learning and Dynamic Recognition for Robot Vision	\$897,597	7/1/89-6/30/92		6 mo		George Mason Univers
2. Other pending proposals, including renewal applications. If none, report none.								
3. Proposals planned to be submitted in near future. If none, report none.								
II. Name of co-principal investigator and/or faculty associate A. _____ B. _____								
III. Transfer of Support If this project has previously been funded by another agency, please list and furnish information for immediately preceding funding period.								
IV. Other agencies to which this proposal has been/will be submitted								

USE ADDITIONAL SHEETS AS NECESSARY

Appendix 1

An Algorithm for Dynamic Recognition

The algorithm consists of three major stages, the REDUCE - that reduces a set of candidates based on the initial information about an unknown object and the context, the INDUCE - that determines necessary and sufficient criteria for recognizing the object in a given context, and INQUIRE - that determines what facts/measurements should be obtained to recognize the object.

1. The REDUCE stage

Given:

1. Characteristic descriptions of concepts (i.e., collections of assertions about each concept).
2. Zero or more initial facts about the object to be recognized.
3. A context that defines the original class of concept descriptions for consideration.

Determine:

A set of candidate concept descriptions (a subset of the original set of concept descriptions) that are consistent with the facts and the context.

The algorithm determines which assertions can be eliminated due to the context, and which are or not satisfied by the initial facts. Assertions that have dissatisfied conditions are removed. Satisfied conditions are removed from assertions. If a sufficient number of facts is provided, only one concept description may be consistent with them, and this will trigger the recognition. Otherwise, the process moves to the next stage.

2. The INDUCE stage

Given:

1. Candidate concept descriptions.
2. Background knowledge (containing a specification of variables and their types, background concepts, constructive generalization rules, and a preference criterion).

Determine:

Discriminant concept recognition rules that maximize the given criterion of preference.

An inductive learning program is applied to descriptions obtained at the REDUCE stage to determine the most preferred discriminant concept recognition rules (Hong, Mozetic and Michalski, 1986). Such a program must be able to use conjunctive concept descriptions as learning examples. Programs such as AQ15 or INDUCE 4 can be used for this purpose after some modification.

3. The INQUIRE stage

Given:

A set of discriminant concept recognition rules.

Determine:

One or more descriptors (variables, relations, tests, measurements, etc.) in the rules, such that if their values are known, then only one candidate rule will remain. Measure or determine the value of these descriptors, and apply them to the rules to trigger the recognition.

The process of determining which descriptors to measure can be done sequentially or in parallel. We have so far developed only an algorithm for a sequential selection. In this algorithm, each step determines which descriptor, if its value is known, will allow the system to maximally reduce the current discriminant concept recognition rules. In a more advanced approach, this may be a descriptor that not only maximally reduces the rules, but also is easy or inexpensive to measure. The latter approach is called the *utility-based* inference control schema, and is described in (Baskin and Michalski, 1988).

Appendix 2

List of Publications of the AI Center for the Last 5 Years

NOVEMBER 8, 1988

PUBLICATIONS
of the
MACHINE LEARNING and INFERENCE LABORATORY

Center for Artificial Intelligence
George Mason University
4400 University Drive
Fairfax, Virginia 22030

(Last Five Years)

- 1983 1 R. S. Michalski and R. Stepp, "LEARNING FROM OBSERVATION: Conceptual Clustering," Chapter in the book, *MACHINE LEARNING: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, 1983, pp. 331-363.
- 2 J. Carbonell, R. S. Michalski and T. Mitchell, "An Overview of Machine Learning," Chapter in the book *MACHINE LEARNING: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, 1983, pp. 3-23.
- 3 R. S. Michalski, "A Theory and Methodology of Inductive Learning," Chapter in the book, *MACHINE LEARNING: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, 1983, pp. 83-134.
- 4 T. Dietterich and R. S. Michalski, "A Comparative Review of Selected Methods for Learning from Examples," Chapter in the book, *MACHINE LEARNING: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, 1983, pp. 41-81.
- 5 R. S. Michalski, J. Carbonell and T. Mitchell (Editors), *MACHINE LEARNING: An Artificial Intelligence Approach*, TIOGA Publishing Company, TIOGA Publishing Co., Palo Alto, 1983.
- 6 R. S. Michalski, J. H. Davis, V. S. Bisht and J. B. Sinclair, "A Computer-Based Advisory System for Diagnosing Soybean Diseases" in *Illinois Plant Disease*, pp. 459-463, April 1983.
- 7 S. Cramm, "ESEL/2: A Program for Selecting the Most Representative Training Events for Inductive Learning," ISG 83-1, UIUCDCS-F-83-901, Department of Computer Science, University of Illinois, Urbana, January 1983.

- 8 J. Pollack, "Relevant Variable Selection for Inductive Learning Programs," ISG 83-2, UIUCDCS-F-83-902, Department of Computer Science, University of Illinois, Urbana, January 1983.
- 9 M. Seyler, "A Knowledge Guided Image Interpretation Program", ISG 83-3, UIUCDCS-F-83-903, Department of Computer Science, University of Illinois, Urbana, January 1983.
- 10 W. Hoff, R. S. Michalski, R. Stepp, "INDUCE 2: A Program for Learning Structural Descriptions from Examples," ISG 83-4, UIUCDCS-F-83-904, Department of Computer Science, University of Illinois, Urbana, January 1983.
- 11 R. S. Michalski and J. B. Larson, "INCREMENTAL GENERATION OF VL1 HYPOTHESES: the underlying methodology and the description of program AQ11," ISG 83-5, UIUCDCS-F-83-905, Department of Computer Science, University of Illinois, Urbana, January 1983.
- 12 R. S. Michalski, "A Theory and Methodology of Inductive Learning" (Modified version of 1983-3), Artificial Intelligence, pp. 111-161, 1983 and Report No. UIUCDCS-R-83-1122, Department of Computer Science, University of Illinois, Urbana, January 1983.
- 13 M. Stauffer and R. S. Michalski, "Determining Computer Architectures and Compiler Structures Through Inductive Inference: A Preliminary Investigation," ISG 83-6, UIUCDCS-F-83-906, Department of Computer Science, University of Illinois, Urbana, January 1983.
- 14 P. Vaidya, "A Representation Scheme for Constructive Induction," ISG 83-7, UIUCDCS-F-83-907, Department of Computer Science, University of Illinois, Urbana, January 1983.
- 15 J. L. Orejel-Opisso, "On the Representation of Rules," ISG 83-8, UIUCDCS-F-83-908, Department of Computer Science, University of Illinois, Urbana, January 1983.
- 16 R. S. Michalski and R. Stepp, "Conjunctive Conceptual Clustering: Classification Using Background Knowledge," Submitted to the North American Branch of the Classification Society, Philadelphia, May 29-31, 1983.
- 17 P. Vaidya, "On Induction Experiments, Their Validity and Constructive Induction," ISG 83-9, UIUCDCS-F-83-909, Department of Computer Science, University of Illinois, Urbana, May 1983.
- 18 R. S. Michalski, "Inductive Learning: A Review of Some Recent Work," Invited paper for the Third Yale Workshop on Applications of Adaptive Systems Theory, June 15-17, 1983.
- 19 R. S. Michalski (Editor), Proceedings of the International Machine Learning Workshop, University of Illinois Allerton House, Urbana, June 22-24, 1983.

- 20 R. S. Michalski and R. Stepp, "How to Structure Structured Objects," Proceedings of the International Machine Learning Workshop, University of Illinois Allerton House, Urbana, June 22-24, 1983, pp. 156-160.
 - 21 T. Dietterich and R. S. Michalski, "Discovering Patterns in Sequences of Events," Proceedings of the International Machine Learning Workshop, University of Illinois Allerton House, Urbana, June 22-24, 1983, pp. 41-57.
 - 22 R. S. Michalski and R. Stepp, "Automated Construction of Classifications: Conceptual Clustering versus Numerical Taxonomy," IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. PAMI-5, No. 4, July 1983, pp. 396-410.
 - 23 B. Biswas, "A Programmer's Guide for CLUSTER, A Program for Conjunctive Clustering," ISG 83-10, UIUCDCS-F-83-910, Department of Computer Science, University of Illinois, Urbana, July 1983.
 - 24 A. Boulanger, "The Expert System Plant/CD: A Case Study in Applying the General Purpose Inference System Advise to Predicting Black Cutworm Damage in Corn," M.S. Thesis, UIUCDCS-R-83-1134, Department of Computer Science, University of Illinois, Urbana, July 1983.
 - 25 R. S. Michalski and A. B. Baskin, "Integrating Multiple Knowledge Representations and Learning Capabilities in an Expert System: The ADVISE System," Proceedings of the 8th IJCAI, Karlsruhe, West Germany, August 8-12, 1983, pp. 256-258.
 - 26 J. G. Carbonell, R. S. Michalski and T. M. Mitchell, "Machine Learning: A Historical and Methodological Analysis," The AI Magazine, 1983, pp. 69-79.
 - 27 R. S. Michalski, A. B. Baskin and K. Spackman, "A Logic-Based Approach to Conceptual Data Base Analysis," (Revised version), Med. Inform., Vol. 8, No. 3, 1983, pp. 187-195.
 - 28 J. Becker, "AQINTERLISP: An INTERLISP Program for Inductive Generalization of VL1 Event Sets, USER'S & PROGRAMMER'S GUIDE," ISG 83-11, UIUCDCS-F-83-911, Department of Computer Science, University of Illinois, Urbana, September 1983.
 - 29 R. Reinke, "PLANT/ds: An Expert System for the Diagnosis of Soybean Diseases Common in Illinois, User's Guide and Program Description," ISG 83-12, UIUCDCS-F-83-912, Department of Computer Science, University of Illinois, Urbana, October 1983.
 - 30 R. Stepp, "A Description and User's Guide for CLUSTER/2, A Program for Conjunctive Conceptual Clustering," Report No. UIUCDCS-R-83-1084, Department of Computer Science, University of Illinois, Urbana, November 1983.
 - 31 K. A. Spackman, "QUIN: Integration of Inferential Operators within a Relational Database," ISG 83-13, UIUCDCS-F-83-917, M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, 1983.
-

- 1984 1 K.-E. Pua, "Translating Annotated Predicate Calculus to PROLOG", ISG 84-1, UIUCDCS-F-918, Department of Computer Science, University of Illinois, Urbana, June, 1984.
- 2 T. D. Channic, "ADVISECORE A Screen Package for Expert Systems," ISG 84-2, UIUCDCS-F-84-919, Department of Computer Science, University of Illinois, Urbana, July, 1984.
- 3 L. E. Rodewald, "BABY: An Expert System for Patient Monitoring in a Newborn Intensive Care Unit," M.S. Thesis, ISG 84-3, UIUCDCS-F-84-920, Department of Computer Science, University of Illinois, Urbana, July, 1984.
- 4 R. E. Reinke, "Knowledge Acquisition and Refinement Tools for the ADVISE META-EXPERT System," M.S. Thesis, ISG 84-4, UIUCDCS-F-84-921, Department of Computer Science, University of Illinois, Urbana, July, 1984.
- 5 P. W. Baim, "Automated Acquisition of Decision Rules: The Problems of Attribute Construction and Selection," M.S. Thesis, ISG 84-5, UIUCDCS-F-84-922, Department of Computer Science, University of Illinois, Urbana, July, 1984.
- 6 R. S. Michalski, "Inductive Learning as Rule-Guided Generalization of Symbolic Descriptions: A Theory and Implementation," Chapter 24, in *Automatic Program Construction Techniques*, Macmillan Publishing Company, New York, 1984, pp. 517-552.
- 7 R. E. Stepp, "Conjunctive Conceptual Clustering: A Methodology and Experimentation," PhD Thesis, UIUCDCS-R-84-1189, Department of Computer Science, University of Illinois, Urbana, September, 1984.
- 8 R. S. Michalski, "Learning Strategies and Automated Knowledge Acquisition: An Overview," ISG 84-6, UIUCDCS-F-84-926, Department of Computer Science, University of Illinois, Urbana, November 1984.
- 9 B. Falkenhainer, "ABACUS Adding Domain Constraints to Quantitative Scientific Discovery," ISG 84-7, UIUCDCS-F-84-927, Department of Computer Science, University of Illinois, Urbana, November 1984.
- 10 K.-E. Pua, "An Annotated Predicate Calculus Inference System (APCIS)," ISG 84-8, UIUCDCS-F-84-928, Department of Computer Science, University of Illinois, Urbana, November 1984.
- 11 A. R. Nowicki, "A Comparison of Performances of a SUN-2 Workstation, VAX 11/780, and the Xerox Dolphin on a Learning Algorithm," ISG 84-9, UIUCDCS-F-84-929, Department of Computer Science, University of Illinois, Urbana, December 1984.
- 12 L. Jacobson and H. Wechsler, A New Theory for Invariant Object Recognition in the Frontoparallel Plane, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. PAMI-6, No. 3, pp. 325-331.

- 13 L. Jacobson and H. Wechsler, Invariant Analogical Image Representations and Pattern Recognition, Pattern Recognition Letters, Vol. 2, No. 5, pp. 289-299.
-
- 1985 1 J. Becker, "AQ-PROLOG: A Prolog Implementation of an Attribute-Based Inductive Learning System," ISG 85-1, UIUCDCS-F-85-930, Department of Computer Science, University of Illinois, Urbana, January 1985.
- 2 L. Rendell, "Substantial Constructive Induction Using Layered Information Compression: Tractable Feature Formation in Search," UIUCDCS-R-85-1198, Department of Computer Science, University of Illinois, Urbana, Revised May 1985, and presented at Proceedings of the Ninth International Joint Conference on Artificial Intelligence, "University of California at Los Angeles, 1985.
- 3 P. Hoffman, "MEL - A Learning Program that Improves by Experience in Playing the Game of MILL," ISG 85-2, UIUCDCS-F-85-931, Department of Computer Science, University of Illinois, Urbana, January 1985.
- 4 P. Newkirk, "TURF: An Expert System for Ground Development," ISG 85-3, UIUCDCS-F-85-933, Department of Computer Science, University of Illinois, Urbana, January 1985.
- 5 T. Channic, "Editing Network-Structured Knowledge Bases in the ADVISE System," ISG 85-4, UIUCDCS-F-85-934, Department of Computer Science, University of Illinois, Urbana, February 1985.
- 6 J. Becker, "Topics in Incremental Learning of Discriminant Descriptions," ISG 85-5, UIUCDCS-F-85-935, Department of Computer Science, University of Illinois, Urbana, February 1985.
- 7 J. Becker, "Word Sense Disambiguation Using Constraints and Supports," ISG 85-6, UIUCDCS-F-85-936, Department of Computer Science, University of Illinois, Urbana, February 1985.
- 8 J. Becker, "Macros and Utilities for Franz Lisp," ISG 85-7, UIUCDCS-F-85-937, Department of Computer Science, University of Illinois, Urbana, February 1985.
- 9 R. S. Michalski, "Understanding the Nature of Learning: Issues and Research Directions," ISG 85-8, UIUCDCS-F-85-938, Department of Computer Science, University of Illinois, Urbana, February 1985 and chapter in *Machine Learning: An Artificial Intelligence Approach* Vol. II, Morgan-Kaufman Publishers, 1986.
- 10 T. Dietterich and R. S. Michalski, "Discovering Patterns in Sequence of Events," *Artificial Intelligence Journal*, Vol. 25, No 2, pp. 187-232, 1985.
- 11 T. Dietterich and R. S. Michalski, "Learning to Predict Sequences," ISG 85-9, UIUCDCS-F-85-939, Department of Computer Science, University of Illinois, Urbana, February 1985.

- 12 R. Stepp and R. S. Michalski, "CONCEPTUAL CLUSTERING: Inventing Goal-Oriented Classifications of Structured Objects," ISG 85-10, UIUCDCS-F-85-940, Department of Computer Science, University of Illinois, Urbana, February 1985 and chapter in *Machine Learning: An Artificial Intelligence Approach*, Vol. II, Morgan-Kaufman Publishers, 1986.
- 13 R. S. Michalski, H. Ko and K. Chen, "SPARC/E(V.2), An Eleusis Rule Generator and Game Player," ISG 85-11, UIUCDCS-F-85-941, Department of Computer Science, University of Illinois, Urbana, February 1985.
- 14 I. Mozetic, "Compression of the ECG Knowledge-base Using the AQ Inductive Learning Algorithm," ISG 85-13, UIUCDCS-F-85-943, Department of Computer Science, University of Illinois, Urbana, March 1985.
- 15 R. E. Reinke and R. S. Michalski, "Incremental Learning of Decision Rules: A Method and Experimental Results," to be presented at the Machine Intelligence Workshop 11, Ross Priory, University of Strathclyde, March 1985.
- 16 R. S. Michalski, "Knowledge Repair Mechanisms: Evolution vs Revolution," ISG 85-14, UIUCDCS-F-85-946, Department of Computer Science, University of Illinois, Urbana, July 1985, and the *Proceedings of the Third International Machine Learning Workshop*, Skytop, Rutgers University, June 1985.
- 17 L. Rendell, "Utility Patterns as Criteria for Efficient Generalization Learning," UIUCDCS-R-85-1206, Department of Computer Science, University of Illinois, Urbana, April 1985, and the *Proceedings of the 1985 Conference on Intelligent Systems and Machines*, Rochester, MI, April 1985.
- 18 L. Rendell, "Induction, of and by Probability," UIUCDCS-R-85-1209, Department of Computer Science, University of Illinois, Urbana, July 1985, and the *Proceedings of the Workshop on Uncertainty and Probability in Artificial Intelligence*, Sponsored by AAAI and IEEE, University of California at Los Angeles, August 1985.
- 19 L. Rendell, "A Scientific Approach to Practical Induction," UIUCDCS-R-85-1211, Department of Computer Science, University of Illinois, Urbana, May 1985, and the *Proceedings of the Third International Machine Learning Workshop*, Skytop Lodge, Rutgers University, June 1985, revised version to appear in *Machine Learning: A Guide to Current Research*, T. Mitchell, J. Carbonell, and R. S. Michalski (Eds.), Kluwer Publishing Company, (forthcoming) 1986.
- 20 L. Rendell, "Genetic Plans and the Probabilistic Learning System: Synthesis and Results," UIUCDCS-R-85-1217, Department of Computer Science, University of Illinois, Urbana, July 1985, and the *Proceedings of the International Conference on Genetic Algorithms and their Applications*, Carnegie-Mellon University, July 1985, revised version to appear in *Machine Learning: A Guide to Current Research*, T. Mitchell, J. Carbonell and R. S. Michalski (Eds.), Kluwer Publishing Company, (forthcoming) 1986.

- 21 T. Fermanian, R. S. Michalski and B. Katz, "An Expert System to Assist Turfgrass Managers in Weed Identification," presented at the *Proceedings of the 1985 Summer Computer Simulation Conference*, Chicago, IL, July 1985.
 - 22 J. Becker, "Inductive Learning of Decision Rules with Exceptions: Methodology and Experimentation," ISG 85-14, UIUCDCS-F-85-945, Master of Science Thesis, Department of Computer Science, University of Illinois, Urbana, August 1985.
 - 23 B. Falkenhainer, "Quantitative Empirical Learning: An Analysis and Methodology," ISG 85-16, UIUCDCS-F-85-947, Master of Science Thesis, Department of Computer Science, University of Illinois, Urbana, August 1985.
 - 24 R. S. Michalski and P. H. Winston, "Variable Precision Logic," in *Artificial Intelligence Memo*, No. 857, MIT, Cambridge, MA, September 1985. An updated version will appear in *AI Journal*, 1986.
 - 25 R. S. Michalski and T. Durham, "Trying to Mimic the Mind," (An article about research done by R.S.M. and his collaborators), in *Computing Magazine*, April 18, 1985
 - 26 R. S. Michalski, "Equipment of Artificial Intelligence Laboratory," at the Department of Computer Science, 1985.
-
- 1986 1 D. L. Medin, Wm D. Wattenmaker and R. S. Michalski, "Constraints and Preferences in Inductive Learning: An Experimental Study Comparing Human and Machine Performance," ISG 86-1, UIUCDCS-F-86-952, Department of Computer Science, University of Illinois, Urbana, February 1986.
 - 2 M. Goldfain, "A Translation System from Clausal Form into Annotated Predicate Calculus," ISG 86-2, UIUCDCS-F-86-954, Department of Computer Science, University of Illinois, Urbana, February 1986.
 - 3 A. Collins and R. S. Michalski, "The Logic of Plausible Reasoning: An Advanced Report," ISG 86-3, UIUCDCS-F-86-951, Department of Computer Science, University of Illinois, Urbana, February 1986.
 - 4 R. S. Michalski, H. Ko and K. Chen, "Qualitative Prediction: A Method and A Program SPARC/G," ISG 86-4, UIUCDCS-F-86-942, Department of Computer Science, University of Illinois, Urbana, 1986, chapter in *Expert Systems*, C. Guetler (Ed.), Academic Press Inc., London (forthcoming).
 - 5 J. Hong, I. Mozetic and R. S. Michalski, "AQ15: Incremental Learning of Attribute-Based Descriptions from Examples, The Method and User's Guide," ISG 86-5, UIUCDCS-F-86-949, Department of Computer Science, University of Illinois, Urbana, May 1986.
 - 6 R. S. Michalski, "Machine Learning Research in the Artificial Intelligence Laboratory at Illinois," *Machine Learning: A Guide to Current Research*,

ISG 86-6, T. Mitchell, J. Carbonell, and R. S. Michalski (Eds.), Kluwer Publishing Company, 1986 (forthcoming).

- 7 H. Ko and L. Lee, "Toward a Practical Planning System for Assembly Task," ISG 86-7, UIUCDCS-F-86-957, Department of Computer Science, University of Illinois Urbana, 1986.
- 8 R. S. Michalski, A. B. Baskin, C. Uhrík, T. Channic, S. Borodkin, A. G. Boulanger, L. Rodewald, and R. E. Reinke, "A Technical Description of the ADVISE.1 Meta-Expert System that Integrates Multiple Knowledge Representations and Learning Capabilities," ISG 86-8, UIUCDCS-F-86-962, Department of Computer Science, University of Illinois, Urbana, March 6, 1986 (in preparation).
- 9 R. S. Michalski and R. Stepp, "INDUCE.3: A Program for Learning Structural Descriptions from Examples," ISG 86-9, UIUCDCS-F-86-960, Department of Computer Science, University of Illinois, Urbana, 1986 (in preparation)
- 10 L. Iwanska, "Critical Issues in Natural Language Understanding and their Importance to Machine Learning," ISG 86-10, UIUCDCS-F-86-953, Department of Computer Science, University of Illinois, Urbana, 1986.
- 11 P. Haddawy, "Implementation of and Experiments with a Variable Precision Logic Inference System," ISG 86-11, UIUCDCS-F-86-959, Department of Computer Science, University of Illinois, Urbana, 1986.
- 12 K. Chen, "Smooth Path Tracking through Symbolic Computations," ISG 86-13, UIUCDCS-F-86-963, Department of Computer Science, University of Illinois, Urbana 1986.
- 13 K. Chen, "The Inductive Acquisition of Temporal Knowledge," ISG-86-14, UIUCDCS-F-86-964, Department of Computer Science, University of Illinois, Urbana 1986.
- 14 G. Thornburg, "Where Expert Systems Intersect with Information Retrieval: Issues and Applications," ISG-86-15, UIUCDCS-F-86-956, Department of Computer Science, University of Illinois, Urbana, 1986.
- 15 R. S. Michalski, "Dynamic Recognition: An Outline of Theory of How to Recognize Concepts without Matching Rules, Intelligence Systems Group," ISG-16, UIUCDCS-F-86-965.
- 16 J. Hong, R. S. Michalski, and C. Uhrík, "An Extension Matrix Approach to the General Covering Problem," in *SPIE Applications of Artificial Intelligence III*, Vol. 635, 1986.
- 17 R. S. Michalski, "Concept Learning," to appear in *AI Encyclopedia*, January 1986.
- 18 R. S. Michalski and R. Stepp, "Clustering," to appear in *AI Encyclopedia*, January 1986.
- 19 R. Stepp and R. S. Michalski, "Conceptual Clustering of Structured Objects: A Goal-Oriented Approach," accepted for "AI Journal," 1986.

- 20 B. Falkenhainer and R. S. Michalski, "Integrating Quantitative and Qualitative Discovery: The ABACUS System," ISG-86-17, UIUCDCS-F-86-967, May, 1986.
 - 21 Skorstad, G and J. C. Skorstad, "Parallel Concept Learning on the Connection Machine," ISG 86-18, UIUCDCS-F-86-966, University of Illinois, Urbana, July 1986.
 - 22 Fermanian, T. W., R. S. Michalski and B. Katz, "Applications of Artificial Intelligence to Agricultural Problems," *International Congress on Computers in Biotechnology Abstracts*, Baltimore Convention Center, Maryland, January 30-31, 1986.
 - 23 R. S. Michalski, I. Mozetic, J. Hong, N. Lavrac: "The AQ15 Inductive Learning System: An Overview and Experiments," ISG 86-20, UIUCDCS-R-86-1260, Department of Computer Science, University of Illinois, Urbana, 1986.
 - 24 R. S. Michalski, "Two-Tiered Concept Meaning, Inferential Matching and Conceptual Cohesiveness," Invited paper for the Allerton Conference on Analogy and Similarity, ISG 86-21, UIUCDCS-F-86-968, University of Illinois, Urbana, IL, June 12-14, 1986.
 - 25 C. T. Uhrík, "A Rule Exerciser for Knowledge Base Enhancement in Expert Systems" ISG Report 86-23, UIUCDCS-F-86-969, M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, September 1986.
 - 26 R. S. Michalski, I. Mozetic, J. Hong, N. Lavrac: The Multi-Purpose Incremental Learning System AQ15 and its Testing Application to Three Medical Domains. Proceedings AAAI, Philadelphia, August 11-15, 1986.
-
- 1987 1 R. S. Michalski, A.B. Baskin, C. Uhrík and T. Channic, "The ADVISE.1 Meta-expert System: The General Design and a Technical Description," ISG Report 87-1, UIUCDCS-F-87-962, Department of Computer Science, University of Illinois January, 1987.
 - 2 J. A. Bentrup, G. J. Mehler and J. D. Riedesel, "INDUCE 4: A Program for Incrementally Learning Structural Descriptions from Examples," ISG Report 87-2, Department of Computer Science, February, 1987.
 - 3 A. B. Baskin, and R. E. Stepp, "The Role of Deterministic and Non-Deterministic Rule Scheduling in Expert Systems," Submitted to the Sixth National Conference on Artificial Intelligence.
 - 4 B. Katz, T. W. Fermanian and R. S. Michalski, "Agassistent: An Experimental Expert System Builder for Agricultural Applications," ISG Report 87-16, UIUCDCS-F-87-978, Department of Computer Science, University of Illinois, Urbana, IL, October, 1987.
 - 5 R. S. Michalski, "Two-Tiered Concept Meaning, Inferential Matching and Conceptual Cohesiveness," Chapter in the Book *Similarity and Analogical Reasoning*, Stella Vosniadou and A. Orton, (Eds.) Cambridge University Press.

- 6 R. S. Michalski, H. Ko and K. Chen, "Qualitative Prediction: the SPARC/G Methodology for Inductively Describing and Predicting Discrete Processes," Chapter in the Book *Current Issues in Expert Systems*, Edited by A Van Lamsweerde and O. Dufour.
- 7 R. S. Michalski, "Learning Strategies and Automated Knowledge Acquisition: An Overview," Chapter in the Book *Computational Models of Learning*, Edited by Leonard Bolc.
- 8 Bradley Lane Whitehall, "Substructure Discovery in Executed Action Sequences", UILU-ENG-87-2256 University of Illinois at Urbana-Champaign, September, 1987.
- 9 H. Wechsler and L. Zimmerman, Fault-Tolerant Recognition Using DAM's, 1st Int. Conference on Neural Networks, San Diego, CA.
- 10 H. Wechsler and L. Zimmerman, Distributed Associative Memories and Data Fusion, 2nd Int. Conference on Neural Networks, Boston, MA.

-
- | | | |
|------|---|--|
| 1988 | 1 | R. S. Michalski and Larry M. Watanabe, "Constructive Closed-Loop Learning: Fundamental Ideas and Examples", Reports of Machine Learning and Inference Laboratory, George Mason University (to appear). |
| | 2 | H. Ko and R. S. Michalski, "Types of Explanation and Their Role in Constructive Closed-loop Learning," Reports on Machine Learning and Inference Laboratory, George Mason University (to appear). |
| | 3 | R. S. Michalski and H. Ko, "On the Nature of Explanation or Why Did the Wine Bottle Shatter," Present Workshop on Explanation-based Learning, Stanford University, March 1988. |
| | 4 | R. E. Reinke and R. S. Michalski, "Incremental Learning of Decision Rules: A Method and Experimental Results", Chapter in Machine Intelligence 11, D. Michie and J. Michie (eds.), 1988 (to appear). |
| | 5 | F. Bergadano, S. Matwin, R. S. Michalski, J. Zhang "Representing and Acquiring Imprecise and Context-dependent Concepts in Knowledge-based Systems, invited paper, International Symposium on Methodologies for Intelligent Systems, Torino, 1988. |
| | 6 | F. Bergadano, S. Matwin, R. S. Michalski, J. Zhang, "Learning Two-Tiered Descriptions of Imprecise Concepts: A Method Employing Examples of Varied Typicality and An Optimized Base Concept Representation: Part I Principles and Methodology", Submitted to the Machine Learning Journal. |
| | 7 | F. Bergadano, S. Matwin, R. S. Michalski, J. Zhang, "Learning Two-Tiered Descriptions of Imprecise Concepts: A Method Employing Examples of Varied Typicality and An Optimized Base Concept Representation: Part II: " Algorithms and Experiments", Submitted to the Machine Learning Journal. |

- 8 Bergadano, F., Matwin, S., Michalski, R. S., Zhang, J., "Measuring Quality of Concept Descriptions", Accepted for presentation to the Third European Working Session on Learning, 1988.
- 9* Robert Stepp, Bradley L. Whitehall, Lawrence B. Holder, "Toward Intelligent Machine Learning Algorithms," University of Illinois at Urbana-Champaign, UILU-ENG-88-2221 May, 1988.
- 10* Lawrence Bruce Holder, Jr. "Discovering Substructure In Examples", University of Illinois at Urbana-Champaign, UILU-ENG-88-2223 May, 1988."
- 11* Robert E. Stepp, "Machine Learning from Structured Objects", University of Illinois at Urbana-Champaign, May, 1988.
- 12* Lawrence B. Holder, "Substructure Discovery in SUBDUE", University of Illinois at Urbana-Champaign, UILU-ENG-88-2220 May, 1988.
- 13* Bradley L. Whitehall, "Substructure Discovery of Macro-Operators", University of Illinois at Urbana-Champaign, UILU-ENG-88-2219, May, 1988.
- 14 Allan Collins and Ryszard S. Michalski, "The Logic of Plausible Reasoning: A Core Theory", to appear in Cognitive Science, Vol. 13, No. 2.
- 15* J. B. Sinclair & R. S. Michalski, "Computer-Based Consulting System For Diagnosing Soybean Diseases", Depts. of Plant Pathology and Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801-4709.
- 16 Anthony Richard Nowicki, "A Methodology for Representing Natural Language Expressions in Variable-Valued Logic," M.S. Thesis, University of Illinois, 1988.
- 17* Gregory Hugh Greene, "Quantitative Discovery: Using Dependencies to Discover Non-Linear Terms." M.S. Thesis, University of Illinois, 1988.
- 18 Jim Kelly, "A Plausible Reasoning System," M.S. Thesis, University of Illinois, 1988.
- 19* H. Ko, "Toward Integration of Planning and Learning from Experience in Robot Assembly," PhD Thesis, University of Illinois at Urbana-Champaign, 1988.
- 20 J. M. Char, V. Cherkassky, H. Wechsler, and L. Zimmerman, Distributed and Fault-Tolerant Computation for Retrieval Tasks Using Distributed Associative Memories (DAMs), IEEE Trans. on Computers, Vol. 37, No. 4, pp. 484-490.
- 21 H. Wechsler and L. Zimmerman, 2-D Invariant Object Recognition Using Distributed Associative Memories, IEEE Trans. on Pattern Analysis and Machine Intelligence.
- 22 H. Wechsler (Ed.), NEURAL NETWORKS for VISUAL and MACHINE PERCEPTION, Oxford University Press, to appear 1989.

- 23 H. Wechsler and L. Zimmerman, Distributed Associative Memories for Bin Picking, IEEE Trans. on Pattern Analysis and Machine Intelligence, to appear, 1989.
- 24 J. M. Char, V. Cherkassky, and H. Wechsler, Fault-Tolerant Database Using Distributed Associative Memories, to appear in Information Sciences, 1989.
- 25 H. Wechsler and L. Zimmerman, REDUNDANT VISION FOR ROBOTICS, in Robots with Redundancy, A. K. Bejczy (Ed.), Springer-Verlag, to appear, 1989.
- 26 H. Wechsler, COMPUTATIONAL VISION, Academic Press, to appear in 1989.

* Research done at the University of Illinois at Urbana-Champaign under a joint grant with the AI Center at George Mason University.

Appendix 3

Selected Publications

October 5, 1988

MACHINE LEARNING:

The State of the Art and Research Directions

Ryszard S. Michalski
Center for Artificial Intelligence
George Mason University
Fairfax, VA 22032

Contents:

1. Motivation and Goals
 2. Research Orientations and Potential Applications
 3. The Essence of Learning Processes
 4. Accomplishments: Selected Examples
 5. Major Research Themes
 6. Special Methodologies
 7. Summary and Promising Research Directions
- References

1. Motivation and Goals

Machine learning, a branch of artificial intelligence, is concerned with developing computational theories of learning, and building machines with learning capabilities. Learning capabilities include self-acquisition of general knowledge from examples or specific facts, self-restructuring of knowledge to perform better a given task, skill acquisition with practice, or discovery of new knowledge by observation or experiment. The research takes a broad view, studying any type of a learning process, no matter whether it occurs in nature or not.

Since 1980, when the first Machine Learning Workshop was organized at CMU (by Carbonell, Michalski and Mitchell), machine learning has experienced an extraordinary growth. Subsequent workshops were held at the University of Illinois at Urbana-Champaign (1983), at Rutgers University (1985), and at the University of California at Irvine (1986). This year in June, the 5th International Conference on Machine Learning was held at the University of Michigan.

There are already several books and a few thousand publications in the area of machine learning. The number of researchers working in the field, both in the US and abroad, is growing rapidly. Following the US lead, there are now regular machine learning conferences on machine learning held in Europe. In July this year, the first summer on Machine Learning was organized in France to provide a review of the field for newcomers as well as researchers. Machine Learning is becoming one of the basic scientific challenges not only to AI, but to the whole information technology.

The reasons for the rapid growth of this field seem to be threefold. First is the realization that machine learning has

the potential to produce a solution to the knowledge acquisition bottleneck plaguing artificial intelligence systems. Second, that this research may give new insights on the improvement of teaching methods and better organization of human education. Third, that it may lead to a fundamentally new machine, a machine with great potential and value to human society.

Any machine developed by man can be viewed as a physical embodiment of an algorithm (idea, design) originally invented by its constructor. The machine so created does just (at best) what it supposed to, what it was designed to do. This idea seems to be deeply ingrained in all our present thinking about a machine. Success in building a truly sophisticated general-purpose learning machine will introduce an intrinsically new dimension to our technology. Such a machine will be capable of goal-oriented self-modification, and self-improvement with experience - features not possessed by any machine built so far.

The output from a learning machine cannot be predicted in advance, as such a machine is not just an executor of a build-in algorithm, but an active creator of new knowledge. A learning machine can improve its performance with practice and experience, make powerful generalizations of observed facts, and assist humans in discovering regularities heretofore unknown to its original human designer. Thus, it can potentially contribute genuinely new elements to the store of human knowledge.

Equipping machine learning programs with adequate natural language understanding capabilities may lead to systems able to acquire knowledge directly from texts and documents, as humans do. This would, of course, produce capabilities for machine acquisition and knowledge processing that are potentially unlimited.

Success of research in all basic research areas of AI is ultimately predicated on the progress in machine learning. This is specifically true for areas such as expert and advisory systems, autonomous intelligent robots, computer vision, speech understanding, natural language processing and intelligent tutoring. For example, it is not possible to build truly autonomous intelligent systems without learning capabilities, because the operating environment of such systems, and potential tasks that these systems face, cannot be completely predicted in advance. It is easy to find other examples indicating the need for machine learning in the areas mentioned above (Michalski, Carbonell and Mitchell, 1986.)

Despite particularly intensive research in machine learning for the last several years, we are still far away from building a universal learning machine. Many important research results, however, have been obtained. Several domain-independent and domain-dependent working learning systems have been implemented. There have been several simple but significant practical applications of these systems. These results, and the lack of theoretical bounds on the possible development of more powerful learning systems, inspire a belief that new achievements and new more powerful machine learning systems are forthcoming.

This article presents an organized critical analysis of some of the most important research developments in the field, a systematic classification of these developments, and a critical review of some of the more promising research directions and applications.

We start with outlining basic research orientations and discussing potential applications of the field.

2. Research Orientations and Potential Applications

The field of machine learning has three interrelated orientations:

- a theoretical one - to explore the space of possible learning mechanisms and their properties, no matter if these mechanisms exist or not in nature (similarly to the study of flight whether manifested in animals or not),
- a cognitive one - to develop comprehensive computational models of learning processes in biological systems, especially in man, and
- an engineering one - to build machines (computer programs or physical systems) capable of various forms of learning and discovery that are useful for solving specific practical tasks.

Better models of learning processes in humans can be expected to have many benefits, especially for improving education and training. As our society is increasingly information-based, the need for improvement in education, and training is gaining strongly in importance. Such an improvement is difficult to achieve, however, without deep understanding of principles and mechanisms underlying learning processes. Such an understanding may help us to resolve many traditional conflicts regarding the preferred educational approaches.

While learning processes have been intensively studied over many years in other fields, primarily in educational psychology and cognitive science, the machine learning approach offers special promise. The reason for this is that machine learning research produces well-defined

computational theories and their computer implementations, which can be tested, verified, and experimented with as objective autonomous entities. For example, the well-known controversy whether teaching methods should emphasize more student's experience or a systematic presentation of organized knowledge (Dewey vs. Hutchins), from the machine learning viewpoint can be viewed as the problem of comparing learning from examples and learning-by-being-told strategies. The machine learning research has been providing many new insights into the relative merits of these strategies.

In this report, our primary focus will be on theoretical and engineering aspects of machine learning. There is a growing technological need for constructing computer systems and devices with learning capabilities. Such systems will be able to improve with experience, ease the "knowledge acquisition bottleneck" for our knowledge-intensive artificial intelligence systems, in particular, expert and advisory systems. Such systems will make possible truly autonomous robots, and will enable us to build better computer vision and speech understanding systems.

With the proliferation of databases, and other means of recording information, we need better methods for extracting useful knowledge from these sources, we need systems capable of discovering regularities in large databases, or in complex behaviors of man-made or natural systems. In the context of an immense proliferation of human technical and other literature, there is a strong need for machines capable of extracting important knowledge contained therein, without a human expert and encoder.

The development of machines with learning capabilities will introduce very powerful new possibilities for better defence technology. The complexity of a modern

battlefield and warfare requires that military leaders have help from machines capable of intelligently interpreting multiple, often contradictory signals. Such machines will be capable of monitoring and detecting important patterns in rapidly changing situations, of summarizing important facts, and of alerting human users of unsuspected or unusual dangers.

The prospect of employing semi-autonomous or even autonomous robots in certain aspects of future battlefield also creates a need for learning machines. Such a machine could be used for automated surveillance, gathering and interpreting information, and creating maps and plans of the changing environment. Given a large amount of information about various aspects of the battlefield, it may help to produce a plan of action that would be difficult to formulate otherwise. It may detect unknown opportunities for action and surprise of the enemy.

The above tasks cannot be realized by programming all the needed knowledge and procedures into a machine in advance. To perform such tasks a machine has to have the capability for generalizing facts, detecting analogies with past situations, and the ability for self-modification and self-improvement. The success of the autonomous land vehicle and automated battlefield projects may eventually hinge upon the implementation of learning capabilities in these systems.

Modern aircraft are continually increasing in capability, as well in complexity. To control them a pilot needs correspondingly better tools. Analysis of accidents in both military and civilian aircraft have revealed that the cause of most of the accidents can be traced to human error. There is need for a pilot's computer assistant that can intelligently monitor the flight and all the procedures,

detect unexpected courses of events, both outside and inside of the plane, and provide support for conducting defensive actions. Again, the success of efforts to develop such a system (e.g., the pilot associate project) will likely require an introduction of learning capabilities into the system. The recent conference on the applications of machine learning to the pilot associate project, organized by the Honeywell team in June 87 at the University of California at Irvine, clearly indicates this understanding.

Naval operations, in response to the growing sophistication of self-guided missiles, need improved information-gathering and signal interpretation capabilities. The recent unsuspected attack on an American-flagged ship by an Iranian silkworm missile seems to be indicative of this. Such an improvement may eventually require systems that can detect regularities or irregularities in courses of events, or can recognize a missile from very limited or confusing signals. To improve target detection and recognition, as well as missile self-guidance, a learning mechanism may have to be employed. Such a mechanism would improve the missile's recognition system in the face of increasing task difficulty and unpredicted changes in the enemy behavior.

The above are just a few examples of potential applications of machine learning. The development of sophisticated general- and special-purpose learning systems will have a profound influence on the whole human technology, and will have significant social implications. These implications will present a new set of issues for statesmen, lawmakers, philosophers, theologians, etc. (e.g., Michalski, 1986; Schwartz, 1987).

3. The Essence of Learning Processes

A commonsense view of learning is that it is a process of improving performance with experience. A system (human, animal or machine) learns if it utilizes observed facts and past experience in performing a given task, so as to do better in subsequent performance of this task.

A more general view is that learning is creating representations of experience and storing them for future use (Michalski, 1986). It involves acquiring systematic knowledge or skill new to the learner from facts, experiences or partial knowledge. Learning thus occurs even if such an acquisition of knowledge or skill may cause no measurable improvement of performance, or even may decrease performance. The latter may happen due to the overhead associated with maintaining knowledge that has no effect on any defined performance.

Thus, a discovery of a new law or relationship, which may have no significance for performing any real task, a theoretically interesting but practically useless generalization of facts, a detection of surprising analogies, or formulation of general theories that are not practically useful are also forms of learning. Similarly, determining an explanation of given observations, which leads to better understanding, but not to an improvement of any defined task is a form of learning. On the other hand, if a system improves its performance by, e.g., increasing the speed of its hardware, rather than by increasing its knowledge or improving its algorithm, then such a process is not considered learning.

Learning processes involve a self-creation or self-modification of learner's knowledge structures. Thus, they require performing operations on the learner's knowledge.

These operations can be classified by the type of inference performed on these knowledge structures in the process of learning.

We can, therefore, classify learning processes by the major type of inference performed on the input information to derive the knowledge that learner actually acquires for future use. This view leads to a well-known classification of learning strategies into

- learning by direct knowledge implantation,
- learning by instruction,
- learning by deduction,
- learning by analogy, and finally,
- learning by induction and discovery.

In this order, these strategies put a decreasing demand on the teacher, and an increasing demand on the learner (Michalski, Carbonell and Mitchell, 1983).

4. Accomplishments: Selected Examples

While traditional computer science and artificial intelligence deals primarily with representing knowledge on a computer, and deriving from it other information, machine learning operates at a metalevel, that is, it deals with problems of self-acquisition and self-modification of that knowledge. This represents an increased level of complexity of research endeavor, and therefore significant results and practical applications are correspondingly more difficult to come by. Here is a sample of results, illustrating some of the achievements, the diversity of applications, and methodological milestones in the development of the field.

- *Samuel's checker playing program (1959)*

This was one of the earliest and most significant programs employing several learning techniques, such as signature tables for rule learning, a parametric learning to improve the evaluation function, and rote learning, to remember the backup value of the evaluation function for selected positions.

- *Learning in the Blocks World (Winston, 1970)*

Winston's research on learning from examples showed how a learning program can compare, modify and generalize complex structural descriptions in the process of learning. The program used a semantic network representation of concepts and exhibited an incremental learning. The network was modified each time a new example contradicted the representation, or was not covered by the representation. The concept of *near-miss*, i.e., of an object that fails to be a concept example because of one feature, has been demonstrated to be an useful tool for teaching a system by examples.

- *Determining Cleavage Rules for Mass Spectrometer by MetaDendral (Buchanan and Mitchell, 1978).*

MetaDendral was a large-scale domain-specific learning program for determining rules for breaking molecules in a mass-spectrometer simulator. The program learned rules from noisy examples by using a model-driven strategy. It showed significant performance.

- *AQ11 and Soybean Rule Diagnosis (Michalski, 1976 and Michalski&Chilausky, 1980).*

AQ11, one of a series of domain-independent learning programs has demonstrated a high capability for learning general decision rules from examples of decisions. The program uses the *star method*, based on a repeated determination of sets (called, stars) of consistent and general descriptions of selected single examples. The descriptions, expressed in an extension of many-valued logic, called VL1 (variable-valued logic one), could be optimized according to a variety of problem-oriented criteria. In one application, diagnostic rules for soybean disease diagnosis determined by the program outperformed rules obtained by interviewing an expert.

- *Version spaces & learning heuristics for integration (Mitchell, 1978 and 1983).*

Assuming that a concept description can be expressed as a unique conjunction, the search space for this description can be usefully bounded by the set of most general descriptions and most specific descriptions. Such a *version space* method was demonstrated to be useful for learning rules for symbolic integration, and other problems. In the system LEX, implementing the method, examples for learning were self-created by applying a standard heuristic search method. The generalization process relied primarily on the "climbing generalization tree" operator.

- *ID3 & chess position classification rules (Quinlan, 1983)*

Given examples of objects from different decision classes, one can classify them into individual classes by a series of tests, each involving measuring one attribute. ID3 is a program for creating such a tree on the basis of maximizing

the information provided by an attribute at any given level. Such trees can be created very efficiently, and therefore this approach has found useful applications in various problems where large sets of examples are easily available (e.g., in chess endgames or certain medical problems). The decision tree knowledge representation is, however, difficult to comprehend, and reason with. It is also not amenable for taking advantage of background knowledge or previously created tree in constructing an improved decision tree. For these reasons, this approach applies to only a limited set of problems.

- *Learning structural descriptions by INDUCE (Michalski & Larson, 1978)*

INDUCE was a large-scale domain-independent program for learning structural descriptions from examples. It used an approach similar to the AQ programs, but employed a more general description language, an extension of predicate calculus. The program was able to utilize user-supplied background knowledge to create variables not present in the input data (such a process is called *constructive induction*). It generated descriptions at high level of abstraction, which were optimized according to a problem - dependent criterion.

- *AM and Eurisko (Lenat, 1976 and 1983)*

AM was a large-scale program demonstrating the capability for creating concept descriptions, in a frame-based language, by running predefined heuristic rules in the domain of elementary mathematics and number theory. The program was able to synthesize a number of non-trivial mathematical concepts, such as the prime number or Goldbach theorem. The continuation of this research led to

Eurisko, a program for determining heuristic rules by conducting a metalevel search.

- *Construction of a cardiac arrhythmia knowledge base by inductive learning (Bratko, Mozetic, and Lavrac, 1986).*

A knowledge base characterizing the relationship between electric malfunctions of the heart and ECGs was created by generating examples of malfunctions using a qualitative model of the heart, and then using an inductive learning program (AQ15) to generate diagnostic and descriptive rules. The rules were in good agreement with the rules found in medical literature. The the resulting knowledge base consisted of about 90 rules (versus initial 900), and took only 1/30 of the memory of the initial base of examples (30 vs. 900) .

- *Acquiring Search Control Rules: the Prodigy System (Minton, Carbonell, Etzioni at al, 1987)*

The system uses a Strips-like problem solver, and an explanation-based learning facility for acquiring control rules from a problem solving trace. The goal of learning was to determine the most effective control rules. The system was tested in domains such as machine shop scheduling and 3-D robot assembly.

- *Learning VLSI Structure Manipulations (Lathrop and Kirk, 1987)*

This is a precedence-based learning system that captures and generalizes procedural knowledge about the design of VLSI structures. By applying a form of analogical reasoning the system permits a designer to simply point out a pattern and specify what design transformation are to be done when a pattern is found in a circuit. The system

facilitates an application of design transformations in common design problems, and helps to improve the design, especially at its near completion stage.

5. Major Research Themes

5.1 Symbolic vs. Numeric Learning

Learning methods can be divided to two broad categories: *symbolic* and *numeric*. Symbolic learning is concerned with acquiring or improving knowledge structures characterizing a set of data. These structures may be in the form of logic expressions, if-then rules, semantic networks, frames, grammars, decision trees or other knowledge representations. In contrast, numeric learning involves tuning weights, modifying the strength of connections or other numerically representable parameters in predefined knowledge structures, e.g., neural networks, threshold logic systems or polynomial equations.

Symbolic learning has been the central focus of the artificial intelligence approach to machine learning (e.g., Michalski, Carbonell and Mitchell; 1983 and 1986). Numeric learning has been primarily studied in pattern recognition, control theory (e.g., Narendra, 1986), and neural networks (e.g., McClelland and Rumelhart, 1986). Computer programs or devices exhibiting various forms of numeric learning are often referred to as *adaptive systems*.

There are also hybrid techniques which attempt to combine symbolic and numeric learning. For example, genetic algorithms (Holland, 1986), and the AQ15 learning algorithm (Hong, Mozetic and Michalski, 1986) include elements of both approaches, as they involve creating and modifying knowledge structures (rules), as well as calculating

numerical weights associated with them. This report addresses primarily symbolic learning, as the progress in symbolic learning seems to be crucial for the advancement of the field.

5.2. Empirical vs. Analytic Approach

Currently, the two most active research orientations in symbolic learning are: *empirical learning* and *analytical learning*. Empirical learning produces general hypotheses from selected observations and data. It employs an inductive learning strategy and does not require much background knowledge before learning can occur. Various forms of empirical learning include learning concepts from examples, quantitative discovery, conceptual clustering, part-to-whole generalization and prediction. (Empirical learning is sometimes called similarity-based learning. We will avoid this term, however, because it is misleading, as it characterizes only a small subset of methods using this approach.)

Analytic learning produces more operational or effective knowledge representations by reformulating its prior knowledge, using an observation or an example for guidance. It employs primarily a deductive learning strategy, and requires a significant amount of domain knowledge before any learning can occur. Its most popular form is called explanation-based learning, which in its pure form, requires complete and consistent domain knowledge before any learning can occur. By creating an explanation of an observation in terms of this knowledge, one can relax the conditions stated in a given observation, and create a more effective concept representation or procedure. Other forms of analytic learning include automatic programming and operationalization.

Empirical and analytic learning are complementary, as they address complementary goals, and are applicable under different assumptions. In empirical learning, the basic assumption is that the system operates in an environment that is not well-known, or deals with problems from a poorly understood domain. Therefore, its primary source of information is a set of data from observing the environment, or examples and facts provided by a teacher. By drawing inductive inferences from information supplied by the external sources, the learner acquires useful general knowledge (concept descriptions, decision rules, etc.).

Examples of empirical learning systems are the AQ family of attribute-based inductive learning program (e.g., Hong, Mozetic and Michalski, 1986) or the BACON programs for quantitative discovery (e.g., Langley, 1983). The system's background knowledge includes primarily the specifications of types and properties of descriptors used for describing objects, and of the preference criteria for choosing among candidate hypotheses. Such a system has the capability of learning at the *knowledge level*, i.e., to create new knowledge (Dietterich, 1986). Because of relatively small amount of background knowledge required, the system is easily applicable to many domains.

Empirical learning has, however, two major limitations. One is that generated descriptions typically involve the same descriptive terms (attributes) as those that have been used in characterizing the input data (usually, a subset of these attributes, or closely related attributes are involved). Therefore, the system's learning capabilities strongly depend on the quality and relevance of descriptive terms used in the original problem formulation. The second is the universal limitation of all processes of acquiring knowledge from data: the descriptions resulting

from learning have the status of hypotheses. They may be incorrect, and have to be tested and evaluated on new data.

In analytic learning, the assumption is that the system already possesses enough knowledge to accomplish a given task in principle, but this knowledge is not directly useful or efficient. By being exposed to an example of a task performance, or a concept, the system utilizes this knowledge to show (explain) how the observed example satisfies its knowledge. The resulting proof (explanation) is then used to create a more efficient or operational knowledge. Since the obtained knowledge is a deductive consequence of prior knowledge, the result has the same truth status as its original knowledge. Thus, if the original knowledge was correct, the derived knowledge is must be correct also.

Basic limitations of this approach are that before any learning can occur, the system must already know in principle what it is supposed to learn (in the sense of deductive closure), and that its application requires a significant effort of codifying the required knowledge to the system. Thus, this approach is limited to learning problems where there exists a well-understood domain knowledge, and it is not too difficult to supply this knowledge to the system, and the task is primarily to restructure the knowledge to a more adequate form.

In an analytic system, a large amount of prior knowledge must be encoded in the system beforehand, so that the system is able to analyze or explain an observed fact or example in terms of this knowledge. The system learns better representations of the knowledge supplied, but cannot in principle create any new knowledge, i.e., knowledge that is not deducible from what it already knows (Dietterich, 1986). It cannot create its own domain

knowledge needed for learning, unless some meta-domain knowledge is inputted to it first. Such learning is oriented toward an improvement of performance, rather than toward the creation of new knowledge.

Many, if not most, learning situations fit neither purely empirical nor purely analytical paradigm. Often a learner starts with some prior knowledge about the concept or task to be learned, and then incrementally improves, corrects, updates or restructures its knowledge in the course of interacting with a teacher or environment.

For the above reasons, more recent research has been trying to extend or combine empirical and analytic paradigms. Researchers in empirical learning have been investigating ways in which the system may use effectively prior general or domain-specific knowledge in the process of developing a hypothesis (e.g., Michalski, 1983). Similarly, recent work on explanation-base learning, that represents the most investigated form of analytic learning, have almost abandoned the pure form of this paradigm, and started to concentrate on various forms of introducing hypothesis formation and knowledge debugging methods (e.g., DeJong, 1988).

An example of early work on an extensive use of background knowledge in inductive learning are various versions of the INDUCE program for inductive learning of structural descriptions from examples (e.g., Michalski and Larson, 1987; Michalski, 1983; Mehler, Bentrup, and Riedesel, 1986). This program includes capabilities for using background knowledge for constructing new attributes and descriptive terms not provided in the input data (the so-called *constructive* induction).

A general theoretical framework for inductive learning (e.g., Michalski, 1987) describes how background knowledge, e.g., domain-specific knowledge, causal or other knowledge, can be used for abductive (or knowledge-rich inductive) learning. It describes such learning as a process of hypothesizing new knowledge (H) that together with background knowledge (BK) makes the observations (O) a logical consequence of H and BK:

$$H \& BK \mid < O$$

where $\mid <$ denotes the logical entailment.

Whenever prior knowledge is incomplete, incorrect or not useful, in order to create new or better knowledge the system must be able to perform some form of inductive inference (generalization or abduction) involving the source data and the prior knowledge. This may be empirical inductive learning, if the prior knowledge is very limited. Or it can be constructive inductive learning that utilizes background knowledge to create new terms and high level concept descriptions.

In general, if the prior knowledge is insufficient for understanding a new concept or performing a given task, the system needs to acquire or hypothesize additional knowledge. If the prior knowledge is inconsistent or incorrect, then the learner tries to improve it. If the prior knowledge is not very useful because of the task complexity (the *intractable domain* problem; Mitchell, Keller and Kedar-Cabelli, 1986), the system develops heuristics to cope with the complexity.

Also, the system must be able to validate any created or corrected knowledge, and then integrate it with previous knowledge, so that this enhanced knowledge can be used in

subsequent learning. A system with the above described capabilities can be called a *constructive closed-loop learning* system (CCL), or briefly, a *constructive learning system*.

If a CCL learning system faces a learning task for which it has sufficient knowledge, but not in the most effective form, then it acts as an analytic learning system. It is easy to implement this feature in such a system, because as a closed-loop system it must have the ability to test if its knowledge satisfies the demands of a task. Thus, to be an analytic system, it needs only the ability to restructure or optimize its knowledge on the basis of this testing.

When the system's knowledge is insufficient or inadequate, then the system must perform inductive learning, because it must learn at the knowledge level. Thus, a constructive learning system integrates empirical learning, analytical learning and constructive induction. Figure 1 illustrates the above described three learning paradigms - empirical, analytic and constructive.

6. Special Methodologies

6.1. Genetic Algorithms

When a learning system operates in an environment that is unknown and very difficult to predict, then it needs to rely on a considerable amount of exploratory search. Genetic algorithms employ for this purpose an analogy from biological evolution as a natural model of adaptation to environments that change in unpredictable ways.

There are currently two primary approaches to applying genetic algorithms to machine learning tasks: the *classifier approach* developed by Holland at the University of Michigan, and the *production system* approach developed by Ken DeJong at George Mason University. The classifier approach explores the analogy between individuals in a population and single rules in the knowledge base. The production system approach treats sets of productions as individuals.

Both approaches seem to be very promising, the second one (the production system approach), however, is less well understood, as it was proposed only recently. The classifier approach, which modifies only single rules, is viewed as better for incremental local tuning of knowledge. However, if the task requires global structural changes of knowledge, or exploration of unknown subareas of the representation space, then it is expected that the production system approach may be more effective.

A general form a genetic algorithm can be characterized as an iterative repetition of the following steps:

1. An initial pool of rules is given. The rules have associated weights reflecting their usefulness, or accumulated value in past applications.
2. The system evaluates the rules in the process of using them for performing some task. The results of evaluation are used to modify the weights of the rules.
3. A subset of "best" rules is selected using the weights as defining a probability distribution (i.e., the probability of selecting a rule is proportional to its weight; thus even rules with small weight have a chance to be selected).
4. Selected rules are modified by various "genetic" operators, such as mutation (a random change in a rule), crossover (a mutual replacement of some parts in a pair of rules) and inversion (changing the location of a part of a rule within the rule, so that it will affect the subsequent crossover operation).
5. New rules replace randomly selected rules in the rule base so to keep the total population of rules more or less constant.

A genetic algorithm can be viewed as a form of inductive learning, in which there are no a priori examples given that guide the process making improvement of current knowledge. Instead, the system makes certain modifications in the current knowledge without any examples, and then tests the modified knowledge on new facts, to see if the modification lead to an improvement. The modifications are not random, but reflect to some extend the past experience of the system. Taking into consideration the past experience (in a statistical form) speeds up the learning process, but, in general, genetic algorithms tend to be rather inefficient, requiring a large

number of iterations and tests. Another weakness of this approach is that it is unclear how to communicate prior general and domain knowledge into such a learning system. Also, the representation of knowledge used in the system is usually rather low level, and difficult to comprehend.

Genetic algorithms are primarily oriented toward dynamic learning situations in which a system must use the current knowledge (even if only partially correct) and at the same time tries to improve it. Such tasks are typical for real-time problem solving (e.g., real-time optimization) and skill improvement. Current research in this area is concerned with making the rule language much richer, beyond the traditional binary strings, and with an appropriate modification of rule transforming operators.

6.2 Neural Networks

Neural networks consists of a large number of interconnected units performing the same neuron-like function, e.g., a sigmoig or a threshold logic function (McClelland and Rumelhart, 1986). The connections between the units may have different strength. These networks perform a generalization of the input signals, so that after the initial training they can react to signals previously unencountered. Thus, a connectionist network can be viewed as performing a form of induction. The resulting generalizations are encoded within the network by appropriately modifying the weights interconnecting the units, rather than represented as logic-like expressions, as in symbolic learning.

These neural models seem to have significant applications in classes of problems involving transformation and

processing of signals and low level information (e.g., low level visual or auditory signals).

Among major weakness of this methodology are the difficulty of encoding any prior general, or domain knowledge within such a model, and the incomprehensibility of the knowledge represented by weights of connections in a network. It is unclear how to represent abstract knowledge in such networks, in particular, knowledge that requires quantifiers.

7. Summary and Promising Research Directions

To characterize most promising research directions, let us use a general model that distinguish five basic components in a learning process (figure 2):

- a source of information (teacher or environment),
- a knowledge base that stores acquired knowledge, or a procedure for performing a skill),
- an inference engine that derives desired knowledge from the information provided by the source,
- a selection module that evokes from the Knowledge Base a subset of knowledge that is relevant to a given learning task, and
- an evaluation module that evaluates, criticizes and tests the knowledge acquired.




Figure 2.
A General Diagram of a Learning System

Most of the past research has been oriented toward a development of methodologies for empirical learning, i.e., to implement an inductive inference engine that utilizes relatively little background knowledge. More recent work was on analytic methods (such as explanation-based learning) that utilize a large amount of background domain-oriented knowledge, and their extensions. New research needs to address problems of integrating these two, complementary paradigms.

Also, most machine learning research assumed that concepts learned are representable by a precise, well-defined and context independent descriptions. In contrast,

human concept are typically flexible and context-dependent. It is important to undertake research on methods for learning such flexible, imprecisely defined and context-dependent concepts. Recent work on two-tiered concept representation and learning represents an example of such research.

Here is a summary of some of the most important general topics for further research:

1. multistrategy learning systems, in particular, integrated empirical and analytic systems (e.g., the above-described constructive learning systems),
2. methods for learning imprecise, flexible and context dependent concepts,
3. methods for determining relevant background knowledge for any given learning task,
4. determining criteria and methods for evaluating and testing learned knowledge,
5. the development of standards for judging the progress of the field, and a theoretical unification and clarification of various approaches,
6. further exploration of various current topics, such as learning by analogy, integrated quantitative and qualitative discovery, conceptual clustering, qualitative prediction, and their relationships.
7. an investigation of abstract concept representation in neural networks, and methods for transforming the knowledge incorporated in such networks into comprehensible logic-style or linguistic expressions, as well as methods for reverse transformation

8. research in specific problem domains to demonstrate the applicability and usefulness of machine learning for solving various practical tasks.

Among many promising problem domains are:

- construction of knowledge bases for expert systems from examples of expert decision making
- learning and improving plans for an autonomous robot
- 3D visual object recognition and scene analysis via learning
- learning control algorithms for automatic assembly
- conceptual analysis of large databases
- self-improvement of robot navigation procedures
- robot object manipulation, grasping and delivery
- quantitative and qualitative discovery of laws in a variety of applied fields, such as economy, business, microbiology, medicine, agriculture, chemistry, and others.

In every problem domain listed above, one may distinguish a number of specific research tasks that would constitute a challenging problem for machine learning research.

References

- Bratko I., Mozetic I. and Lavrac, N., Automatic Synthesis and Compression of Cardiological Knowledge, in Expert Systems: Automating Knowledge Acquisition, D. Michie and I. Bratko (eds.), Addison-Wesley Publishing Company, 1986.
- Brazdil, P. (ed.) Machine Learning, Meta-Reasoning and Logics, Workshop Proceedings, Sesimbra, Portugal, February 1988.
- Carbonell, J., Derivational Analogy: A Theory of Reconstructive Problem solving and Expertise Acquisition, in *Machine Learning: An Artificial Intelligence Approach volume II*, R. S. Michalski, J. G. Carbonell, T. M. Mitchell (eds.), Morgan Kaufmann, 1986.
- Cohen, P.R. & E. A. Feigenbaum (eds.), The Handbook of Artificial Intelligence, vol. III, chap. 14, William Kaufmann, Inc.
- DeJong, K. A., On Using Genetic Algorithms to Search Program Spaces, Proceedings of the 2nd Intern. Conference on Genetic Algorithms, Lawrence Erlbaum Publishing, July 1987.
- DeJong, G. and Mooney, R. Explanation-based Learning: An Alternative View, Machine Learning Journal, 2, 1986.
- DeJong, G. (ed.), Proceedings of the Symposium on Explanation-based Learning, Stanford University, March 22-24, 1988.
- Dietterich, T. G. & Michalski, R. S. A comparative review of selected methods for learning structural descriptions. In *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. G. Carbonell, T. M. Mitchell (eds.), Tioga Publishing Co., January 1983.
- Dietterich, T. G. (1986). Learning at the knowledge level. Tech Report 86-30-1, Dept. of Computer Science, Oregon State University.
- Haase, K. W. (1986). Discovery systems. AI Memo 898, Artificial Intelligence Laboratory, MIT .

Grefensette, John J. , Genetic Algorithms and their Applications: Proceedings of the Second International Conference on Genetic Algorithms, July 28-31, 1987, Massachusetts Institute of Technology, Cambridge (Lawrence Erlbaum Associates, Publishers, Hillsdale, New Jersey 1987).

Holland, J. H., Escaping Brittleness: The of general Purpose Learning Algorithms Applied to Parallel Rule Based Systems, In *Machine Learning: An Artificial Intelligence Approach volume II*, R. S. Michalski, J. G. Carbonell, T. M. Mitchell (eds.), Morgan Kaufmann, Inc., 1986.

Hong, J., Mozetic, I., and Michalski, R.S., AQ15: Incremental Learning of Attribute-based Descriptions from Examples, Reports of Intelligent Systems Group, ISG 86-5, UIUCDCS-F-86-949, Department of Computer Science, University of Illinois, 1986.

Keller, R. M. (1987). Defining operationality for explanation-based learning. Proceedings of the Sixth National Conference on Artificial Intelligence (pp. 482-487). Seattle, Washington: Morgan Kaufmann.

Ko, H., Michalski, R. S., Types of explanation and their role in constructive closed-loop learning, Reports of Machine Learning and Inference Laboratory, January 1988.

Lenat, D. B. The role of heuristics in learning by discovery: three case studies. In *Machine Learning: An Artificial Intelligence Approach, Volume I*, Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (Eds.) Morgan Kaufmann, 1983.

McClelland, J.L. and Rumelhart D. E., *Parallel Distributed Networks*, MIT Press, 1986.

Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (eds.), *Machine Learning: An Artificial Intelligence Approach (volumes I and II)*, Morgan Kaufman Publishing Co., 1983 and 1986.

Michalski, R. S., Theory and Methodology of Inductive Learning, *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. G. Carbonell, T. M. Mitchell (eds.), Tioga Publishing Co., January 1983.

Michalski, R. S., Concept Learning, *Encyclopedia of Artificial Intelligence*, E. Shapiro (ed.), John Wiley & Sons, January, 1987.

Michalski, R. S., Ko, H. and Chen, K., Qualitative Prediction: The SPARC/G Methodology for Describing and Predicting Discrete Processes, in *Expert Systems*, P. Dufour and A. Van Lamsweerde (eds), Academic Press Incorporated, pp 125-158, 1987.

Michalski, R.S. and Ko, H., On the nature of explanation or Why did the wine bottle shatter, A paper presented at the AAAI Workshop on Explanation-based Learning, Stanford University, March 1988.

Michalski, R. S., Watanabe, L., Constructive Closed-loop Learning: Fundamental Ideas and Examples, Reports of Machine Learning and Inference Laboratory, January 1988.

Minton S., Carbonell J. G., Etzioni O., Knoblock C. A., and Kuokka D. R., Acquiring Effective Search Control Rules: Explanation-Based Learning in the Prodigy System, Proceedings of the Fourth International Workshop on Machine Learning, University of California, June 22-25, 1987.

Mitchell, T. M. (1977). Version spaces: a candidate elimination approach to rule learning, Fifth International Joint Conference on Artificial Intelligence, (pp 305-310). Morgan Kaufmann Publ, Cambridge, Massachusetts.

Mitchell, T. M., Keller, T., and Kedar-Cabelli, S., Explanation-Based Generalization: A Unifying View, *Machine Learning Journal*, vol 1, 11-46, January 1986.

Narendra, K.S., Adaptive Learning Systems: Theory and Applications, Plenum, New York, 19886.

Quinlan, R. J., Learning Efficient Classification Procedures and their Application to Chess En Games, in *Machine Learning: An Artificial Intelligence Approach, Volume I*, Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (Eds.) Morgan Kaufmann, 1983.

Rumelhart, D. E. & Norman, D. A. (1978). Accretion, tuning, and restructuring: three modes of learning. In J. W. Cotton and R. Klatzky

Learning Two-tiered Descriptions of Flexible Concepts

*A Method Employing Examples of Varied Typicality
and a Two-staged Construction of the Base Concept Representation*

Part I: Principles and Methodology

F. Bergadano¹, S. Matwin², R. S. Michalski, J. Zhang

Artificial Intelligence Center
George Mason University,
Fairfax, VA 22030

Abstract

A method for learning flexible concepts is described, that is concepts that are imprecise and context dependent. The method is based on a *two-tiered concept representation*. In such a representation the first tier, called the *Base Concept Representation*, describes typical properties of a concept in an explicit, comprehensible, and efficient form. The second tier, called the *Inferential Concept Interpretation*, contains inference rules and metaknowledge that define allowable transformations of the concept under different contexts, and handle exceptional instances.

In the method, the first tier is created in two stages. In the first stage, a complete and consistent description of the concept is learned by applying the inductive learning methodology (AQ and INDUCE) to examples of varying typicality. In the second stage, so obtained description is optimized through a heuristic search, employing a *description quality criterion*. The second tier is defined by an expert under the guidance from the system, which asks the expert to explain the context-dependent meaning or special cases of the concept. Alternatively, the second tier can be inherited from more general concepts. This part of the paper concentrates on basic ideas behind the method, and gives illustrative examples. Part II of the paper describes algorithms, their implementation, and experimental results.

¹ On leave from Universita di Torino, Torino, Italy

² On leave from University of Ottawa, Ottawa, Canada

1. Introduction

Most methods of machine learning research assume that concepts are precise entities, representable by a single symbolic description. In such a representation, the boundaries of a concept are well-defined and context-independent. All instances of a concept are assumed to be equally representative. If an instance satisfies the given concept description, then it belongs to the concept, otherwise it does not. Some of these assumptions are relieved in works that assume that the concepts are defined by a probability distribution (e.g., Cheeseman et al., 88) or set membership function (Zadeh, 74). However, once such a probability distribution or a set membership function is defined explicitly for a given concept, the concept again has a fixed meaning, independent of the context in which it is used. Moreover, such concept representations remain inadequate for handling exceptional cases, or for capturing increases of knowledge about the properties of the concept.

In contrast, most human concepts have a context-dependent meaning and lack precisely defined boundaries. The imprecision of the boundaries seems to have a logical rather than probabilistic character. That means that the classification of instances of imprecise concepts typically involves logical, rather than probabilistic inference. For example, a human trying to resolve the problem if a given object is a desk or a table will evoke logical arguments rather than a probability value.

Examples of human concepts are usually not all equivalent. They may have different *degrees of typicality* in representing the concept. For example, a robin is conventionally viewed as a more typical bird than a penguin or an ostrich. Also, under different contexts the "bird" concept may apply to a live, flying bird, a picture or a sculpture, a chick hatching out of the egg, or even an airplane. Thus human concepts are *flexible*, as they adapt to the context in which they are used. It is clear that in order to handle such flexible concepts, machine learning systems need to employ richer concept representations than are currently used. Developing methods for acquiring flexible concepts and reasoning with them is thus an important goal in the new phase of machine learning research.

One problem that arises is how to represent such concepts. There have been a number of attempts to deal with representing imprecise concepts. Multiple-valued logic (e.g. Rine, 77) introduces additional truth-values, which represent different degrees of certainty with which an instance is believed to belong to a given concept.

The fuzzy set approach (Zadeh, 74) introduces a numerical degree of membership for examples in the concept. This set-membership function is described by people describing the concept, and thus is subjective. The fuzzy set approach allows one to express varying degree of membership of instances in the concept, for example, in the concept "tall person." It does not provide, however, appropriate mechanisms for expressing and handling the context-dependence of the concept meaning.

When the knowledge is incomplete rather than imprecise, it may be appropriate to make default assumptions about the world. In non-monotonic reasoning systems (e.g., Doyle, 79; McCarthy, 80; deKleer 86; Reiter, 87) these assumptions are revised if they can be refuted, i.e., if their negation can be proven. The non-monotonic approach to knowledge representation usually adds significant complexity and computational cost to knowledge representation system, and does not appropriately handle the context-dependency nor express degrees of uncertainty.

The starting point of research presented here is the idea of the *two-tiered concept representation* (Michalski et al. 86). In this representation the total meaning of a concept consists of two components, the *Base Concept Representation* (BCR) and the *Inferential Concept Interpretation* (ICI). The BCR defines the most typical or ideal properties of the concept. The ICI makes the boundaries of the concept flexible by describing allowed modifications of the concept's features in different contexts. Consequently, learning any concept requires constructing the appropriate BCR and ICI.

In many cases, the ICI is common for a class of concepts: e.g., a method for matching an instance of a specific liver disease with descriptions of different liver diseases will usually be the same for every liver disease. The ICI may also be inherited from more general concepts. For example, many properties of trees are the same as those of plants. The process of inferential matching goes beyond the standard partial matching based on the probability that an instance covers a given concept. It may involve any type of reasoning: deductive, analogical, or inductive inference.

Early ideas on learning two-tiered concept representations were presented in (Michalski, 88a) and closely related earlier work (Michalski et al., 86; Michalski, 87). An intriguing result of that research was that a substantial reduction of the description complexity can be achieved using even a very simple version of such a representation, without affecting its performance. This paper is an extension and continuation of these early ideas. Important advances are the introduction of a very general

description quality measure, the use of a rule base for performing the ICI, and the development of a heuristic search procedure that explores the trade-offs between the BCR and the ICI.

The *general description quality* (GDQ) measure takes into consideration the accuracy, the simplicity of the total description (BCR+ICI), the computational cost and the comprehensibility. The introduction of such a general evaluation measure allows us to redefine the concept of learning (Bergadano et al., 88a). Namely the learning activity is seen as the multistep process of improving the initial concept description (e.g., a set of instances of the concept, or an incomplete or inconsistent description) in terms of the above mentioned description quality measure.

To demonstrate these ideas we have built a system and then applied it to selected problems. Two problems in particular have been investigated: the problem of learning the concept of a "chair" and the problem of learning the concept of a labor contract.

2. Two-tiered Concept Representation

Traditional work on concept representation assumes that the whole meaning of a concept resides in a single stored structure, e.g. a semantic network that captures all relevant properties of the concept. (e.g. Collins and Quinlan, 72, Minsky, 75, Sowa, 84). The process of recognizing a concept involves simple matching of the stored representation with the perceived facts. Such matching may include comparing features of concept descriptions, or tracing links in networks of concepts, but has not been assumed to involve any complex inferential processes.

On the other hand, our knowledge can be viewed as a combination of two components, recorded knowledge and inferential extension, i.e. knowledge that can be created from recorded knowledge by conducting inference. This view leads us to the proposition that the meaning we assign to a concept in any given situation is a result of an interplay between two parts. The first part represents what the agent knows, or remembers. The second part represents what the agent can infer from his knowledge, using rules of inference.

The same principle of two-tiered representation seems to apply to individual concepts. In order to investigate the consequences of this conjecture, Michalski (87) has proposed a two-tiered knowledge representation. A concept description is split into two-parts: a Base Concept Representation (BCR) and an Inferential Concept Interpretation (ICI). The

BCR defines the concept explicitly, by giving a description of the concept in terms of either the attributes observed in the example, or in terms constructively learned during concept formation. The prototypical instances of the concept are classified by matching with the BCR. Characteristics of the concept represented in the BCR tend to capture the principle, the ideal or the intention behind the concept.

Anomalies, exceptions and context-dependencies are covered by a reasoning process that uses information contained in the ICI. The ICI deals with exceptions by inferring that they are instances of the concept (concept extending), or that they ought to be excluded from the description supplied by the BCR (concept shrinking). The ICI is used in the process of assigning the meaning to a concept using the BCR and the context. This process involves the background knowledge and relevant inference methods that allow the recognition, extension, or modification of the concept meaning according to context. When an unknown entity is matched against the BCR, it may satisfy it directly, or it may satisfy some of its inferential extensions. During the process of interpreting the ICI, one may use a probabilistic inference based on a simple distance measure (so called flexible matching (Michalski et al. 86)), analogical reasoning, inductive reasoning, or deductive reasoning to classify "special" uses of concepts.

Let us illustrate the idea of two-tiered representation with the concept of chair. A two-tiered representation of the chair concept could have the following form:

BCR: A piece of furniture.

Purpose: to seat one person.

Structure: seat, four legs, and a backrest.

(A picture of a typical chair, or a description of the relationship among the parts may be included).

ICI: no-of legs may vary from 1 to 4

the shape, the size, the color and the material of all components can vary as long as the function defined in the BCR is preserved

(chair without the backrest) ---> (stool rather than chair)

(chair with arm-rests) ---> (chair specializes to armchair)

(context = museum exhibit) --> (chair is not used for seating
any more)

(context = capital punishment) --> (specializes to electric_chair)

(context = toys) --> (dimensions can be much smaller, but other
structural properties are preserved.

Does not serve for sitting by normal

persons, but by correspondingly small
dolls)

(a part of the chair is broken) --> (a broken chair)

This simple example illustrates several important features of the two-tiered representation. If recognition time is important, only BCR will be used to match an example. If more time can be allocated, or if a more precise classification is required for a given event, ICI is used. When interpreting the ICI, one relies on background and general knowledge, and on the context in which the concept operates. Contexts can have hierarchical organization. Finally, ICI rules may chain, although it is not shown in this simple example.

Some systems that generate and use two-tiered representations have been described in the literature (Michalski et al., 86; Bergadano et al., 88b; Bergadano, Giordana, to appear). Also, the work on pruning ID3 trees in order to avoid overfitting (e.g. Cestnik et al., 87) can be viewed as related to this approach. However, pruning the trees results in the loss of coverage of concept instances. The approach presented here does not experience this problem, since the instances removed from the BCR are covered by the ICI.

Two-tiered concept descriptions are usually simpler, easier to understand and more efficient to use than the conventional ones. They also exhibit performance improvement on a testing set. In the systems developed so far, the ICI includes only a flexible matching function. More importantly, in their quality evaluation measures, these systems do not take into account the inferentially covered parts of concept descriptions. Improvement in quality is therefore measured only by the improvement in the first tier.

3. The learning system: a general overview.

3.1. The General Architecture.

The learning system presented here produces two-tiered concept descriptions by performing inference on examples obtained from the source. Fig. 1 presents the general architecture of such a system.




Fig. 1. General architecture of a system for learning two-tiered concept descriptions.

The core of the system is an inference engine that learns concept descriptions in a two-tiered format, consisting of the Base Concept Representation (BCR) and the Inferential Concept Interpretation (ICI). Different types of inference are involved in obtaining a two-tiered representation.

The BCR describes typical and easy-to-define meaning of the concept. Inductive inference is appropriate to perform this task. The inductive inference performed here is knowledge intensive. First the background knowledge can be used as the semantic criterion to restrict the search space. For example, we could determine that one attribute is more important than the others through the deductive inference performed on the background knowledge. Second, the background knowledge can be used to guide the constructive induction. The knowledge intensive method can help the system to learn the BCR from a small set of typical examples.

The ICI is used mainly to handle the special and rare cases. It is obtained from exceptions and nontypical events. Their number is usually limited, and therefore deductive or analogical inference is used to acquire some of the deductive rules. One of the methods to deductively obtain the rules can be outlined as follows: the background knowledge is used to produce the explanations of those exceptions and the ICI rules are generated from the explanations. The explanations can also be obtained from human experts.

Concepts in a given domain usually form a hierarchy, ordered by the

relation of one concept being a specialization of another concept. Some ICI rules of a concept may then be inherited from a higher-level concept.

Moreover, as indicated with the arrows in Fig. 1, both parts of the two-tiered description obtained by the system can contribute to the improvement and enrichment of the background knowledge.

3.2. Architecture of the Experimental System.

An experimental version of the system implements the architecture presented in Fig. 1. Table 1 specifies the input, output, and the function of the system.

Input:

Examples obtained from the source.

Output

BCR and ICI for the concept.

(The BCR is an efficient and comprehensible representation of typical instances of the concept. The ICI provides inferential means for assigning instances to classes, and interpreting concept exceptions.)

Function;

Initial concept description is provided by an inductive learner

This description becomes the root of a tree-like search space

Non-root nodes of the search space are obtained by simplification of the ancestor nodes

Search is guided by the General Description Quality measure.

A specification of the input, output and the function of the system

Table 1

Fig. 2 shows the design of the experimental version of the system. There are two differences between the general architecture presented in Fig. 1, and the specific experimental system presented in Fig. 2 that implements it. First, the experimental system relies on an inductive learning system, such as AQ15 (Michalski et al., 86) or INDUCE (Hoff et al., 82), to learn a concept description from examples. This description is treated as the BCR of the initial two-tiered representation, whose ICI is empty. Second, the experimental system relies on the expert to provide rules that explain special events.




Fig. 2. Design of the experimental version of the system

The BCR is represented in disjunctive normal form described in VL_1 or VL_2 notation (Michalski, 83). The ICI consists of two parts: a flexible matching function and deductive rules. The system learns the BCR and the deductive rules. The flexible matching function is predefined. Corresponding to the three kinds of knowledge, three types of matching are defined:

1. Strict matching: the event matches the BCR exactly,
2. Flexible matching: the event matches the BCR through a flexible matching function and
3. Deductive matching: the event matches the concept through deductive reasoning by using the deductive rules.

The system attempts to build a better quality two-tiered representation of the concept under consideration, using a heuristic search procedure. The search space is a potentially infinite set of all possible two-tiered descriptions of a given concept. The search procedure searches only a finite fragment of the space. Only those descriptions that can be obtained from the initial description by applying search operations are involved in the search. The search process is driven by the General Description Quality

(GDQ) of the descriptions obtained at each step of the search.

3.3. General Description Quality.

The notion of general description quality of concepts (Bergadano et. al. 88a) is formalized and discussed in more detail in Part II. In order to understand how it influences the search, it suffices to say here that the quality of a concept description depends on the accuracy, the comprehensibility, and the cost of the description. Covering the typical instances of the concept by the BCR, and the less typical ones by the ICI, also improves the quality of a description.

The accuracy of a concept description reflects the degree to which the description relates to the concept it describes. In the case of concept learning from examples, accuracy depends on the completeness and consistency of the description with regard to learning examples. It also depends on the typicality of the examples it covers, and the justification that can be constructed for the description. If the description can be plausibly justified in terms of the domain knowledge, the confidence in its correctness will increase.

The comprehensibility of the acquired knowledge is related to subjective and domain dependent criteria. An important requirement of an AI system is that knowledge has to be explicit and easily understandable by human experts. This is important for improving or modifying the knowledge, and for communicating with experts. Therefore knowledge acquired automatically should be easy to understand, should contain the descriptors most frequently used by experts, and should not be syntactically too complex. In practice, only the last feature is easy to ensure.

The cost captures the properties of a description related to its storage and use. Other things being equal, descriptions which are easier to store and easier to use for recognizing new examples are preferred. When considering the cost of a description, two characteristics are of primary importance. The first one is the cost of measuring the values of variables occurring in the description. In some application domains, e.g. in medicine, this may be a very important consideration. The second one is the computational cost of evaluating the description. Again, certain applications in realtime environment, e.g. speech or image recognition, may impose constraints on the evaluation time of a description.

Both the BCR and the ICI are parts of the concept description and they are used together in concept recognition. They influence each other. It

is not necessarily true that if a BCR performs well with one ICI, then it also performs well with a different ICI. For example, the experiment described in Part II showed that a BCR performed poorly with an empty ICI, but it performed well with a flexible matching function. Furthermore, in order to learn a better two-tiered concept description, the distribution between the BCR and the ICI should be adjusted during learning. Therefore they should not be learned separately and they should be related in some way during learning. Most of the current learning systems learned the BCR and the ICI independently, and flexible matching is only applied in the performance element. In our approach, when computing GDQ of a concept description, both the BCR and the ICI are considered. When we compute the completeness and consistency, the type of matching (strict matching, deductive matching and flexible matching) of an event is taken into account. The comprehensibility and cost of the ICI are also considered.

3.4. Learning the Base Concept Representation.

An elementary search operation may either specialize or generalize a description. The heuristics used at a given step of the search (Part II) decide which operation is applied. In the experimental version of the system, generalization is implemented as selector truncation, and specialization is implemented as complex truncation. Another operation, referent modification, simplifies the range of a selector, and may behave either as a generalization, or a specialization, depending on the selector relation (see table II). For instance, if the selector is

[size = 1..5, 7]

then referent modification giving

[size = 1..7]

is a generalization, since the cover is extended. On the other hand, if the selector is

[size $\leq$ 1..5, 7]

then the same referent modification represents a specialization, since the cover shrinks. Table II summarizes the implementation of generalization and specialization operators in the existing system.

generalization:	selector truncation	referent modification
specialization:	complex truncation	referent modification

Table 2. Implementation of the search operators

3.5. Learning the Inferential Concept Interpretation.

After a search operation is applied on the BCR of a description d , referred to as BCR_d , a new BCR may be either more specialized or more general than BCR_d . If the description is more specialized, some positive events previously covered by BCR_d may not be covered any more. The coverage by the concept description is arranged by building ICI rules that will cover them.

On the other hand, when a description obtained from BCR_d is more general than BCR_d , some new events, previously not covered by BCR_d , may have been added. These events could be positive as well as negative. If negative events have been added as a result of generalization, they will have to be excluded from the set of events covered by the concept description by means of the ICI rules. There are two types of rules: rules that cover a positive example otherwise left out of the BCR, and rules that eliminate a negative example from the BCR. Rules of the first type are referred to as extending rules, and rules of the second type are called shrinking rules.

In order to exclude or cover an event by the ICI part of a concept description, one has to obtain rules that will match the event, and perform the action necessary for the exclusion or coverage of the event. These rules, or their chains that ultimately lead to a conclusion regarding the membership of an event in the concept, are treated as an explanation of the event. The rules used in an explanation can be obtained in several ways.

First, suppose that the system is given an instance of a broken chair. More precisely, the description of the instance matches neither the BCR nor the ICI of the chair, because one of the four legs is shorter than the other three. Suppose that the ICI of the leg concept will deduce from this description that the shorter leg is broken. The chair concept is a specialization of the higher level concept furniture. Suppose further that

the ICI for furniture has the following rule: if some part of an object that otherwise is an instance of furniture is broken, then the object still is an instance of furniture. The following rule may now be inherited by the ICI of the chair concept: if some part of an object, that otherwise matches the BCR of chair, is broken, then the object is a chair.

Second, an analytic learning system, such as described in (Mitchell et al., 86; DeJong 86), can be applied to obtain an explanation. In these systems, a typical event is explained by deductively inferring it from the underlying background theory. The result is an operational rule for the concept. This rule is then generalized, e.g. using techniques described in (Mitchell et al. 86; Prieditis and Mostow 87). In our approach, only nontypical events will be subject to explanation. The purpose of the explanation here is to justify the special character of the event explained, rather than to operationalize the proof of its membership in the concept.

Finally, explanation of an event may be obtained from an expert, as is the case in knowledge acquisition for expert systems. The experimental system described here uses this method in its initial version.

3.6. A Summary of the Design of the Experimental System.

Once the ICI rules associated with the new description are known, its GDQ is computed. The GDQ value directs the best-first search. A heuristic, described in sec. 5, stops the search process. When the search process stops, the selected node defines the BCR and the ICI of a good concept description. As in the general architecture of Fig. 1, if these give rise to interesting generalizations that impact the general knowledge of the system, this knowledge will be modified. On the other hand, the general knowledge is used in the process of explaining the special cases to be covered by the ICI. Moreover, the general knowledge is necessary to perform constructive induction in the inductive learning phase that produces the initial description.

The architecture presented above involves both analytic and empirical learning, and may support learning in a multi-concept environment. Moreover, its constructive induction feature implements a feedback between the previously learned knowledge and the future learning. Such an architecture satisfies therefore the basic tenets of the constructive learning approach as understood in (Michalski and Ko, 88). Given the encouraging results with the experimental version of the system, a more complex design of an integrated learning system, that does not rely on an inductive learner for its input, is considered.

4. Examples

This section presents some examples of two-tiered representations obtained using the experimental system. Learning the concept of a labor-management contract provides a suitable and interesting application of learning a two-tiered concept representation. The nature of the domain is such that the example can be described using an attribute-based language, such as VL₁. Consequently, the learning program AQ15 is used to obtain the initial concept description. The application, discussed in Part II, is the natural extension of this example.

The second example, learning the concept of a chair from structural examples, requires the use of structure-based language, such as VL₂, since relations between objects of given types have to be represented. Consequently, the learning program INDUCE is used to produce the initial concept description.

4.1. Learning two-tiered Description of Labor-management Contracts.

Labor-management contracts usually show a number of typical characteristics. Among those are: general wage increase, job security, and pensions. Meeting all those demands by the management would result in an "ideal" contract, from the labor point of view. In practice, labor demands are usually scaled down during negotiation, which results in a contract.

It is not uncommon, however, to see contracts that exhibit very nontypical characteristics. These exceptional contracts may be explained by the context and background knowledge. For instance, a contract which is highly unsatisfactory in both wages and pension areas, but offers some job security, may be accepted during a deep recession. On the other hand, given an exceptionally good economic environment in an industry where labor supply is scarce, any contract proposal that is not highly satisfactory in all three areas may be unacceptable. Furthermore, a contract that is recognized as acceptable by the first tier representation may actually be evaluated negatively in the second tier: it may imply micro-economic consequences overwriting initial values of its attributes. For all these reasons, a two-tiered representation seems appropriate when learning the concept of a contract.

The example space is divided, from the labor point of view, into acceptable and unacceptable contracts. Both are agreements negotiated between a trade union and the management of an organization. The

former have furthermore been ratified by the general union membership, while the latter have been rejected.

In sequel we present a set of simple examples: they describe specific contracts, where only some selected characteristics of a contract are given. Those characteristics pertain to seven chosen attributes of a contract: general wage increase (gwi), cost of living allowance (cola), job security (job_sec), retirement age (ret_age), extent of pension (p_ext), pension for overtime work (p_ovt), and fringe benefits. The following is therefore an example of an acceptable contract:

gwi	cola	job_sec	ret_age	p_ext	p_ovt	fringes
2%	inflation+.5%	some	60	part	false	maintained

A number of other examples have been described in the same way, and are shown in Table. 3. The typicality of the last two positive and the last two negative examples was 0.5; the typicality of all the remaining examples was 1.0 (the highest value of typicality).

 acceptable-events

#	gwi	cola	job_sec	p_age	p_ext	p_ovt	fringes
1	2	ilpos	some	60	part	false	main
2	10	inf	good	65	none	false	main
3	2	inf	some	60	part	false	main
4	1	zero	good	65	full	false	main
5	10	zero	good	65	full	true	inc
6	3	inf	good	64	part	true	loss
7	6	zero	good	65	full	true	loss
8	14	ilpos	good	62	part	false	inc
9	4	ilpos	good	64	full	false	main
10	3	zero	some	58	full	false	inc
11	5	zero	some	60	full	true	loss
12	2	inf	good	61	part	false	main
13	11	zero	good	65	none	false	main
14	14	zero	good	65	none	false	main
15	11	ilpos	good	65	none	false	main
16	15	zero	some	65	none	false	loss
17	15	zero	some	65	none	false	inc
18	12	zero	some	61	part	false	loss
19	10	zero	none	65	full	false	loss
20	11	zero	none	65	full	false	loss
21	15	ilneg	none	58	part	false	inc
22	14	inf	none	61	part	true	loss
23	10	inf	none	60	full	false	main
24	7	ilpos	none	55	full	true	loss
25	6	inf	none	56	full	true	loss

unacceptable-events

#	gwi	cola	job_sec	p_age	p_ext	p_ovt	fringes
1	1..8	zero	none	60..65	none	false	loss
2	14	zero	none	58..65	none	false	loss
3	11	zero	none	65	none	false	loss
4	10	inf	none	64	none	true	main
5	2	inf	some	62	none	false	loss
6	0..4	zero	some	65	none	false	main
7	4	zero	good	60	none	false	inc
8	6	inf	some	65	none	false	inc
9	2	zero	good	65	none	false	main
10	14	zero	none	65	none	false	main
11	2...7	inf	none	60..65	full	false	main
12	0	ilneg	none	60	part	false	main
13	0	zero	none	65	part	false	main
14	3	zero	none	65	full	false	loss
15	2	zero	some	64..65	part	false	loss

Table 3. Cases of acceptable and unacceptable contracts.

The examples were submitted to AQ15, which produced a discriminant description of a concept of an acceptable contract (see Fig. 4).

```
[gwi <> 0..4 v 6] & [job_sec = some v good] v
[gwi <> 0..2 v 7] & [p_ext = full v part] v
[job_sec = some v good] & [p_ext = full v part] & [fringes = incr v maint] v
[gwi <> 10] & [p_ovt = true]
```

Fig. 3. The initial concept description for the labor-management contract

The quality of this description, according to the specific quality measure defined in Part II, is 0.989. This description is modified by the system in the following way:

1. Referent modification is performed on selectors. Close interval operation is applied on the first selector of the first two complexes, so that the complex [gwi <> 0..4 v 6] becomes [gwi > 6], and the complex [gwi <> 0 v 2..7] becomes [gwi > 7].

In this example, completeness and inconsistency are left unchanged by range modification, and simplicity (and comprehensibility) are improved.

```
[gwi > 6] & [job_sec = some v good] v
[gwi > 7] & [p_ext = full v part] v
[job_sec = some v good] & [p_ext = full v part]
```

Fig. 4. The improved concept description obtained by the system (BCR only)

2. The algorithm, described above, gives a modified description, shown in Fig. 5. The quality of this description is 0.976. This represents a small deterioration of quality, since two negative examples are covered by the modified description. Moreover, two positive examples are lost because the last complex of the description shown in Fig. 4 is truncated. The two negative examples are explained as exceptions. The reasoning about this exception is the following: even if the three attributes occurring in the rule of Fig. 5 (gwi, job_sec, p_ext) have acceptable values, when all the

other attributes have the worst values possible, the contract is not acceptable. This shrinking rule may be expressed as follows:

```
[gwi = very_low] & [cola = zero] & [p_age = very_high] & [p_ovt = false] & [fringes = loss]
--> unacceptable_contract
```

Is there a good extending explanation rule for the two events e_{24+} , e_{25+} covered uniquely by complex 4? These events may be covered by applying the following reasoning: if the pension offer is extremely good, and the state of the economy is good, the value of the gwi attribute does not matter anymore. The following rules convey this reasoning:

```
[p_age = very_low] [p_ext = full] & [p_ovt = true] -->
exceptional(pension)
exceptional(pension) ^ good_economy --> irrelevant(gwi).
```

Therefore, complex 4 is truncated and the above rules are added to the ICI. The final BCR obtained in this example is shown in Fig. 5. The BCR obtained represents a concept of acceptable contract which has high values of two of the three areas important for the union. The quality of this description is 1.0: it is complete and consistent, and the partition of events between the BCR and the ICI corresponds to the typicality of these events.

4.2. Learning a two-tiered description of the *chair* concept.

The instances of visual concepts present a high degree of variability, are affected by noise and are subject to modifications related to context. For this reason visual concepts can be better represented through a two-tiered scheme, allowing the system to capture the stable characteristics and reason about the special cases in a unified framework.

As in other pattern analysis problems, instances of visual concepts can be segmented into elementary components, and each component can be described by a list of attributes. Moreover, there can be spatial or functional relationships among the components, and attributes related to the whole instance can be present. It is important to point out that functional attributes have an important role in the representation of the instances, since they are the basis for the learning of a classification knowledge which is truly meaningful and understandable in a given domain. Human learning also seems to interleave the acquisition of concept descriptions and the detection of functional elements in the examples, without a well-defined separation or a clear ordering of the two activities in time.

Below we will give a simple example, where a two-tiered

representation of the concept of "chair" is learned from examples in a partially automated way. Fig. 5 gives an example of how a concept instance is described to the system. First of all, the example is divided into components (b1-b6) and the instance is said to contain them. Then every component is described through its shape attribute and its space relationships to the other components (ontop, attached from above). Moreover the predicate `person_can_sit_on` is used to specify a functional property of the component b5.




Fig. 5 Symbolic representation of a chair

Some other examples have been described in a similar way. These examples are given, in a pictorial form, in Fig. 5. There are 7 positive (1 thru 7) and 8 negative examples (8 thru 15) of the "chair" concept. The reader should realize that some of the information contained in the drawings is actually lost in the symbolic description and cannot possibly be obtained in the learned knowledge.




Fig. 6. Visual representation of examples used to obtain a two-tiered description of the chair concept.

These examples were given as input to INDUCE, together with the following constructive rule, as part of the background knowledge:

$[type(x)=line] \ \& \ [ontop(x,floor)] \Rightarrow [type(x)=leg]$

and the following discriminant descriptions were obtained, for the "chair" concept:

- 1: [type(b1)=leg] & [type(b2)=leg] & [type(b3)=leg] &
[person_can_sit_on(b4)]&[ontop(b4,b1^b2^b3)]&
attached_from_above(b5,b4)] & [type(b4)≠circle]
events covered: 1,3,5,6 - events uniquely covered: 1,5,6
- 2: [person_can_sit_on(b2)] & [ontop(b2,b1)] &
[type(b1)≠line] & [type(b2)=square]
events covered: 7 - events uniquely covered: 7
- 3: [type(b1)=semicircle]
events covered: 3,4 - events uniquely covered: 4
- 4: [attached_from_above(b2,b1)]&[ontop(b3,b2)]&[type(b3)=rectangle]
events covered: 2 - events uniquely covered: 2

These four conjunctive descriptions (or complexes) form a disjunctive normal form expression (or cover), that is consistent and complete with respect to the examples given in Fig. 6. In particular, complex 1 covers the positive examples 1,3,5 and 6, among which examples 1,5 and 6 are uniquely covered by this complex; complex 2 covers uniquely positive example 7; complex 3 covers positive examples 3 and 4, and covers uniquely example 4; complex 4 covers uniquely positive example 2. The induction process has effectively used the available background knowledge and the complexes that were produced also contain some functional description, but they still do not capture the most relevant characteristics of the concept of a chair. The reason is that the given examples contain a lot of information which was not really important, and an inductive learning system uses this information if it helps discriminating among the concepts. For example the presence of a component whose shape is semicircle actually covers examples 3 and 4 and no counterexamples in Fig. 6, but is by no means a plausible description of a chair. This is a fundamental problem in many applications of inductive learning, because the acquired knowledge will cover the given events but sometimes will not be understood and accepted by experts. This is often the case when only a limited number of examples is available.

In the following we will show how a two-tiered description of the concept "chair" can be automatically acquired from the description generated by INDUCE in the system described in sec. 3. The obtained concept description will then have two parts: the BCR and the ICI. The BCR for the "chair" concept will not necessarily be consistent nor complete, and some exceptional examples and counterexamples will need to be taken into account by a set of rules in the ICI.

Complex 1 is the most important one, since it covers four positive examples out of seven, and covers uniquely three of them. Since INDUCE only generates perfectly consistent complexes, complex 1 does not cover

any negative examples, but this may change when the two-tiered description is constructed, since some inconsistency might be introduced during the truncation procedure. The other complexes cover uniquely only one example each. Moreover if a semantic measure of comprehensibility is available, complex 1 could be evaluated as closer to our intuitive idea of what a chair is, since it contains the requirement that there have to be at least three legs, on top of which there is something a person can sit on.

The description generated by INDUCE, with an empty ICI is then used as a root node in the search process. The system will try to build a better quality two-tiered description of the concept "chair" from the initial description through the heuristic search discussed in sec. 3. The two-tiered description obtained by the system is shown in Fig. 7. It is obtained as follows. First, the last selector $[type(b4) \neq circle]$ of the complex 1 is truncated since it will improve the GDQ of the description most. In this way examples 2 and 4 will also be covered by the complex, which will now cover all the positive examples except for example 7, but one counterexample (event 13) will also be covered. Nevertheless, the overall quality of the complex is increased, after truncation of its last selector, also because it will then be possible to truncate complexes 3 and 4, since they will not uniquely cover any positive example. This will improve the simplicity of the description. Event 13 is then taken into account by the following ICI rules, that are introduced by a human user:

$$\begin{aligned} & \exists y [person_can_sit_on(y)] \ \& \ [attached_from_above(x,y)] \ \& \ [big(x)] \Rightarrow \\ & \quad [backrest(x)] \\ & \quad \neg \exists x [backrest(x)] \Rightarrow \neg chair \end{aligned}$$

these rules will allow the system to reason about "special" counterexamples such as event 13, by understanding that, in a chair, the backrest has to be sufficiently large. At this point, complex 2 can also be truncated, since it only covers one example, and the following two-tiered representation for the "chair" concept will be produced:

BCR: $\exists x, z \exists (\geq 3) y [person_can_sit_on(x)] \ \& \ [type(y)=leg] \ \& \ [ontop(x,y)] \ \& \ [attached_from_above(z,x)]$

ICI: $\exists y [person_can_sit_on(y)] \ \& \ [attached_from_above(x,y)] \ \& \ [big(x)] \Rightarrow [type(x)=backrest]$
 $\neg \exists x [type(x)=backrest] \Rightarrow \neg chair$
 $\exists (2)x [type(x)=wheel] \Rightarrow Irrelevant(\exists (\geq 3)y [type(y)=leg])$

Fig 8. Two-tiered representation of the "chair" concept

The last rule in the ICI has been introduced in order to cover the positive example that was covered by complex 2, before the truncation. In fact example 7 represents an exceptional condition and should not be included in the BCR; on the contrary, the rule in the ICI captures the relevant information that was present in the example. Event 7 is still a chair, but it is a special kind of chair, where legs have been replaced by two wheels. By using this rule as a rewriting rule, in conjunction with the BCR, event 7 will be covered as a positive example. BCR and ICI, if taken together, are a complete and consistent representation of the concept.

The description is now more intuitive, and the BCR seems to capture our natural understanding of the concept of a chair, while the ICI deals with special cases in an explicit way. The quality of the two tiered description would be higher because its accuracy is the same as the one of the description generated by INDUCE, but its comprehensibility is significantly better. This learning method seems to be appropriate in applications where a limited number of examples is available, and general and domain knowledge is needed in order to capture the most relevant aspects of the learning events.

5. Related Work.

The research presented here is different in several ways from the recent work in machine learning that investigates the effects of simplifying concept descriptions, e.g. (Fisher and Schlimmer, 88; Iba et al., 88). First, the method described here does not experience any loss of coverage as a result of description modification. This is a major difference between experimental results reported in Part II, and the findings of both (Iba et al., 88) and (Fisher and Schlimmer, 88). The reason is that in our approach events that lose their strict cover as the result of BCR simplification, become then covered by the ICI. Moreover, unlike (Fisher and Schlimmer, 88) and (Iba et al., 88), this approach to concept description simplification takes into account the typicality of events covered by the simplified description, thus preventing loss of coverage of typical events.

The experiments of (Fisher and Schlimmer, 88) in truncating the ID3's decision trees are based on a statistical attribute dependence measure that determines the attributes to be pruned. Because of its statistical character, there is a loss of predictive power when simplifying

descriptions learned on small training sets. As the experiments described in Part II indicate, the approach presented here does not seem to suffer from this problem.

The system developed by (Iba et al. 88) uses a trade-off measure that is similar to the GDQ measure proposed in this paper. The GDQ measure, defined in detail in Part II, considers more factors. Besides taking into account the typicality of the instances covered by the description, it considers the type of matching between an instance and a description. Moreover, the simplicity measured by the GDQ depends not only on the number of disjuncts in the description, as in (Iba et al. 88), but also on the different syntactic features of the terms in the description.

An important difference between the approach presented here and pruning of decision trees (Cestnik et al., 87) is lack of constraints on the part of the representation that is truncated when learning a two-tiered concept description. In post-pruning of decision trees, only paths ending in leaves may be truncated, which may improve the efficiency at the expense of the description quality. Moreover, pruning decision trees involves only generalization of the concept description, while the above method performs both generalization and specialization of the description.

Truncation of the BCR, obtained inductively from a small learning set does not affect predictive power if an adequate typicality measure is available. The existence of an adequate ICI further alleviates the problems resulting from induction with few examples.

The problem of defining and using the typicality of examples has been considered in the past both in machine learning and cognitive science. Negative examples of low typicality are referred to as *near misses* in Winston's system (Winston, 1975). Such examples, labeled by the user as near misses, are used in Winston's system to delineate the borders of a concept. Michalski and Larson, (78) introduced the idea of an outstanding representative of a concept. The concept of prototypical examples has been also studied by Smith and Medin (81) and by Roch and Mervis (75). Prototypical examples are fully specified instantiations of a given concept. In the method described here such prototypical examples do not need to exist, prototypical properties of a concept are learned automatically from examples of different typicality.

To summarize, there are four major differences between the work presented here and related research described in the literature. First, the above method does not experience a loss of coverage although it still yields a simpler description with improved predictive power. Second, it simplifies

the description by performing both generalization and specialization. Third, any part of the description may be truncated in the simplification process. Finally, the method takes into account the typicality of the examples.

6. Conclusions and Future Work.

The paper describes a method of learning two-tiered concept descriptions. The method is based on transforming an initial Base Concept Representation. The transformed BCR covers groups of events characterized by high typicality. It is also syntactically simpler, and therefore more comprehensible. A more complete coverage of the events from the learning set by the whole two-tiered description is achieved through inference. The method presented in this paper relies not only on the probabilistic inference, implemented as flexible matching in (Michalski et al., 86). It uses also a rule base for deductive inference. Deductive inference has the additional advantage of explaining why a given event is to be included (or excluded) from the cover.

In order to achieve this objective, transformations of BCR are implemented as truncations of the cover. The cover is provided by AQ15 in a standard, disjunctive normal form. The truncations either specialize the description (complex truncation), or generalize it (selector truncation). A search process, guided by a heuristic quality measure, is used to obtain a "good" description. The measure takes into account not only the explicit part of the description, but also the implicit one. The algorithms and heuristics are discussed in more detail in Part II of this paper.

The experiments that we have performed on real data indicate that the system is capable of learning good quality descriptions that are easy to comprehend and efficient to evaluate. In Part II, we describe an application in which a two-tiered description of a labor-management example has been learned from real data. The results indicate that recognition rate of two-tiered concept descriptions on the testing set is better than the performance of the normal, inductively learned descriptions on this set. This provides the evidence about the adequate predictive power of a two-tiered concept description of good quality. On the other hand, the BCR of the two-tiered description is much simpler than the cover learned by AQ15.

A number of problems remain to be addressed in the future. First, an integrated system that learns two-tiered descriptions needs to be designed and build. Such a system will have to satisfy a number of design goals: it will have to be incremental, learn descriptions of good quality, exhibit good predictive power, and be efficient. The incremental behavior

of the system would involve remembering examples that are neither strictly matched nor close (in the sense of the measure used by flexible matching) to the BCR. Then an explanation of such an example using the existing background and context knowledge, and the existing deductive ICI rules, is attempted. If no explanation is obtained automatically or from an expert, the example is stored. When the ICI grows significantly, there will be a repeated attempt to explain such an example.

Second, the problem of automatic acquisition of the ICI has to be investigated and tackled. The methods developed in Explanation-Based Learning will provide a good starting point. It has to be observed, however, that since the events to be explained are usually exceptions, the knowledge necessary to explain them may be lacking from the system.

Third, the system does not address the problems of dynamically emerging hierarchies of concepts. In the existing version the system only learns one concept at a time, and concepts do not change or split as new examples become available.

Finally, the system should be able to self-reorganize. The distribution of knowledge between the BCR and the ICI will be determined by the performance of the system on large testing sets. If it turns out, e.g., that some ICI rules are used very often, then these rules could be compiled into explicit BCR assertions. The representation obtained in this manner will be faster, but it will occupy more memory. It seems, therefore, that in concept representation one can trade one parameter against the other, within certain limits. This interesting research problem merits further investigation.

Acknowledgements

This research was done in the Artificial Intelligence Center of George Mason University. The activities of the Center are supported in part by the Defence Advanced Research Projects Agency under grant No. N00014-87-K-0874, , administered by the Office of Naval Research, and in part by the Office of Naval Research under grant No. N00014-88-K-0226. The first author was supported in part by the Italian Ministry of Education (ASSI) project, and the second author was supported in part by the NSERC grant number A2480. The authors thank G. Thornburg, K. Kaufman, H. Ko, A. Giordana, and F. Oppacher for useful criticism.

References

- Bergadano, F., Matwin, S., Michalski, R. S., Zhang, J., "Measuring Quality of Concept Descriptions", Proc. Third European Working Sessions on Learning, pp. 1-14, Glasgow, 1988a.
- Bergadano, F. , Giordana, A. , Saitta, L., "Automated Concept Acquisition in Noisy Environments", IEEE Transactions on PAMI, Vol 10, NO. 4, pp. 555-578, July 1988.
- Bergadano, F., Giordana, A., "Pattern Classification: An Approximate Reasoning Framework", International Journal of Intelligent Systems, (To appear).
- Cheeseman, P., Kelly, J., Self, M., Stutz, J., Taylor, W., Freeman, D., "AutoClass: A Bayesian Classification System", Procs. of the Fifth Int'l. Conf. On Machine Learning, Ann Arbor, pp. 54-64 (1988).
- Cestnik, B., Kononenko, I., Bratko, I., "ASSISTANT 86: A Knowledge-elicitation Tool for Sophisticated Users", Procs. of the 2nd European Workshop on Learning, pp. 31-45 (1987).
- Clark, K., "Negation as Failure", pp. 293-322 in Logic and Data Bases, ed. Gallaire, H. and Minker, J., Plenum Press, New York (1978).
- Collins, A. M. and Quillian, M. R., "Experiments on Semantic Memory and Language Comprehension" in Cognition, Learning and Memory, L. W. Gregg ed., John Wiley, 1972.
- de Kleer, J., "An Assumption Based TMS", Artificial Intelligence 28, pp. 127-162 (1986).
- DeJong, G., Mooney, R., "Explanation-based Learning: An Alternative View", Machine Learning, vol. 1, No. 2, 1986.
- Doyle, J., "A Truth Maintenance System", Artificial Intelligence 12, pp. 231-272 (1979).
- Fisher, D. H., Schlimmer, J. C., "Concept Simplification and Prediction Accuracy", Procs. of the Fifth Int'l. Conf. On Machine Learning, Ann Arbor, pp. 22-28 (1988).
- Hoff, B., Michalski, R. S., Stepp, R., "INDUCE-2 - A Program for Learning Structural Descriptions from Examples", TR 82-5, ISG, University of

Illinois, 1982.

Iba, W., Wogulis, J., Langley, P., "Trading off Simplicity and Coverage in Incremental Concept Learning", Procs. of the Fifth Int'l. Conf. On Machine Learning, Ann Arbor, pp. 73-79 (1988).

McCarthy, J., "Circumscription: a form of Non-monotonic Reasoning", Artificial Intelligence 13, 27-39 (1980).

Michalski, R.S., Larson, J. B., "Selection of Most Representative Training Examples and Incremental Generation of VL₁ Hypotheses: the Underlying Methodology and the Description of Programs ESEL and AQ11", TR 867, Dept. of Computer Sci., UIUC, 1978.

Michalski, R.S., "A Theory and Methodology of Inductive Learning", Chapter in the book "Machine Learning, an Artificial Intelligence Approach", Michalski, R. S., Carbonell, J. G., Mitchell, T. M. (Eds.), Tioga Pub. Co., Palo Alto, Ca, 1983.

Michalski, R. S., Mozetic, Hong, J.,I., Lavrac, "The Multi-purpose Incremental Learning System AQ15 and its Testing Application to Three Medical Domains", Proc. 5th AAI, pp. 1041-1045 1986.

Michalski, R. S., "Two-Tiered Concept Meaning, Inferential Matching and Conceptual Cohesiveness", Chapter in the Book "Similarity and Analogy", Stella Vosniadou and A. Orton, (Eds), 1988.

Michalski, R. S. and Ko, H., "On the Nature of Explanation, or Why Did the Wine Bottle Shatter", AAAI Symposium: Explanation-Based Learning, Stanford University, pp. 12-16, 1988a.

Michalski, R. S., "How to Learn Imprecise Concepts: A Method Employing a Two-Tiered Representation for Learning", Procs. of the Fourth International Workshop on Machine Learning, Irvine, CA, pp. 50-58,1987.

Minsky, M., "A Framework for Representing Knowledge", in The Psychology of Computer Vision, ed. P. Winston, (1975)

Mitchell, T. M., Keller, R. & Kedar-Cabelli, S., "Explanation-based Generalization: a unifying view," Machine Learning Journal, 1, 11-46 1986.

Prieditis, A. E. and Mostow, J., "PROLEARN: Towards A Prolog Interpreter

that Learns," Proc. of IJCAI 87, Milan, pp 494-498 1987.

Reiter, R., "Non-monotonic Reasoning", pp. 147-186 in Annual Review of Computer Science, (1987).

Rine, D. C., ed., "Computer Science and Multiple-valued Logic. Theory and Applications", North Holland, 1977.

Smith, E. E., Medin, D. L., "Categories and Concepts", Harvard University Press, 1981.

Rosch, E., and Mervis, C. B., "Family Resemblances: Studies in the Internal Structure of Categories", Cognitive Psychology, vol. 7, pp. 573-605, 1975.

Sowa, J. F., "Conceptual Structures", Addison Wesley, 1984.

Zadeh, L. A., "Fuzzy Logic and its Applications to Approximate Reasoning", Information Processing, North Holland, pp. 591-594, 1974.

Learning Two-tiered Descriptions of Flexible Concepts

*A Method Employing Examples of Varied Typicality
and a Two-staged Construction of the Base Concept Representation*

Part II: Algorithms and Experiments

F. Bergadano¹, S. Matwin², R. S. Michalski, J. Zhang

Artificial Intelligence Center
George Mason University
Fairfax, VA 22030

Abstract

This paper describes algorithms, their implementation, and experimental results from a system learning two-tiered concept descriptions. A powerful inference method for approximate matching and a general measure for evaluating concept description quality are presented. The learning process is implemented as a two stage procedure. In the first stage, a concept description is obtained from examples using an inductive learning system. The second stage of improving concept descriptions is implemented as a heuristic search in a search space controlled by the quality measure. The heuristics used to trim the search are presented. Experimental results provide evidence that two-tiered concept descriptions are not only simpler than the ones obtained from the existing inductively learning programs, but also have a better predictive power.

¹ On leave from Universita` di Torino, Torino, Italy

² On leave from University of Ottawa, Ottawa, Canada

1. Introduction

A theoretical framework and a method for representing and learning imprecise concepts using a two-tiered concept representation are described in Part I of this paper (Bergadano et al., 88a). In a two-tiered representation, the first tier, called the *Base Concept Representation* (BCR), describes typical properties of the concept in an explicit, comprehensible and efficient format. The second tier, called the *Inferential Concept Interpretation* (ICI), contains inference rules and metaknowledge that define allowable transformations of the concept under different contexts, handle special cases, and cover exceptional instances.

Part I also introduces a system for learning two-tiered concept descriptions. The learning process is organized as a search, guided by a *general description quality* measure. The search operations are truncations of the explicit descriptions. Instances of the concept that are not covered by the truncated descriptions are interpreted by different kinds of inference. The experimental system supports two kinds of such inference. The first kind allows for flexible, probabilistic matching of instances that are not strictly matched by the explicit description. The second kind of inference allows for deductive matching of exceptional instances, using a rule base.

This paper describes the algorithms implementing the above ideas, the two-tiered knowledge representation used, and the experiments with the implemented programs. The knowledge representation schemes used to express the BCR and the ICI are specified. A probabilistic inference method for dealing with approximate matching is introduced. This method is more powerful than the one used in (Michalski et al., 86), since it introduces the idea of the distance between a selector and an example. The deductive inference mechanism using a rule base is also discussed, and the kinds of rules supported by our system are defined. The General concept Description Quality measure (GDQ) is specified in detail, with emphasis on the quality of two-tiered concept descriptions. Finally, the paper presents experimental results of the implemented programs. In these experiments, the system has learned a two-tiered description of a labor-management contract from real-world data.

This study was motivated by the initial work of (Michalski et al., 86), which has demonstrated that by truncating the description obtained from an inductive learning program, one can obtain a significant reduction of memory requirements of a concept representation, without decreasing

its performance accuracy.

This paper explores the problem of optimizing an inductively-obtained concept description in greater detail, and it presents a more advanced method of its implementation. Our study confirms the earlier findings of (Michalski et al., 86). Experimental evidence described in sec. 5 of this paper shows that by applying a two-tiered knowledge representation, one can significantly decrease the complexity of a description, and at the same time maintain or even improve its accuracy in recognizing new instances of the concept.

2. Knowledge Representation

In this section, we describe the notation used to represent each of the tiers of a two-tiered concept description. Such concept descriptions are represented by two parts: the base concept representation (BCR) and the inferential concept interpretation (ICI). The ICI consists of a flexible matching function and a set of deductive rules.

2.1 Base Concept Representation

The BCR is represented in Variable-valued Logic System VL (VL_1 or VL_2) notation (Michalski 83). VL is a multiple-valued logic calculus with typed variables. The BCR for a concept is a disjunctive normal form which is called a cover. A cover is a disjunction of complexes. A complex is a conjunction of selectors. A selector is a form:

$$[L \# R] \text{ or } [L]$$

where

L is called the referee. In the first form, it is a variable or a function, and in the second form, it is a predicate.

R is called the referent. It is a set of values in the domain of the function or the variable of L.

is one of the following relational symbols : = < > ≤ ≥ ≠

Finally quantifiers can be applied to cover, complex and selector.

2.2 Inferential Concept Interpretation: Flexible Matching Function

A flexible matching function F is used as a part of the ICI and it is predefined. The flexible matching function F which is used in our current implementation matches concept descriptions from the set D with events from the set E:

$$F: D \times E \rightarrow [0, 1].$$

The value of F of a cover c and an event e is defined as the probabilistic sum of F of its complexes. If c consists of a disjunction of two complexes cpx_1 and cpx_2 , we have:

$$F(e, c) = F(e, cpx_1) + F(e, cpx_2) - F(e, cpx_1) * F(e, cpx_2)$$

There is one problem with this definition. Suppose that the cover c consists of many complexes and all values of F of the complexes are very small, say 0.2. Then the value of $F(e, c)$ is close to 1. This is not what we want. We solved the problem by providing a threshold t , such that any value of $F(e, cpx)$ which is smaller than t is treated as 0.

F of a complex cpx and event e is defined as the average of the F s for a conjunction of its constituent selectors, weighted by the proportion of positive examples covered by the complex:

$$F(e, cpx) = (\sum F(e, sel_i) / n) * \#cpxpos / (\#cpxpos + \#cpxneg)$$

where n is the number of the selectors in cpx and $\#cpxpos$ and $\#cpxneg$ are the number of positive examples covered by cpx and the number of the negative examples covered by cpx respectively.

F of an event e and a selector sel is defined by the degree of match between the selector and the event weighted by the coverage of positive and negative examples of the selector:

$$F(e, sel) = (1 - \text{DegMatch}(e, sel)) * (1 + (\#selpos/\#pos - \#selneg/\#neg))/2$$

where $\#selpos$ and $\#selneg$ are the numbers of positive examples and negative examples covered by the selector respectively. $\#pos$ and $\#neg$ are the numbers of the positive and negative examples, respectively. Suppose that selector sel is $[x = a_{j_1} \vee \dots \vee a_{j_m}]$. $\text{DegMatch}(e, sel)$ is then defined as follows:

$$\text{DegMatch}(e, sel) = \begin{cases} 1 & \text{if } x \text{ is nominal and } e \text{ is covered by } sel, \\ 0 & \text{if } x \text{ is nominal and } e \text{ is not covered by } sel, \\ \text{dis}(a_k, sel) / \max_{i=1, \dots, n} (\text{dis}(a_i, sel)), & \text{if } x \text{ is linear} \end{cases}$$

$$\text{dis}(a_k, sel) = \min_{k=j_1, \dots, j_m} (l_i - k_l)$$

The domain of x is the ordered list $(a_1, a_2, \dots, a_n)$, and a_k is the value of x of the event e . For example, if the domain of x is $[0 \dots 10]$ and the value of x for the event e is 4, then $\text{DegMatch}(e, [x = 2 \vee 5]) = (5-4) / \max_{i=0, \dots, 10}(\text{dis}(i, \text{sel})) = 1/5$.

We will not force the system to make a decision when the difference between the values of flexible matching function for two concepts is very small. If the difference is smaller than the preset threshold, the result will be no match.

2.3 Inferential Concept Interpretation: Deductive Rules

The ICI also includes a set of deductive rules, allowing the system to recognize transformed or special cases. In fact, the flexible matching is most useful to cover instances that are close to the typical case. For example, flexible matching could allow us to recognize a sequoia as a tree, although it does not match the typical size requirements, while deductive reasoning would be required to recognize a tree without leaves (in the winter time) or to include in the concept of tree some metaphorical meaning (e.g. a genealogical tree or a search tree).

The deductive rules in the ICI are expressed as Horn clauses. Inference on these rules is implemented using the LOGLISP inference system (Robinson and Sibert, 1982). Numerical quantifiers and internal connectives are also allowed (Michalski 1983).

Finally, the conclusion of a deductive rule can be the special form "Irrelevant(formula)", meaning that "true" can be substituted for "formula" when matching the BCR, if the antecedent of the rule holds. This is especially useful when expressing some transformation of the concept (e.g. $\text{season}=\text{winter} \Rightarrow \text{Irrelevant}([\text{has_leaves}])$). Similar transformation rules have been used in (Bergadano et al., 87). Other rules in the ICI can have as a conclusion either a general predicate (e.g. $\text{month}=\text{december}.. \text{march} \ \& \ \text{northern_hemisphere} \Rightarrow \text{season}=\text{winter}$) or the classification of the instance, i.e. the predicate $\text{classification}=\text{concept_name}$ or the predicate $\text{classification} \neq \text{concept_name}$ (e.g. $\text{context}=\text{Knuth's_book} \ \& \ \text{acyclic_graph} \Rightarrow \text{classification}=\text{tree}$). Rules in the ICI may chain, but simpler deductions are preferred, in order to make the classification easy to understand for the human users. Although the implementation supports recursion, non-recursive rules should be used when possible, and the number of rule activations should be limited.

2.4 Types of Matching

An event can then be covered by a two-tiered description through the following three types of matching:

1. Strict matching: the event matches the BCR exactly, in which case we say that the event is S-covered,
2. Flexible matching: the event matches the BCR through a flexible matching function, and we say that the event is F-covered.
3. Deductive matching: the event matches the concept through deductive reasoning by using the ICI Rules, and we say that the event is D-covered.

These three sets are made mutually exclusive: if an event is S-covered, then it is not D-covered or F-covered, and if an event is D-covered, then it is not F-covered. Thus, S-covered events are explicitly covered, and F-covered and D-covered events are implicitly covered.

3. Quality of concept descriptions

Our objective is to obtain concept descriptions of good quality, so the notion of quality has to be introduced, and an operational definition useable in our system has to be given. In sec. 3.1 we discuss the important aspects of description quality, and in sec. 3.2 we focus on these aspects of description quality which are relevant for measuring the quality of two-tiered concept descriptions. Finally, Appendix 1 contains the details of the General Description Quality measure that was implemented in the experimental system. Appendix 2 contains the details of the preference-based evaluation criterion that combines different characteristics into a single measure.

3.1 Criteria for Determining the Quality of Concept Descriptions

The quality of a concept description is influenced by three basic characteristics: the accuracy, the comprehensibility, and the cost. This section discusses these three components, as well as a mechanism for combining them into a single measure.

The accuracy represents the description's ability to produce correct classifications. A common way to prefer more accurate descriptions is to require that they be complete and consistent with respect to the training events (Michalski, 73; Mitchell, 77; Michalski, 80). Even if a description is

incomplete and inconsistent, the number of positive and negative examples it covers provides important information for evaluating its quality. In this case, we can measure the degree of completeness and consistency of a given description. If the description is also sufficiently general and does not depend on the particular characteristics of the training events, these measures can be a meaningful estimate of the accuracy of the description. In order to achieve completeness and consistency in presence of noise, one may generate overly complex and detailed descriptions. Such descriptions, however, may not perform well in future cases and examples. This is the well known phenomenon of overfitting (Watanabe, 69; Sturt, 81).

The comprehensibility of the acquired knowledge is related to subjective and domain dependent criteria. An important requirement of an AI system is that knowledge has to be explicit and easily understandable by human experts. This is important for improving or modifying the knowledge, and for communicating with experts. Since a black box classifier will not be accepted by experts verifying the knowledge of a performance element, therefore knowledge acquired automatically should be easy to understand, should contain the descriptors most frequently used by experts, and should not be syntactically too complex. In practice, only the last feature is easy to obtain.

The cost captures the properties of a description related to its storage and use. Other things being equal, descriptions which are easier to store and easier to use for recognizing new examples are preferred. When considering the cost of a description, two characteristics are of primary importance. The first one is the cost of measuring the values of variables occurring in the description. In some application domains, e.g. in medicine, this may be a very important consideration. The second one is the computational cost of evaluating the description. Again, certain applications in real-time environment, e.g. speech or image recognition, may impose constraints on the evaluation time of a description.

These criteria need to be combined into a single evaluation procedure that can be used to compare different concept descriptions. A possible solution is to have an algebraic formula that, given the numeric evaluations of single criteria, produces a number that represents their combined value. Examples of such approaches are multiplication, weighted sum, maximum/minimum, t-norm/t-conorm (Weber, 1983). Although these approaches are often appropriate, some of them may present disadvantages. Firstly, they usually combine a set of heterogeneous evaluations into a single number, and the meaning of this final number is hard to understand for a human expert. Secondly, they may force the

system to evaluate all the criteria, even if it would be sufficient to compare two given descriptions on the basis of the most important one, if one is much better than the other.

In order to overcome some of these problems, we use a lexicographic evaluation functional (LEF) (Michalski, 1972; Michalski, 83) that combines the above mentioned criteria. Appendix 2 discusses in detail the LEF, as well as its modification used in the described system.

The criteria discussed above can also be applied to two-tiered descriptions. The accuracy of the acquired knowledge does not only depend on the explicit information, but also on the implicit reasoning abilities. Inferential Concept Interpretation also affects cost, since it allows the performance system to use a simpler BCR, and reason about special details only in exceptional cases. Finally, the comprehensibility of a two-tiered representation must be carefully evaluated, since one of its implied goals is to state a clear and simple concept description in the BCR and to account for meaningful special cases through a reasoning process.

3.2 The quality of two-tiered concept descriptions

In the previous section, we proposed a general framework for evaluating the quality of concept descriptions:

$$\text{Quality}(\text{description}) = \langle (\text{Accuracy}, t_1) (\text{Comprehensibility}, t_2) (\text{Cost}, t_3) \rangle$$

which is evaluated using LEF with tolerances t_1 , t_2 , t_3 .

We will now apply the results of this discussion to the task of defining a quality measure appropriate to compare the quality of two-tiered concept descriptions.

The accuracy is the first criterion of concept description quality. Accuracy depends linearly on completeness and consistency of the description, as well as on the typicality of the events covered by the two parts of the description. In evaluating the accuracy of a two-tiered representation, we have to take into account the fact that degree of confidence in the results of inference decreases from deduction to induction (Michalski, 87). These requirements are met by making completeness and consistency dependent on the typicality of the covered examples and on the way these examples are covered. We assume that an expert can provide typicality of examples at the time they are presented to the system responsible for building the initial description. The experts are usually quite good at determining the typicality of events in their area of

expertise.

Completeness and consistency of a two-tiered description brings up additional requirements: a good representation should cover the typical examples explicitly, and the non-typical ones implicitly. Moreover, the coverage of typical negative examples in the BCR is particularly detrimental to the quality of the representation. This is important to accuracy because the BCR is mainly obtained or justified by the training events, on an inductive basis. Therefore, one can be confident in the information contained in the BCR only if a sufficient number of examples is available, or if the examples are really typical or representative for the domain. On the contrary, the ICI, being generated by experts or with the available domain knowledge, is appropriate when dealing with rare or exceptional events.

In general, descriptions that cover many typical positive events are most preferred. Completeness is therefore proportional to the typicality of the events covered. Moreover, if negative events are covered, the consistency of the description is inversely proportional to the typicality of the negative events covered.

It is also preferred that the typical events are covered by the BCR, and non-typical, or exceptional events are covered by the ICI. In fact, the BCR is inductively learned from the events provided by user, and it is more reliable when the training events are typical. The ICI, on the contrary, is deductively obtained from the background knowledge, or from a human expert, and relies more on general and domain knowledge. Generally, the ICI is more reliable when dealing with the special or rare cases, since experts often have difficulty in explaining large quantities of typical events. For these reasons, a typical positive explicitly-covered event should contribute to completeness more than implicitly-covered. And vice-versa, nontypical positive implicitly-covered events contribute to completeness more than explicitly-covered. These assumptions are reflected by weights w_S , w_F , w_D , used in the definitions of completeness and consistency (Appendix 1).

Furthermore, since ICI rules are obtained from background knowledge or from a human expert, they are more reliable than the flexible matching function. Consequently, a positive D-covered event should contribute to completeness more than F-covered. We may also observe that flexible matching is not very useful for exceptions whose typicality is very small. A similar argument holds for consistency.

Comprehensibility of a two-tiered representation takes into account

the operators occurring in both BCR and ICI, and has to weigh the relative contribution of each part to the comprehensibility of the whole description.

Finally, the notion of cost of a description as introduced in the previous section extends directly to two-tiered descriptions. Details are discussed in Appendix 1.

4. Improving the Quality of a Concept Description

Learning two-tiered concept descriptions is performed in two stages. In the first stage, a complete and consistent concept description is obtained from an inductive learning system. In our approach, we have relied on AQ15 (Michalski et al., 86a), and INDUCE (Hoff et al., 82) to obtain such descriptions. In this paper, we describe mainly the second stage, which improves the description obtained in the first stage with respect to its GDQ.

4.1 Search Strategies and Operators

The second stage of improving concept descriptions is seen as a state space search (Bergadano et al., 88a). This process is guided by the general description quality discussed in the previous section, and is implemented as a best first search, i.e., the descriptions of better quality are considered first. According to the nature of the quality measure, descriptions can be improved mainly by increasing their accuracy or by decreasing their complexity. For this reason the operators in the search simplify a given description by dropping (truncating) some of its components or by modifying the argument of some predicates. This does not always result in a loss of accuracy, especially when measured on a testing set of new examples, since simpler features might be more stable and depend less on the set of training examples. Table 1 describes the search process:

Search space:	a tree structure, in which the nodes are two-tiered descriptions (BCR + ICI) of a given concept.
Search operators:	selector truncation, complex truncation, referent modification.
Search strategy:	controlled by the quality measure, as described in sec. 4.2

Table 1

Top-level specification of the search algorithm performing optimization of concept descriptions

The goal of this procedure is not necessarily to find an optimal

solution, i.e. the description with the highest GDQ value, because this would require a combinatorial search. On the contrary, the system tries to improve the given concept description by expanding a limited number of nodes in the search tree and is guided by a heuristic measure.

One more characteristic of the system should be mentioned: only one operator is applied to the selected (best-quality) node at any one time; therefore, the new node can be selected if its quality is better than the quality of the father node. This is different from standard search procedures, where all the applicable operators are used for the selected node (node expansion). This choice was introduced because the creation of a new node involves the computation of its quality, which, in some cases, can be time-consuming. On the contrary, we try to avoid generating bad quality nodes by selecting the best applicable operator on a heuristic basis, and applying only that operator. The other operators will be used only if the results obtained on this search branch turn out to be unsatisfactory. The heuristics used for selecting the operators will be discussed shortly.

The operators in the search correspond to generalizations or specializations of a given description. In particular, selector truncation is a generalization operator, making the new description cover more positive and negative examples, while complex truncation is a specialization operator, making the set of examples covered by the modified description smaller. Referent modification can be either a specialization or a generalization operator, depending on the type of modification that is being used and on the type of selector involved. The search starts from an initial description, and in this implementation the initial description is supplied by a previous inductive learning phase (using AQ15 or INDUCE). The system can also be applied to the disjunction of all the training events, and the search process can use this as an initial description, although this choice could make the whole system slower. An important issue for further research is related to the integration of the first inductive learning phase and the search in the space of two-tiered concept descriptions.

4.2. Example.

An abstract example of the search process is given in Fig. 3. The nodes contain BCR, ICI, and a graphical representation of the covered examples. The tree is kept in memory throughout the search. The BCR is expressed in disjunctive normal form (it is a "cover"). In this example accuracy is computed according to the formula discussed in sec. 3 and given in Appendix 1, assuming the same typicality of all the examples. The initial description is represented in node 1, and contains two disjuncts

(complexes). The complexes cover the two corresponding rectangular areas in the graphical representation, containing five positive examples out of eight, and one negative example out of five. The ICI extends this coverage by recognizing one more positive example. By eliminating conjunct (selector) s_5 in the second complex we obtain node 3 in the search tree. The accuracy of the description is now improved since all the positive examples are covered. Finally, by truncating the first complex we obtain node 5. It does not cover negative examples any more, and is definitely simpler. This node is then accepted as the improved description resulting from the search. The other nodes lead to inferior concept representations, with respect to GDQ, and are discarded. The quality has been computed with $w_1=w_2=0.5$ (see Appendix 1).




Fig. 1. Example of the search space organized as a tree. For clarity, the cost is omitted, and the simplicity of the ICI is not taken into account. The simplicity of the BCR depends on the number of complexes and selectors.

4.3. The Search Algorithm.

The general search procedure is more precisely presented by the following Search Algorithm:

1. Identify in the search tree the best description D (one with the highest GDQ). Initially, D is the complete and consistent description obtained in stage 1.
2. Apply to D this operator among C_i , S_{ij} , R that potentially improves GDQ of D the best, based on the *Potential Accuracy Improvement* (PAI) heuristics described in sequel:
 - C_i : Remove the i -th complex from D .
 - S_{ij} : Remove the j -th selector from the i -th complex in D .
 - R : Modify the referent in a selector of a complex in D .
3. Compute the GDQ of the node obtained in step 2. If this GDQ is smaller than the GDQ of D , then proceed to step 1.
4. Ask for an explanation of
 - (a) the positive examples that are not covered any more
 - (b) the negative examples that are now covered
 If such an explanation is found, augment ICI accordingly.
5. Update the GDQ value of the new node, by taking into account the added ICI rules.
6. If the *stopping criterion* is satisfied, then STOP, otherwise proceed to step 1.

We shall now discuss the motivation and details of the algorithm, and explain the search strategy.

In step 1, the nodes are chosen on a best first basis, that is the node in the search space with the highest GDQ value is expanded first. This is not always an optimal choice, since apparently "bad" nodes can lead to better descriptions after a number of truncations. Whether the search will behave in this manner will depend on the adequacy of the GDQ as the measure of concept quality.

In step 2, a search operation on the description is chosen heuristically and applied to the description. Only one operator is applied at one time. The heuristic here is to choose the operation which has the best chance of improving the GDQ of the description. The specific choice of the operations C_i and S_{ij} is determined based on the Potential Accuracy Improvement heuristic (PAI). The idea behind the PAI heuristic is to truncate first a complex which covers uniquely a small number of examples. Then, a specific choice of the operator S_{ij} is made so that the truncation of a selector improves the completeness of the description, while the consistency measure has an acceptable value. Finally, when no other operator is recommended, referent modification can be selected. Referent modification can improve both consistency and completeness measures.

In the worst case, this algorithm will perform a full search of the search space. We can observe, however, that the search is controlled by the heuristic PAI, which is computationally much less expensive than the GDQ.

The complex and selector truncation heuristics are implemented together in the following way. Let us first define the PAI. The PAI of truncating a complex is computed as follows:

$$PAI = \#CNI/\#NEG - \#CPI/\#POS$$

where #CNI (#CPI) is the number of negative (positive) examples no longer covered by the concept description after truncating the complex, respectively. #NEG and #POS are the numbers of negative and positive examples, respectively. The PAI of truncating a selector is more complex and is defined as follows:

$$PAI = (\#SPI/\#POS) * P - (\#SNI/\#NEG) * N$$

where #SNI (#SPI) is the number of additional negative (positive) examples covered by the concept after having truncated the selector, respectively. N is the proportion of the negative examples which are not covered by the description and, P is the proportion the positive examples which are not covered by the description.

The operation is chosen based on the value of the PAIs. The operation with the largest PAI is chosen. Finally, the PAIs of selector and complex truncation can be weighted differently. More weight can be assigned to PAI of complex truncation, since the complex truncation simplifies the description more than selector truncation.

In step 3, the system computes the GDQ of the new node. It should be noted that, in the GDQ measure, the typical examples covered directly by the BCR can weigh more than those covered through flexible matching. The examples covered by ICI rules should weigh more than the ones covered through flexible matching but less than the ones covered by the BCR.

In step 4, the "explainer" module (Bergadano et al. 88) is used in order to improve the description even further: the BCR description is extended or shrunk by adding ICI rules. Firstly, complex truncation might have caused some positive examples, that were previously covered, to be lost. In this case some new rules could be introduced in the ICI, that would allow the system to reason about such "special" positive examples, and

understand why they should still be classified as instances of the concept under consideration. On the other hand, selector truncation might have caused some negative examples to be covered, and new rules in the ICI may be added in order to "shrink" the BCR and avoid these erroneous classifications. Another issue, concerning step 3, is whether an explanation should be required at all, since, in some cases, the chosen truncation operator is not an appropriate one, and will lead to a very poor description. In this case it is not even worth to ask for an explanation, and search can continue in other directions. The current strategy is as follows. Suppose the relation $<$ denotes the GDQ ordering among two-tiered descriptions, n is the node we are expanding and m is the node we obtain after the selected truncation. If $m \ll n$, then no explanation is even tried, otherwise the explainer is asked for an explanation and is told how m compares to n with respect to $<$, in order to know how important the request for the explanation is for the search procedure.

In step 5, the GDQ of the obtained two-tiered description is updated after the new ICI rules have been added. Since ICI rules are taken into consideration in the GDQ, new ICI rules will change the GDQ value for a concept.

In step 6, the system decides whether to stop or continue the search. The *stopping criterion* is satisfied when the search space obtained is "very large", or when no qualitative improvement has been obtained for a "long time". The particular values of "very large" and "long time" are parameters of the algorithm. They depend on the size of the initial description. When the system stops, the best node in the search space is produced and becomes the modified two-tiered concept description.

5. Experiments: Learning a two-tiered description of a labor-management contract.

This section describes how the system learns a two-tiered description of a labor-management contract from real data. Unlike the example in Part I, which has been simplified to illustrate the method, rather than its power, the example in this section relies on real data.

The data used in this section comes from *Collective Bargaining* - a review of current collective bargaining issues published by the Government of Canada through its Department of Labor. The data given in *Collective Bargaining* describes labor-management contracts which have been currently negotiated between organizations and those union locals that count at least 500 members. The raw data is divided geographically, as well as by economic sectors. The format of the raw data is pretty

standard. Each contract is described by a number of attributes. Since the attributes vary between economic sectors, we have decided to focus on a single sector: personal and business services. This sector includes unions representing hospital staff, teachers, university professors, social workers, and certain classes of administrative personnel of different organizations. With this kind of data, it was natural to represent concepts using the VL₁ formalism.

Our data describes contracts finalized in the second half of 1987 and first half of 1988. Each contract is described by sixteen attributes, belonging to two main groups: issues related to salaries (e.g. pay increases in each year of contract, cost of living allowance, stand-by pay, etc.), and issues related to fringe benefits (e.g. different kinds of pension contributions, holidays, vacation, dental insurance, etc.).

We have run three experiments. In each experiment we were dealing with two descriptions: a description of a contract, and a description of a contract proposal deemed unacceptable by one of the parties. In each experiment we have looked at the number of events correctly and incorrectly covered by both descriptions, and at the number of events that were not covered by either concept. This was done both on the training set and on a testing set of examples not previously seen by the system. The results of the three experiments are given in Tables 2, 3, 4, and 5 are discussed below. In each experiment, the same training and testing sets were used. The training set consisted of 18 positive and 9 negative examples of contracts; the testing set consisted of 19 positive and 11 negative examples.

Factual Knowledge (27 complexes and 432 selectors)

		Correct	Incorrect	No_Match
<i>Strict Match</i>				
	Training Set	100%	0%	0%
	Testing Set	0%	0%	100%
<i>Flexible Match</i>				
	Training Set	100%	0%	0%
	Testing Set	37%	0%	63%

Table 2 Results of Experiment 1
A concept description as a disjunction of training examples.

Initial Description (11 complexes and 28 selectors)

		Correct	Incorrect	No_Match
<i>Strict Match</i>	Training Set	100%	0%	0%
	Testing Set	80%	17%	3%
<i>Flexible Match</i>	Training Set	100%	0%	0%
	Testing Set	80%	17%	3%

Table 3 Results of Experiment 2.
 Concept descriptions derived by AQ15. Completeness and consistency
 were required.

Optimized Description (9 complexes and 12 selectors)

		Correct	Incorrect	No_Match
<i>Strict Match</i>	Training Set	63%	0%	37%
	Testing Set	43%	3%	54%
<i>Flexible Match</i>	Training Set	85%	0%	15%
	Testing Set	83%	13%	4%
<i>Inferential Match</i>	Training Set	96%	0%	4%
	Testing Set	90%	10%	0%

Table 4 Results of Experiment 3.
 Concept descriptions were obtained using the search procedure described
 in this paper. In the case of inferential matching, expert-provided rules
 were used in combination with flexible matching.

In the first experiment, we have used the events from the training set as a purely factual concept description (Factual Knowledge): the concept was just the disjunction of the training examples. This description is obviously

complete and consistent on the training set but has no predictive power, i.e. it always produces a No_Match answer. This happens because in our experiment, as it often happens when dealing with real data, no testing examples were exactly equal to some training event.

The factual knowledge was also used with a flexible matching technique, based on a measure of the distance between an event and a concept description (see Section 2). Consistency and completeness are achieved in the training set, and this is again obvious since flexible matching is only used when either no match or multiple match has occurred, and this never happens if the factual knowledge is used to classify the training examples. On the contrary, the flexible matching causes some performance improvement for the test set, where 37% of the events are now correctly classified, and we still do not obtain any erroneous classification. This method is similar to statistical techniques in Pattern Recognition, e.g. the nearest neighbor method (Watanabe 85).

In the second experiment, we have used the descriptions learned by AQ15 (Initial Description). Since AQ15 only generates consistent and complete descriptions, classification is 100% correct for the training set, and flexible matching does not affect this performance. For the testing set, the number of correct classifications is still high (80%), and flexible matching does not improve the result in this case. This is partly related to the fact that the descriptions generated by AQ are detailed and specify many alternatives, leaving little space for the no_match case (3%). Moreover, the multiple match case was impossible because AQ15 was run with the "disjoint cover" parameter, causing the generated concept descriptions to have disjoint extensions. In general the flexible matching can improve the performance of the initial description on the test set.

If compared with the factual knowledge, the descriptions generated by AQ are definitely superior. Not only their performance is higher, but they are also much simpler (28 selectors compared to 432). Simplicity is closely related to comprehensibility in the given domain, and allows the performance system to recognize new events more efficiently. The descriptions generated by AQ15 are given in Fig. 2.

```
[duration = 1] & [wage_incr_yr2 = 3.0] & [holidays = 10] v
[wage_incr_yr1 = 2.0% v 2.5% v 2.8% v 3.0% v 4.0% v 4.5%] v
[wage_incr_yr1 = 2.0% v 2.5% v 2.8% v 4.0%] & [wage_incr_yr2 = 2.0% v 4.0%] v
[wage_incr_yr1 = 2.0% v 2.5% v 3.0% v 4.0% v 4.5%] & [holidays = 9] v
[wage_incr_yr1 = 2.0% & [vacation = above_average] ::> acceptable contract
```

```

[wage_incr_yr1 = 2.0% v 2.5% v 4.0%] & [holidays = 10] &
    [vacation = below_average v average] v
[wage_incr_yr1 = 2.0% v 2.5% v 3.0%v 4.0% v 4.5%] & [wage_incr_yr2 = 2.0% v 4.0%] &
    [holidays = 10] & [vacation = below_average v average] v
[duration = 1] & [wage_incr_yr1 = 2.0% v 2.5% v 2.8%v 4.0%] &
    [holidays = 9] & [vacation = below_average v average] v
[wage_incr_yr1 = 2.0% v 2.5% v 4.0%] & [wage_incr_yr2 = 3.0%] &
    [vacation = below_average v average] v
[duration = 1] & [wage_incr_yr1 = 2.0%v 2.5% v 4.0%] &
    [vacation = below_average v average] v
[wage_incr_yr1 = 2.0%] & [wage_incr_yr2 = 3.0%] ::> unacceptable contract

```

Fig. 2. Descriptions Generated by AQ15

The third experiment allows us to evaluate empirically the method presented in this paper. We have used the description generated by the search process (Optimized Description), and evaluated its performance both with the flexible matching alone and with the combination of flexible and deductive matching (Inferential Match). For the sake of completeness we also present the performance of the generated descriptions with strict matching, although this would never be used. Strict matching alone yields restricted coverage and poor performance. In fact, the power of the modified description is due to a combination of all three types of matching (strict, flexible and deductive), and all three contribute to the quality measure of a description as computed during the learning process. This represents a new feature of this system, since inferential matching is usually introduced only after the learning phase is completed (Bergadano and Giordana ,to appear; Michalski et al., 86). Fig 2 gives the descriptions output by the system.

```

[wage_incr_yr2 = 3.0%] v [wage_incr_yr1 = 2.0% v 2.5% v 2.8% v 3.0%v 4.0% v 4.5%] v
[holidays = 9] v [vacation = above_average] ::> acceptable contract

[wage_incr_yr1 = 2.0% v 2.5% v 4.0%] & [holidays = 10] v
[wage_incr_yr2 = 2.0% v 4.0%] & [vacation = below_average v average] v
[holidays = 9] v
[duration = 1] & [wage_incr_yr1 = 2.0% v 2.5% v 4.0%] v
[wage_incr_yr2 = 3.0%] ::> unacceptable contract

```

Fig. 3. Optimized Descriptions.

The modified descriptions are simpler than the ones generated by AQ15,

and should represent the most important characteristics of the labor management concepts. The performance of this descriptions is slightly better if the flexible matching is used (83% correct classifications for the testing set, compared to 80% for the initial description with flexible matching). It should be noted that, on the contrary, the descriptions generated by AQ15 performed better on the training set (100% vs 85%), suggesting that the data might have been overfitted.

This cannot usually be avoided if complete and consistent descriptions have to be obtained, as is required in the implementation of AQ15 used in the experiment.

The training events that were not correctly classified by the description, as it was modified step by step during the search, were analyzed by a domain expert, and the following rules were acquired, as part of the ICI:

```
[wage_incr_yr1 < 3.0%] & [wage_incr_yr2 < 3.0%] ==> low_wages
low_wages & [wage_incr_yr2 < wage_incr_yr1] ::> unacceptable contract
low_wages & [hours ≥ 40]
    & [pension_type = none v retirement_allowance] ::> unacceptable contract
[wage_incr_yr1 > 5.5%] & [vacation = above_average] ::> acceptable contract
```

Fig. 4. Deductive rules of the ICI of the optimized description

These rules allow the system to classify almost all the training events (one of them could not be explained by the expert). The combination of modified description and inferential matching (rules plus flexible matching) produces the best results. The description is still simple, although it now includes the ICI rules, and the number of correct classifications is 90%. Moreover, some of the examples that were previously recognized by flexible matching or strict matching are now also correctly recognized by the ICI rules, and this might suggest that the description is more robust, and could perform even better on a larger test set.

Results of this experiment are summarized in Table 5. They indicate that a two-tiered approach to learning allows a system to learn concepts from a small number of examples, and produce descriptions that are simpler than the ones provided by a typical inductive learning system. In Table 5, the Factual Description denotes again the disjunction of training examples, plus the flexible matching function. The Initial Description is provided by AQ15 and also uses the flexible matching; the missing 3% of the training set represent the "no match" situation. The highest performance is realized by

the Optimized Description, in which the concept meaning is divided into the BCR, generated by the search, a flexible matching procedure, defined a-priori, and a set of deductive rules given by the expert on the basis of misclassified examples selected automatically during the search process. There are four deductive ICI rules with nine conditions (see Fig. 4); their complexity corresponds to four complexes and nine selectors.

	#Complexes	#Selectors	Performance (correct/incorrect)
<i>Factual Description</i>	27	432	37%/0%

<i>Initial Description</i>	11	28	80%/17%

<i>Optimized Description</i>			90%/10%
<i>BCR</i>	9	12	
<i>ICI</i>	4	9	

Table 5. Summary of the experimental results.

6. Conclusion

In this paper, a two-tiered knowledge representation formalism has been developed, and a system that is able to learn two-tiered concept descriptions has been described. Learning is viewed as a state space search guided by a measure of quality that is applied to concept descriptions. The operators in the search modify a given description by removing some of its components or by simplifying the referent of the selectors. The goal of the search process is to obtain simpler but still accurate descriptions. In this way, the comprehensibility and the predictive power of the acquired knowledge are improved.

The system was applied to the problem of learning the concept of an acceptable labor management contract. The experimental results that we obtained confirm the hypothesis that two-tiered descriptions can be more accurate and easier to understand. The ICI used in the experiments included a flexible matching function and a set of logical rules. The performance of the descriptions produced by the search process on the test set is influenced by the use of the inferential matching. This is due to the fact that ICI is used during learning, in order to choose and modify the best descriptions. This property represents an important difference between the presented system and previous approaches, that tend to apply flexible matching only after the learning process is completed.

Some of the motivations behind the system come from previous work (Michalski et al, 1986), that produced some preliminary results in which flexible matching function was applied during the testing phase. The same research investigated the effect of truncating concept descriptions. In the system presented here, though, the flexible matching function is augmented by a set of rules defining how to extend or modify a concept description at the "knowledge level", by describing symbolically its possible transformations. Selector truncation and referent modification are also introduced (besides complex truncation). Truncations are applied automatically during a search process, whereas in (Michalski et al., 1986) they were applied manually, with a limited number of tries. The search for better two-tiered descriptions is the main process during the learning phase implemented in this system.

An important issue for future research and improvements of the implemented system is the integration of the search procedure with the inductive learning system used to generate the initial description (AQ or INDUCE). The first step in this direction is being experimented with: it allows the two systems to share the same heuristics and the same measure of quality. Further progress is related to the possibility of obtaining partially incomplete and inconsistent description also during the generation of the initial description. More experimentation is also needed in order to evaluate more precisely the performance of the implemented system.

Acknowledgements

This research was done in the Artificial Intelligence Center of George Mason University. The activities of the Center are supported in part by the Defence Advanced Research Projects Agency under grant, administered by the Office of Naval Research, N00014-87-k-0874, and in part by the Office of Naval Research under grant No. N00014-88-K-0226. The first author

was supported in part by the Italian Ministry of Education (ASSI) project, and the second author was supported in part by the NSERC grant number A2480. The authors thank Zbig Koperczak for his help with acquiring the contract data, and the MLI group of the Artificial Intelligence center for the comments.

References

Bergadano, F., Matwin, S., Michalski, R.S., Zhang, J., "Learning Two-tiered Descriptions of Imprecise Concepts: A Method Employing Examples of Varied Typicality and A Two-staged Construction of the Base Concept Representation", Part 1: Principles and Methods, Research Report, Machine Learning and Inference Group, Artificial Intelligence Center, George Mason University, 1988 (submitted to Machine Learning Journal).

Bergadano, F. , Giordana, A. , Saitta, L., "Automated Concept Acquisition in Noisy Environments", IEEE Transactions on PAMI, July 1988.

Bergadano, F., Giordana, A., "Pattern Classification: An Approximate Reasoning Framework", International Journal of Intelligent Systems, (To appear).

Michalski, R. S., "A Variable-Valued Logic System as Applied to Picture Description and Recognition", in "Graphic Languages", Nake, F. and Rosenfield, A. (Eds.), North Holland, 1972.

Michalski, R. S., "AQVAL/1--Computer Implementation of a Variable Valued Logic System VL1 and Examples of its Application to Pattern Recognition", Proc. of the 1rst International Joint Conf. on Pattern Recognition, Washington, D.C., pp. 3-17, 1973.

Michalski, R. S., "Pattern Recognition as Rule-guided Inductive Inference", IEEE Transactions on PAMI, vol. 2, no. 2,3,4, pp. 349-361, 1980.

Michalski, R.S., "A Theory and Methodology of Inductive Learning", Chapter in the book "Machine Learning, an Artificial Intelligence Approach", Michalski, R. S., Carbonell, J. G., Mitchell, T. M. (Eds.), Tioga Pub. Co., Palo Alto, Ca, 1983.

Michalski, R. S., Carbonell, J. G., Mitchell, T. M., "Machine Learning: An Artificial Intelligence Approach", Tioga Publishing Co., Palo Alto, Ca, 1983.

chalski, R. S., Mozetic, Hong, J.,I., Lavrac, "The Multi-purpose Incremental Learning System AQ15 and its Testing Application to Three Medical

Domains", Proc. 5th AAAI, pp. 1041-1045 1986.

Michalski, R. S., "Two-Tiered Concept Meaning, Inferential Matching and Conceptual Cohesiveness", Chapter in the Book "Similarity and Analogy", Stella Vosniadou and A. Orton, (Eds), 1987.

Mitchell, T. M., "Version Spaces: An Approach to Concept Learning", Ph. D. dissertation, Stanford University, December 1977.

Robinson J. A. and Sibert E. E., "LOGLISP: an alternative to Prolog" in Machine Intelligence 10, Hayes, J. E. and Michie, D. (Eds.), pp. 399-419, 1982.

Sturt, E., "Computerized Construction in Fortran of a Discriminant Function for Categorical Data", Applied Statistics, vol. 30, pp. 213-222, 1981.

Watanabe, S., "Knowing and Guessing - a Formal and Quantitative Study", Wiley Pub. Co., 1969.

Watanabe, S., "Pattern Recognition Human and Mechanical", John Wiley, 1985.

Weber, S., "A General Concept of Fuzzy Connectives, Negations and Implications based on t-norms and t-conorms", Fuzzy Sets and Systems, vol. 11, pp. 115-134, 1983.

Appendix 1.

A detailed General Description Quality Measure.

The purpose of the Appendix is to define in detail the GDQ measure that was implemented in our experimental system.

First, we have to define the typicality-dependent completeness (TCOM) and typicality-dependent consistency (TCON) of a description:

$$TCOM = \frac{\sum_{e^+ \text{ is } S\text{-covered}} w_S * Typicality(e^+) + \sum_{e^+ \text{ is } F\text{-covered}} w_F * Typicality(e^+) + \sum_{e^+ \text{ is } D\text{-covered}} w_D * Typicality(e^+)}{\sum_{e \in PosQov} Typicality(e)}$$

$$TCON = \frac{\sum_{e^- \text{ is S-covered}} w_s * Typicality(e^-) + \sum_{e^- \text{ is F-covered}} w_f * Typicality(e^-) + \sum_{e^- \text{ is D-covered}} w_d * Typicality(e^-)}{\sum_{e \in NegCov} Typicality(e)}$$

where:

PosCov -- set of positive events covered by two-tiered concept description,

NegCov -- set of negative events covered by two-tiered concept description,

typicality(e) -- typicality of the event e specified by the expert when the event is given.

w_s -- if $typicality(e) \geq t_2$ then 1 else w ,

w_f -- if $t_2 \geq typicality(e) \geq t_1$ then 1 else w ,

w_d -- if $t_2 \geq typicality(e)$ then 1 else w ,

where, t_1 and t_2 are thresholds, and $1 \geq t_2 \geq t_1 \geq 0$, $1 \geq w > 0$.

Now *accuracy* can be defined in terms of TCOM and TCON:

$$Accuracy(description) = w_1 * TCOM(description) + w_2 * TCON(description)$$

where $w_1 + w_2 = 1$.

A measure of comprehensibility of a concept description is difficult to define. We will approximate this measure by a syntactic complexity, defined as:

$$v_1 \sum_{op \in BCR(dsp)} C(op) + v_2 \sum_{op \in ICI(dsp)} C(op)$$

where:

- BCR(dsp) -- a set of all operator occurrences in the BCR

- ICI(dsp) -- a set of all operator occurrences in the ICI

- C(op) -- the complexity of an operator. The complexity of operator on the list <interval, internal disjunction, =, <>, not, &, v, implication, predicate> increases with its position on the list. When an operator is a predicate, C increases with the number of the arguments in the predicate.

- v_1 and v_2 are weights, $v_1 + v_2 = 1$. The BCR should describe the general and easy-to-define meaning of the concept, while the ICI is mainly used to handle nontypical or exceptional events, therefore the BCR should be easier to comprehend than the ICI. v_1 should therefore be larger than v_2 .

The cost consists of two parts:

Measure-Cost -- the cost of measuring the values of variables used in the concept description, it is defined as the function MC

Evaluation-Cost -- the computational cost of evaluating the concept description, it is defined as the function EC.

$$MC(description) = \sum_{e \in Pos + Neg} \sum_{v \in vars(e)} mc(v) / (|Pos| + |Neg|)$$

$$EC(\text{description}) = \sum_{e \in \text{Pos} + \text{Neg}} ec(e) / (|\text{Pos}| + |\text{Neg}|)$$

where

- vars(e) -- set of all occurrence variables used to evaluate the concept description to classify the event e.
- mc(v) -- the cost of measuring the values of the variable v,
- ec(e) -- computational cost of evaluating concept description to classify the event e. This could depend on computation time or on the number of operators involved in the evaluation.

We now define the cost of a description:

$$\text{Cost}(\text{description}) = u_1 * \text{MC}(\text{description}) + u_2 * \text{EC}(\text{description})$$

where u_1 and u_2 are weights.

With the exception of the weights which can be determined experimentally, we have already defined all three components of the quality measure of concept descriptions. In the next section, we will show how the quality measure evaluates a simple concept description. This quality measure has been experimented with two non-trivial examples, acceptable labor-management contract and the concept of "chair", the results are satisfactory. Currently, we are implementing the quality measure in a two-tiered concept learning system and using it to guide the search for a better two-tiered concept description.

Appendix 2

The Preference-based Evaluation Criterion

The problem that we are addressing here is how to combine different criteria used to compare concept descriptions. A possible solution is to have an algebraic formula that, given the numeric evaluations of single criteria, produces a number that represents their combined value. Examples of such approaches are multiplication, weighted sum, maximum/minimum, t-norm/t-conorm (Weber, 1983). Although these approaches are often appropriate, some of them may present disadvantages. Firstly, they usually combine a set of heterogeneous evaluations into a single number, and the meaning of this final number is hard to understand for a human expert. Secondly, they may force the system to evaluate all the criteria, even if it would be sufficient to compare two given descriptions on the basis of the most important one, if one is much better than the other.

In order to overcome some of these problems, we use a *lexicographic evaluation functional* (LEF) (Michalski 1972, Michalski 83) that combines the above mentioned criteria. The general description quality measure is thus defined as:

$$\text{GDQ}(\text{description}) = \langle (\text{Accuracy}, \tau_1), (\text{Comprehensibility}, \tau_2), (\text{Cost}, \tau_3) \rangle$$

where τ_1 , τ_2 , and τ_3 are tolerance thresholds (which will be discussed later).

In this evaluation scheme, the criteria are ordered according to their importance, and a tolerance threshold is associated with each criterion. If the difference of the evaluation of two expressions under a given criterion is less than the corresponding

tolerance, the two descriptions are considered equivalent with respect to that criterion. The most important measure in the LEF is evaluated first, and the subsequent measure is evaluated only if the previous one is a tie.

The LEF evaluation scheme is not affected by the main problems related to algebraic formulas, which we have discussed above, but it may be useful to extend it in some cases. In fact, it can be difficult to determine the tolerance. If the tolerance is too small, we have very little chance of using the other criteria. If the tolerance is too large, some important criterion might be underestimated. Suppose, for example that two descriptions d_1 and d_2 are such that $\text{Accuracy}(d_1) \gg \text{Accuracy}(d_2)$ but the difference is within the tolerance, and suppose that $\text{Comprehensibility}(d_2)$ is slightly better than $\text{Comprehensibility}(d_1)$ and not in the tolerance. In this case d_2 would be preferred although d_1 is probably better, since its accuracy is much greater and the comprehensibility of the two descriptions is approximately the same. In order to avoid this problem, the LEF measure can be extended in the following way: first LEF is applied with larger tolerances, in such a way that all the relevant criteria are taken into account; then, if the comparison still results in a tie, a Weighed Evaluation Functional (WEF) is used to combine the measures (i.e. the description having the maximum weighted sum of the measures is preferred).

The above criteria can also be applied to two-tiered descriptions. The accuracy of the acquired knowledge does not only depend on the explicit information, but also on the implicit reasoning abilities. Inferential Concept Interpretation also affects cost, since it allows the performance system to use a simpler BCR, and reason about special details only in exceptional cases. Finally, the comprehensibility of a two-tiered representation must be carefully evaluated, since one of its implied goals is to state a clear and simple concept description in the BCR and to account for meaningful special cases through a reasoning process.

REDUNDANT VISION FOR ROBOTICS

Harry Wechsler
Department of Computer Science
George Mason University
Fairfax, VA 22030

Lee Zimmerman
Department of Electrical Engineering
University of Minnesota
Minneapolis, MN 55455

Abstract: We suggest the distributed associative memory (DAM) model for distributed and fault-tolerant computation as it relates to object recognition tasks. The fault-tolerance is with respect to geometrical distortions (scale and rotation), noisy inputs, occlusion/overlap, and memory faults. We have developed an experimental system for fault-tolerant structure recognition which shows the feasibility of such an approach. We further extended our approach to the problem of multisensory data integration and applied it successfully to the recognition of colored polyhedral objects.

2

1) Introduction

The challenge of the visual recognition problem stems from the fact that the projection of an object onto an image can be confounded by several dimensions of variability such as uncertain perspective, changing orientation and scale, sensor noise and occlusion, and non-uniform illumination. A vision system must not only be able to sense the identity of an object despite this variability, but must also be able to characterize such variability- because the variability inherently carries much of the valuable information about the world.

Our goal is to derive the functional characteristics of image representations suitable for invariant recognition using distributed processing methods. One has to seek appropriate transformations such that interactions between the internal structure of the resulting representations yield invariant recognition. As Simon (1982) points out, all mathematical derivation can be viewed simple as a change of representation, making evident what was previously true but obscure. Solving a problem then means transforming it so as to make the solution transparent.

The recognition process should match a derived input representation of the short term memory (STM) against long term memory (LTM). We suggest the LTM organization in terms of distributed associative memory (DAM). The DAM is related to generalized match filters (GMF) (Caulfield and Weinberg, 1982) and synthetic discriminant functions (SDF) (Hester and Casasent, 1981). Like them, DAMs attempt to capture a distributed representation which averages over the variations belonging to the same class. DAMs allow for the implicit representation of structural relationships and contextual information, helpful constraints for choosing among different interpretations. Finally, because information is distributed in the memory, the overall function of the memory becomes resistant to noise, faults in memory, and degraded stimulus key vectors.

2) System Outline

This section describes our invariant object recognition system in detail. We begin by examining the preprocessing system which produces the vectors associated by the distributed associative memory. Then the distributed associative memory and our method of fusing multiple sources of information are described and analyzed. A block diagram of the system is shown in Figure 1.

2.1) Invariant Representation

The goal of this subsection is to examine the various components used to produce the vectors which are associated in the distributed associative memory. The block diagram which describes the various functional units involved in obtaining an invariant image representation is shown in Figure 1. The image is complex-log conformally mapped so that rotation and scale changes become translation in the transform domain. Along with the conformal mapping, the image is also filtered by a space variant filter to reduce the effects of aliasing. The conformally mapped image is then processed through a Laplacian in order to solve some problems associated with the conformal mapping (See 2.1.3). The Fourier transform of both the conformally mapped image and the Laplacian processed image produce the four output vectors. The magnitude output vector $|e|_1$ is invariant to linear transformations of the object in the input image. The phase output vector ϕ_2 contains information concerning the spatial properties of the object in the input image.

2.1.1) Complex-Log Mapping and Space Variant Filtering

The first box of the block diagram given in Figure 1 consists of two components: complex-log mapping and space variant filtering. Complex-log mapping transforms an image from rectangular coordinates to polar exponential coordinates. This transformation changes rotation and scale into translation. If the image is mapped onto a complex plane then each pixel (x,y) on the Cartesian plane can be described mathematically by $z = x + jy$. The complex-log mapped points w are described by

$$w = \ln(z) = \ln(|z|) + j\theta_z \quad (1)$$

where $|z| = (x^2 + y^2)^{1/2}$ and $\theta_z = \tan^{-1}(y/x)$.

Our system sampled 256x256 pixel images to construct 64x64 complex-log mapped images. Samples were taken along radial lines spaced 5.6 degrees apart. Along each radial line the step size between samples increased by powers of 1.08. These numbers are derived from the number of pixels in the original image and the number of samples in the complex-log mapped image.

An excellent examination of the different conditions involved in selecting the appropriate number of samples for a complex-log mapped image is given in (Massone et. al., 1985). The non-linear sampling can be split into two distinct parts along each radial line. Toward the center of the image the samples are dense enough that no anti-aliasing filter is needed. Samples taken at the edge of the image are large and an anti-aliasing filter is necessary. The image filtered in this manner has a circular region around the center which corresponds to an area of highest resolution. The size of this region is a function of a number of angular samples and radial samples. The filtering is done, at the same time as the sampling, by convolving truncated Bessel functions with the image in the space domain. The width of the Bessel functions main lobe is inversely proportional to the eccentricity of the sample point.

A problem associated with the complex-log mapping is sensitivity to center misalignment of the sampled image. Small shifts from the center causes distortions in the complex-log mapped image. Our system assumes that the object is centered in the image frame. Slight misalignments are considered noise. Large misalignments are considered as translations and could be accounted for by changing the gaze in such a way as to bring the object into the center of the frame (see 3.2). The decision about what to bring into the center of the frame is an active function and should be determined by the task. An example of a system which could be used to guide the translation process was developed by Anderson et al (1985). Their pyramid system analyzes the input image at different temporal and spatial resolution levels. Their smart sensor was then able to shift its fixation such that interesting parts of the image (i.e. something large and moving) was brought into the central part of the frame for recognition.

2.1.2) Fourier Transform

The second box in the block diagram of Figure 1 is the Fourier transform. The Fourier transform of a 2-dimensional image $f(x,y)$ is given by

$$F(u,v) = \iint_{-\infty}^{\infty} f(x,y) \exp[-j(ux+vy)] dx dy \quad (2)$$

and can be described by two 2-dimensional functions corresponding to the magnitude $|F(u,v)|$ and phase $\angle F(u,v)$. The magnitude component of the Fourier transform which is invariant to translation, carries much of the contrast information of the image. The phase component of the Fourier transform carries information about how things are placed in an image. Translation of $f(x,y)$ corresponds to the addition of a linear phase component. The complex-log mapping transforms rotation and scale into translation and the magnitude of the Fourier transform is invariant to those translations so that $|F(u,v)|$ will not change significantly with rotation and scale of the object in the image.

2.1.3) Laplacian

The Laplacian that we use is a difference-of-Gaussians (DOG) approximation to the $\nabla^2 G$ function as given by Marr (1982).

$$\nabla^2 G = \frac{1}{\pi\sigma^4} [1 - r^2/2\sigma^2] \exp(-r^2/2\sigma^2) \quad (3)$$

The result of convolving the Laplacian with an image can be viewed as a two step process. The image is blurred by a Gaussian kernel of a specified width σ . Then the isotropic second derivative of the blurred image is computed. The width of the Gaussian kernel is chosen such that the conformally mapped image is visible -- approximately 2 pixels in our experiments. The Laplacian sharpens the edges of the object in the image and sets any region that did not change much to zero. Below we describe the benefits from using the Laplacian.

The Laplacian eliminates the stretching problem encountered by the complex-log mapping due to changes in object size. When an object is expanded the complex-log mapped image will translate. The pixels vacated by this translation will be filled with more pixels sampled from the center of the scaled object. These new pixels will not be significantly different than the displaced pixels so the result looks like a stretching in the complex-log mapped image. The Laplacian of the complex-log mapped image will set the new pixels to zero because they do not significantly change from their surrounding pixels. The Laplacian eliminates high frequency spreading due to the finite structure of the discrete Fourier transform and enhances the differences between memorized objects by accentuating edges and de-emphasizing areas of little change.

2.2) Distributed Associative Memory (DAM)

The particular form of distributed associative memory that we deal with in this paper is a memory matrix which modifies the flow of information. Stimulus vectors are associated with response vectors and the result of this association is spread over the entire memory space. Distributing in this manner means that information about a small portion of the association can be found in a large area of the memory. New associations are placed over the older ones and are allowed to interact. This means that the size of the memory matrix stays the same regardless of the number of associations that have been memorized.

The above discussion illuminates several properties of distributed associative memories which are different from the more traditional ones about memory. Because the associations are allowed to interact with each other an implicit representation of structural relationships and contextual information can develop, and as a consequence a very rich level of interactions can be captured. Since there are few restrictions on what vectors can be associated there can exist extensive indexing and cross-referencing in the memory. Since the information is

distributed, the overall function of the system is resistant to faults in the memory and degraded stimulus vectors. Distributed associative memory captures a distributed representation which is context dependent. This is quite different from the simplistic behavioral model (Hebb, 1949).

2.2.1) Construction and Recall

The construction stage assumes that there are n pairs of m -dimensional vectors that are to be associated by the distributed associative memory. This can be written as

$$M\vec{s}_i = \vec{r}_i \quad \text{for } i = 1, \dots, n \quad (4)$$

where $\vec{s}_i$ denotes the i th stimulus vector and $\vec{r}_i$ denotes the i th corresponding response vector. We want to construct a memory matrix M such that when the k th stimulus vector $\vec{s}_k$ is projected onto the space defined by M the resulting projection will be the corresponding response vector $\vec{r}_k$. Specifically we want to solve the following equation:

$$MS = R \quad (5)$$

where $S = [\vec{s}_1 | \vec{s}_2 | \dots | \vec{s}_n]$ and $R = [\vec{r}_1 | \vec{r}_2 | \dots | \vec{r}_n]$. A unique solution for this equation does not necessarily exist for any arbitrary group of associations that might be chosen. Usually, the number of associations n is smaller than m , the length of the vector to be associated, so the system of equations is underconstrained. The constraint used to solve for a unique matrix M is that of minimizing the square error, $\|MS - R\|^2$, which results in the solution

$$M = RS^+ \quad (6)$$

where S^+ is known as the Moore-Penrose generalized inverse of S (Kohonen, 1984)

The recall operation projects an unknown stimulus vector $\vec{s}$ onto the memory space M . The resulting projection yields the response vector $\vec{r}$.

$$\vec{r} = M\vec{s} \quad (7)$$

If the memorized stimulus vectors are independent and the unknown stimulus vector $\vec{s}$ is one of the memorized vectors $\vec{s}_k$, then the recalled vector will be the associated response vector $\vec{r}_k$. If the memorized stimulus vectors are dependent, then the vector recalled by one of the memorized stimulus vectors will contain the associated response vector and some crosstalk from the other stored response vectors. The resulting noise or crosstalk in the output is due to the similarity of the memorized vectors.

The recall can be viewed as the weighted sum of the response vectors. The recall begins by assigning weights according to how well the unknown stimulus vector matches with the memorized stimulus vector using a linear least squares classifier. The response vectors are multiplied by the weights and summed together to build the recalled response vector. The

recalled response vector is usually dominated by the memorized response vector that is closest to the unknown stimulus vector. The distributed associative memory will have interactions between the different associations and this allows some generalization of responses to previously unknown stimulus.

3 Fault tolerance is a byproduct of the distributed nature and error correcting capabilities of the distributed associative memory. By distributing the information, no single memory cell carries a significant portion of the information critical to the overall performance of the memory. Assume that there are n associations in the memory and each of the associated stimulus and response vectors have m elements. This means that the memory matrix has m^2 elements. Also assume that the noise that is added to each element of a memorized stimulus vector is independent, zero mean, with a variance of σ_i^2 . The recall from the memory is then

$$\vec{r} = \vec{r}_k + \vec{y}_o = M(\vec{s}_k + \vec{y}_i) = \vec{r}_k + M\vec{y}_i \quad (8)$$

9 where $\vec{y}_i$ is the input noise vector and $\vec{y}_o$ is the output noise vector. The ratio of the average output noise variance to the average input noise variance is

$$\frac{\sigma_o^2}{\sigma_i^2} = \frac{1}{m} \text{Tr}[MM^T] \quad (9)$$

10 For the autoassociative case this simplifies to

$$\frac{\sigma_o^2}{\sigma_i^2} = \frac{n}{m} \quad (10)$$

12 This says that when a noisy version of a memorized input vector is applied to the memory the recall is improved by a factor corresponding to the ratio of the number of memorized vectors to the number of elements in the vectors. For the heteroassociative memory matrix a similar formula holds as long as n is less than m (Stiles and Denz, 1985).

$$\frac{\sigma_o^2}{\sigma_i^2} = \frac{1}{m} \text{Tr}[RR^T] \text{Tr}[(S^T S)^{-1}] \quad (11)$$

13 Another way of viewing this error correcting process is to notice that the memory matrix is the orthogonal projection matrix for the set of stimulus vectors. The noise vector in this m -dimensional space will be projected onto the space spanned by the n memorized vectors. The parts of the noise vector that are orthogonal to the n memorized stimulus vectors will be lost and this accounts for the noise reduction in the output recall vector.

2.3) Data Fusion

We suggest next using distributed associative memory to integrate visual information from multiple cues or sources for performing image interpretation tasks. A source amounts to a 2-dimensional function extracted from an image which carries some distinct information about the 3-dimensional scene. In our particular case the sources of information can take on two forms, direct or derived. A direct source is one where the information is directly in registration with the iconic image. Examples would be infra-red images, range images, spectral band images, etc.. A

derived source is one which requires some type of preprocessing before the spatial information is made explicit. Preprocessing examples are stereopsis and optical flow derivation. Integration means combining information from several distinct sources to produce a response which is possibly quite different from the input information. An example of integration would be the ability of human vision to extract reliable viewer-centered depth information from the binocular and motion information contained in a set of images. In this case the 'sensors', stereopsis and optical flow, are distinct but both carry geometric information about the scene. In order for information from distinct sources to be integrated at some point they must speak the same language.

Looking for a general method for combining multiple distinct sources into a unified interpretation is a problem which has been examined from many angles. Statistical approaches such as Garvey and Lowrance's (1983) threat assessment for battle management which uses Dempster/Shaffer model for source integration, need much a-priori information about the statistical variations between the different source measurements and are computationally very expensive. Hybrid approaches such as Waxman and Duncan's (1986) stereomotion, use extensive knowledge about how the sources vary with each other in order to develop their algorithm, and as a result are not general or easily extensible to the addition of other sources. Both of these methods need higher level processing to handle conflicts in source interpretation.

Our approach attempts to overcome the above problems. We assume that the output of the sources vary consistently for a given input.

Let $S = [\vec{s}_1 | \vec{s}_2 | \dots | \vec{s}_n]$ and $R = [\vec{r}_1 | \vec{r}_2 | \dots | \vec{r}_n]$ be the stimulus and response matrices respectively. Each stimulus vector is made up of

elements coming from the different sources $\vec{s}_i = [\vec{s}_{i1} | \vec{s}_{i2} | \dots | \vec{s}_{ij}]$ where $\vec{s}_{ij}$

consists of elements from the j^{th} source. The memory matrix is constructed in the same fashion as before, $M = RS^+$.

Integration of sources in this manner has several positive properties. First, it is quite general and as such it is easily extensible to multiple sources. If new sources need to be added they simply augment the previous stimulus vector structure. Second, prioritizing the influence that a single source can have on the output can be done by adjusting the size and/or the quantization of elements that source donates to the stimulus vectors.

3) Experiments

In this section we discuss the result of computer simulations of our system. Images of objects are first preprocessed through the subsystem outlined in subsection 2.2. The output of such a subsystem is four vectors: $|o|_1$, ϕ_1 , $|o|_2$, and ϕ_2 . We construct the memory by associating the stimulus vector $|o|_1$ with the response vector ϕ_2 for each object in the database. To perform a recall from the memory the unknown image is preprocessed by the same subsystem to produce the vectors $|\tilde{o}|_1$, $\tilde{\phi}_1$, $|\tilde{o}|_2$, and $\tilde{\phi}_2$. The resulting stimulus vector $|\tilde{o}|_1$ is projected onto the memory matrix to produce a response vector which is an estimate of the

memorized phase $\hat{\phi}_2$. The estimated phase vector $\hat{\phi}_2$ and the magnitude $|\tilde{\phi}_1|$ are used to reconstruct the memorized object. The difference between the estimated phase $\hat{\phi}_2$ and the phase $\tilde{\phi}_2$ is used to estimate the amount of rotation and scale experienced by the object.

3.1) Invariant Recognition

The database of images consists of twelve objects: four keys, four mechanical parts, and four leaves. The objects were chosen for their essentially two-dimensional structure. Each object was photographed using a digitizing video camera against a black background. We emphasize that all of the images used in creating and testing the recognition system were taken at different times using various camera rotations and distances. The images are digitized to 256x256, eight bit quantized pixels, and each object covers an area of about 40x40 pixels. This small object size relative to the background is necessary due to the non-linear sampling of the complex-log mapping. The objects were centered within the frame by hand. This is the source of much of the noise and could have been done automatically using the object's center of mass or some other criteria determined by the task. The orientation of each memorized object was arbitrarily chosen such that their major axis was vertical. The 2-dimensional images that are the output from the invariant representation subsystem are scanned horizontally to form the vectors for memorization.

The first example of the operation of our system is shown in Figure 2. Figure 2a) is the image of one of the key as it was memorized. Figure 2b) is the unknown object presented to our system. The unknown object in this case is the same key that has been scaled. Figure 2c) is the recalled, reconstructed image. The rounded edges of the recalled image are artifacts of the complex-log mapping. Notice that the reconstructed recall is the unrotated memorized key with some noise caused by errors in the recalled phase. Figure 2d) is a histogram which graphically displays the classification vector which corresponds to S^3 . The histogram shows the interplay between the memorized images and unknown image. The "2" on the bargraph indicates which of the twelve classes the unknown object belongs. The histogram gives a value which is the best linear estimate of the image relative to the memorized objects. Another measure, the signal-to-noise ratio (SNR), is given at the bottom of the recalled image. SNR compares the variance of the ideal recall after processing with the variance of the difference between the ideal and actual recall. This is a measure of the amount of noise in the recall. The SNR does not carry much information about the quality of the recall image because the noise measured by the SNR is due to many factors such as misalignment of the center, changing reflections, and dependence between other memorized objects -- each affecting quality in a variety of ways. Rotation and scale estimates are made using vector $\tilde{D}$ corresponding to the difference between the unknown vector $\tilde{\phi}_2$ and the recalled vector $\hat{\phi}_2$. In an ideal situation $\tilde{D}$ will be a plane whose gradient indicates the exact amount of rotation and scale the recalled object has experienced. In our system the recalled vector $\hat{\phi}_2$ is corrupted with noise which means rotation and scale have to be estimated. The estimate is made by letting the first order difference $\tilde{D}$ at each point in the plane vote for a specified range of rotation or scale.

Figure 3 is the result of randomly setting the elements of the memory

5

9.

3.

objects being viewed the response is correct. For viewpoints between two objects the response was dominated by one of the objects present in the input. The drawing 4e) shows the placement of the sampled center points.

4) Conclusion

Our experiments demonstrate the feasibility of our system for 2-dimensional invariant recognition and data fusion. Information, implicit in the input patterns, is made explicit through the preprocessing subsystem. This is followed by the distributed associative memory which stores and retrieves information. The DAM, because of its distributed nature, lends robustness to the overall system function in the presence of noise, occlusion, and other factors which degrade performance. There are some properties of the DAM which are crucial for data integration. In the DAM, the parameters necessary for combining multiple sources are extracted from the environment and recall is context sensitive. We feel an important aspect of our system is the natural way it processes iconic information which is in contrast to AI symbolic approaches.

The DAM, like most other NN models, does not exploit the topographical power naturally present in visual information. Examinations of the visual context have shown that visual information is processed through multiple mappings of the visual fields. Incorporation of this type of information into a NN model could lead to dramatic compression of the necessary number of connections and more processing power for visual tasks. All pattern recognition techniques, including NN, are sensitive to the nature of the training set. If the training set is not representative of the classes which will be encountered in the environment then the abilities of the system to classify and generalize will be hampered. Research in the area of NN is in its infancy. We expect the future to bring better understanding of these characteristics along with new methods of learning.

The computer vision system presented in this paper was designed with only 2-dimensional metric distortions in mind. We have shown some ability to deal with clustered 3-dimensional polyhedra and are presently extending our work toward 3-dimensional object recognition. We propose to use an approach based on characteristic views (Chakravarty and Freeman, 1982) or aspects (Koenderink and Van Doorn, 1979) which suggests that the infinite 2-dimensional projections of a 3-dimensional object can be grouped into a finite number of topological equivalence classes. An efficient 3-dimensional recognition system would require a parallel indexing method to search for object models in the presence of geometric distortions, noise, and occlusion. Our object recognition system using distributed associative memory can fulfill those requirements with respect to characteristic views. The strength of biological vision and the weaknesses in computational vision when it comes to analyzing, sensing, and integrating information from the environment, suggest that certain aspects need to be incorporated into future systems. First, the interpretation of the environment should be done in 3-dimensions because the world is 3-dimensional. Second, maintaining spatial integrity of the scene and its projection is extremely important for recognition and integration. Third, the integration of information should take place in an active manner. Finally, and most importantly, methods of integration need to include methods for verifying assumptions and learning from the environment.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16

1

2


References

1. Anderson, C. H., P. J. Burt, and G. S. Van Der Wal: Change detection and tracking using pyramid transform techniques, Proc. of the SPIE Conf. on Intelligent Robots, and Computer Vision, Vol. 379, pp. 72-78, 1985.
2. Caulfield, H. J. and M. H. Weinberg: Computer recognition of 2-D pattern using generalized matched filters, Applied Optics, Vol. 21, no. 9, 1982.
3. Chakravarty, I., and H. Freeman: Characteristic views as a basis for 3-D object recognition, Proc. SPIE of Robot Vision, Vol. 336, pp. 37-45, 1982.
4. Faugeras, O. D., and M. Hebert: The representation, recognition, and positioning of 3-D shapes from range data, Techniques for 3-D Machine Perception, A. Rosenfeld (Ed.), North-Holland, pp. 13-52, 1986.
5. Hartline, P. L. Kass, and M. Loop: Merging of modalities in the optic tectum: infrared and visual integration in the rattlesnake, Science, Vol. 199, pp. 543-548, 1978.
6. Hebb, D. O.: The Organization of Behavior, New York: Wiley.
7. Hester, C., and D. Casasent: Interclass discrimination using synthetic discriminant functions (SDF), Proc. SPIE on Infrared Technology for Detection and Classification, Vol. 302, 1981.
8. Koenderink, J. J., and A. J. Van Doorn: Internal representation of solid shape with respect to vision, Biological Cybernetics, Vol. 32, no. 4, pp. 211-216, 1979.
9. Kohonen, T.: Self-Organization and Associative-Memories, Springer-Verlag, 1984.
10. Massone, L., G. Sandini, and V. Tagliasco: Form-invariant topological mapping strategy for 2D shape recognition, Computer Vision Graphics and Image Processing, Vol. 30, 169-188, 1985.
11. McClelland, J. L., D. E. Rumelhart, and the PDP Research Group (Eds.): Parallel Distributed Processing, Vol. 1, 2, MIT Press, 1986.
12. Simon, H. A.: The Sciences of the Artificial (2nd ed.), MIT Press, 1982.
13. Stiles, G. S. and D. L. Denq: On the effect of noise on the Moore-Penrose generalized inverse associative memory, IEEE Trans. on PAMI, Vol. 7, no. 3, pp. 358-360, 1985.
14. Waxman, A. M., and J. H. Duncan: Binocular image flows: steps toward stereo motion fusion, IEEE Trans. on PAMI, Vol 8, no. 6, 715-729, 1986.

15. Wechsler, H. and G. L. Zimmerman: 2-D invariant object recognition using distributed associative memory, IEEE Trans. on PAMI (to appear), 1989.


Acknowledgement: This work was supported in part by National Science Foundation under grant NSF-EET8713563.

Figure 1: Block Diagram of the System.






Original



Unknown




Estimated Rotation: 90°
SNR = -3.37 Db



Memory:6

Figure 2. Recall Using Rotated and Scaled Key



a) Ideal recall



b) 30% Removed



c) 50% Removed



d) 75% Removed

Figure 3. Recall for Memory Matrix Randomly Set to Zero

20b - 20 sided medium blue; rotation approx. 30 right
 20g - 20 sided gray; rotation 15 left
 12r - 12 sided red; smaller with rotation; very dim
 12y - 12 sided yellow; smaller
 10b - 10 sided light brown; smaller rotated 20
 10o - 10 sided orange; rotation 40
 8b - 8 sided light blue; rotation 30
 8bl - 8 sided black; smaller; extremely dim
 8o - 8 sided orange; larger
 6b - 6 sided dark blue; larger
 6g - 6 sided green; rotation 45
 6y - 6 sided yellow; larger
 4b - 4 sided dark blue; rotation 45
 4g - 4 sided light blue; smaller rotation 10
 4p - 4 sided lavender; rotation 30

Table 1. - Description of 15 Classes of Colored Polyhedral
 Objects and Their Corresponding Transformations.

10b = scaled and rotated
 20b = -0.037122
 20g = 0.112992
 12r = 0.074970
 12y = -0.002770
 10b = 0.589933
 10o = 0.389303
 8b = -0.008726
 8o = -0.038870
 8b1 = -0.260333
 6b = -0.039565
 6g = -0.011397
 6y = -0.034024
 4b = 0.010866
 4g = 0.018852
 4p = 0.083681
 maximum: 10b = 0.589933
 minimum: 8b1 = 0.260333

20b = rotated
 20b = 0.669728
 20g = 0.085861
 12r = -0.063536
 12y = -0.021761
 10b = 0.027957
 10o = 0.051344
 8b = 0.038651
 8o = -0.027160
 8b1 = -0.206991
 6b = 0.189514
 6g = 0.076489
 6y = -0.013071
 4b = 0.001625
 4g = -0.000555
 4p = 0.080724
 maximum: 20b = 0.669728
 minimum: 8b1 = -0.206991

4g - scaled and rotated
 20b = 0.099462
 20g = -0.005444
 12r = 0.101725
 12y = -0.092328
 10b = 0.102195
 10o = -0.020118
 8b = -0.020353
 8o = 0.156823
 8b1 = 0.239798
 6b = 0.177184
 6g = 0.055344
 6y = 0.007695
 4b = 0.112544
 4g = 0.451811
 4p = 0.023719
 maximum: 4g = 0.451811
 minimum: 12y = -0.092328

6b - scaled
 20b = 0.121183
 20g = 0.018289
 12r = 0.078140
 12y = -0.041012
 10b = 0.082051
 10o = -0.084216
 8b = 0.009179
 8o = -0.085697
 8b1 = 0.274767
 6b = 0.330043
 6g = 0.152556
 6y = 0.040148
 4b = 0.006260
 4g = 0.005951
 4p = 0.091934
 maximum: 6b = 0.330043
 minimum: 8o = -0.085697

Table 2. Responses of the Fusion Memory

8b - rotated
 20b = 0.149574
 20g = 0.226339
 12r = 0.091095
 12y = 0.033431
 10b = 0.072598
 10o = -0.084049
 8b = 0.553082
 8o = 0.277205
 8b1 = 0.332871
 6b = -0.155819
 6g = -0.043836
 6y = -0.031867
 4b = -0.051010
 4g = 0.070805
 4p = -0.201123
 maximum: 8b = 0.553082
 minimum: 4p = -0.201123

10o - rotated
 20b = -0.107486
 20g = 0.016014
 12r = -0.019775
 12y = 0.026568
 10b = 0.213287
 10o = 0.519179
 8b = 0.014449
 8o = -0.049224
 8b1 = -0.481672
 6b = 0.266687
 6g = 0.055487
 6y = -0.006067
 4b = 0.075286
 4g = 0.061362
 4p = 0.208430
 maximum: 10o = 0.519179
 minimum: 8b1 = -0.481672

12y - scaled
 20b = -0.016374
 20g = -0.030442
 12r = 0.007098
 12y = 0.684478
 10b = 0.068539
 10o = 0.002150
 8b = 0.072255
 8o = -0.128156
 8b1 = -0.156196
 6b = 0.258971
 6g = -0.085213
 6y = 0.036109
 4b = 0.044082
 4g = 0.116001
 4p = -0.021450
 maximum: 12y = 0.684478
 minimum: 8b1 = -0.156196

14b - rotated
 20b = 0.033306
 20g = -0.043632
 12r = 0.102113
 12y = -0.053745
 10b = -0.021697
 10o = 0.045749
 8b = 0.077244
 8o = 0.127757
 8b1 = 0.302109
 6b = 0.038464
 6g = -0.060469
 6y = 0.057145
 4b = 0.347842
 4g = 0.089829
 4p = 0.012316
 maximum: 4b = 0.347842
 minimum: 6g = -0.060469

16
 15
 14
 13
 12
 11
 10
 9
 8
 7
 6
 5
 4
 3
 2
 1
 0
 6g - rotated
 20b = -0.049967
 20g = 0.068956
 12r = -0.166670
 12y = 0.003740
 10b = 0.118696
 10o = 0.046837
 8b = -0.007574
 8o = -0.077262
 8b1 = 0.009852
 6b = 0.012869
 6g = 0.461146
 6y = 0.210418
 4b = 0.080270
 4g = 0.052528
 4p = -0.078049
 maximum: 6g = 0.461146
 minimum: 12r = -0.166670

8b1 - scaled
 20b = -0.009174
 20g = -0.018269
 12r = 0.046817
 12y = 0.016891
 10b = -0.014401
 10o = 0.012176
 8b = -0.000099
 8o = -0.006983
 8b1 = 0.445833
 6b = 0.099586
 6g = 0.017941
 6y = -0.013268
 4b = 0.025530
 4g = 0.020772
 4p = 0.087969
 maximum: 8b1 = 0.445833
 minimum: 20g = -0.018269

12r - scaled and rotated
 20b = -0.101214
 20g = 0.061621
 12r = 0.333425
 12y = -0.011372
 10b = 0.090512
 10o = 0.107102
 8b = 0.042359
 8o = -0.090681
 8b1 = 0.138120
 6b = 0.119825
 6g = -0.071757
 6y = 0.077901
 4b = -0.016150
 4g = 0.032028
 4p = 0.053942
 maximum: 12r = 0.333425
 minimum: 20b = -0.101214

20g - rotated
 20b = 0.074942
 20g = 0.789082
 12r = 0.209688
 12y = -0.032170
 10b = 0.009353
 10o = 0.001790
 8b = 0.097038
 8o = 0.113085
 8b1 = 0.083690
 6b = 0.187161
 6g = -0.066753
 6y = 0.073837
 4b = -0.145125
 4g = -0.046533
 4p = -0.283062
 maximum: 20g = 0.789082
 minimum: 4p = -0.283062

Table 2. Responses of the Fusion Memory

4p - rotated
 20b = 0.135660
 20g = -0.034477
 12r = -0.036622
 12y = -0.020434
 10b = 0.015881
 10o = 0.087427
 8b = -0.062108
 8o = 0.044841
 8b1 = 0.134441
 6b = 0.063891
 6g = 0.063415
 6y = -0.022305
 4b = 0.059183
 4g = 0.042149
 4p = 0.329339
 maximum: 4p = 0.329339
 minimum: 8b = -0.062108

6y - scaled
 20b = 0.026048
 20g = -0.044931
 12r = -0.226118
 12y = -0.033081
 10b = 0.107014
 10o = 0.043923
 8b = -0.078311
 8o = 0.035680
 8b1 = -0.225961
 6b = 0.066302
 6g = 0.645227
 6y = 0.683283
 4b = 0.088817
 4g = 0.024370
 4p = 0.032338
 maximum: 6y = 0.683283
 minimum: 12r = -0.226118

8o - scaled
 20b = 0.067367
 20g = 0.056499
 12r = 0.108363
 12y = -0.003583
 10b = -0.117320
 10o = 0.009010
 8b = -0.051742
 8o = 0.986751
 8b1 = -0.184686
 6b = -0.202276
 6g = 0.031471
 6y = 0.017449
 4b = 0.034253
 4g = 0.124651
 4p = 0.371463
 maximum: 8o = 0.986751
 minimum: 6b = -0.202276

20b = -0.016374
 20g = -0.030442
 12r = 0.007098
 12y = 0.684478
 10b = 0.068539
 10o = 0.002150
 8b = 0.072253
 8o = -0.128156
 8bl = -0.156196
 6b = 0.258971
 6g = -0.085213
 6y = 0.036109
 4b = 0.044082
 4g = 0.116001
 4p = -0.021450
 maximum: 12y = 0.684478
 minimum: 8bl = -0.156196

without occlusion

20b = -0.061517
 20g = 0.055189
 12r = 0.036854
 12y = 0.638085
 10b = 0.084086
 10o = 0.059582
 8b = 0.086944
 8o = -0.140270
 8bl = -0.566859
 6b = 0.114629
 6g = -0.098479
 6y = 0.043607
 4b = 0.118768
 4g = 0.107182
 4p = 0.100319
 maximum: 12y = 0.638085
 minimum: 8bl = -0.566859

with occlusion

6g - rotated

20b = -0.049967
 20g = 0.068956
 12r = -0.166670
 12y = 0.003740
 10b = 0.118696
 10o = 0.046837
 8b = -0.007574
 8o = -0.077262
 8bl = 0.009852
 6b = 0.012869
 6g = 0.461146
 6y = 0.210418
 4b = 0.080270
 4g = 0.052528
 4p = -0.078049
 maximum: 6g = 0.461146
 minimum: 12r = -0.166670

without noise

20b = -0.039071
 20g = 0.098584
 12r = -0.141735
 12y = -0.008002
 10b = 0.110771
 10o = 0.091263
 8b = -0.051865
 8o = -0.045028
 8bl = 0.070492
 6b = -0.042376
 6g = 0.413561
 6y = 0.204100
 4b = 0.077749
 4g = 0.093187
 4p = -0.097903
 maximum: 6g = 0.413561
 minimum: 12r = -0.141735

noise (00b, mean=5)

Table 3. Invariant Recognition to Occlusion and Noise








Figure 4 Response of memory at multiple viewpoints for the overlapped configuration