4 Rieger, C. An organization of knowledge for problem solving and language comprehension. *Artificial Intelligence*, *7*, 1976.

5 Rieger, C. Spontaneous computation in cognitive models. TR-459, University of Maryland Computer Science Technical Report Series, 1976..

6 Srinivas, S. Error recovery in a robot system. Ph.D. Thesis, C.I.T., 1976.

7 Yakimovsky, Y., and Cunningham, R. DABI--A data base for image analysis with nondeterministic inference capability. *Pattern Recognition and Artificial Intelligence*, Chen, C.H. (Ed.), Academic Press, Inc., New York, 1976, 554-592.

Cognitive Systems Based on Adaptive Algorithms

John H. Holland and Judith S. Reitman

Dept. of Computer Science and Dept. of Psychology (respectively)

University of Michigan Ann Arbor, MI 48104

The type of cognitive system (CS) studied here has four basic parts: (1) a set of interacting elementary productions, called *classifiers*, (2) a performance algorithm that directs the action of the system in the environment, (3) a simple learning algorithm that keeps a record of each classifier's success in bringing about rewards, and (4) a more complex learning algorithm, called the *genetic algorithm*, that modifies the set of classifiers so that variants of good classifiers persist and new, potentially better ones are created in a provably efficient manner.

The genetic algorithm can be shown to provide a selective pressure toward classifiers of appropriate generality by attending to each classifier's frequency of use (too particular classifiers will be infrequently used) and its consistency in predicting rewards (too general classifiers will predict inaccurately). Also, the genetic algorithm can be shown to build new classifiers based on the estimated value not of the entire classifier but of particular combinations of attributes within the classifiers. This information about the value of combinations of attributes can be found in the set of M classifiers in memory (each length n) through counting (ranking) the number of classifiers that have a particular combination of attributes. The counting never actually takes place, however; the genetic algorithm automatically selects high valued (high-ranking) attribute combinations. Furthermore, each time the genetic algorithm operates, it reranks $M \cdot 2^{n/2}$ combinations appropriately and uses these adjusted rankings to build new classifiers. This efficient maintenance and use of information is a feature of the genetic algorithm called *intrinsic parallelism*.

Two "proof-of-principle" experiments are reported. One experiment shows CS's performance in a maze when it has only the ability to adjust the predictions about ensuing rewards of classifiers (similar to adjusting the 'weight' of each classifier) vs. when the power of the genetic algorithm is added. Criterion was achieved an order of magnitude more rapidly when the genetic algorithm was operative. A second experiment examines transfer of learning. Placed in a more difficult maze, CS with experience in the simpler maze reaches criterion an order of magnitude more rapidly than CS without prior experience.

The CS is constructed with classifiers that act as condition-action productions; they broadcast *messages* and respond to certain classes of messages broadcast by other classifiers. The resulting system is therefore *computationally complete*, a feature which allows it to develop appropriately in a wide range of environments. And, through the *additional* use of the predictions attached to the classifiers, the system automatically generates an experience-based cognitive map, allowing it the capacity to look ahead and to apportion credit during non-rewarded intervals.

Knowledge-Directed Learning[1]

Elliot M. Soloway and Edward M. Riseman

Computer & Information Science Department

University of Mass. Amherst, MA 01003

Abstract

A system embodying a knowledge-directed approach to unsupervised learning is examined in this paper. This approach is based on the premise that knowledge of new situations is acquired and interpreted in terms of the previous knowledge brought to the learning situation. In particular, our system is provided with a general characterization of action-oriented competitive games. This frame of reference is used to construct an interpretation for the patterns of human activity that are observed in games of baseball.

Multiple levels of knowledge and processing are used to proceed through various levels of description of the observed human behavior. Hypothesis Generation shifts the pattern description from observed physical actions such as "catch" and "run" to inferred goals and causal relationships of the players executing those actions. Hypothesis Generalization abstracts generalized classes of events and schemata that represent concepts such as "hit" and "out". Hypothesis Evaluation closes the loop in the learning process by verifying or rejecting the various hypotheses. Knowledge encoded as schemata direct these processes; there are schemata for inferring competitive and cooperative goals and causal relationships of players.

An important aspect of the system is its ability to use acquired knowledge. The multi-level organization facilitates the integration of the new information into the existing knowledge structure. Also, both the initial knowledge and the acquired knowledge are represented uniformly as schemata (production rules). Acquired schemata, then, are available to assist in interpreting and predicting future events. This ability demonstrates the effectiveness of our knowledge-directed approach to learning.

I. Introduction

In this paper we outline the major points of a computer system embodying a knowledge-directed approach to learning. The motivation for this approach comes from our daily experience; it seems that when faced with a learning situation (e.g., understanding sequences of apparently novel events) one does not rely solely on statistical learning techniques. Rather, one uses various levels of knowledge and processing to focus in a highly directed fashion on what is important in the observations. This direction is provided by the predispositions, or frames [1, 2], used to interpret those observations.

In particular, our system is provided with a general characterization of action-oriented competitive games. It uses that frame of reference in order to construct an interpretation for the patterns of human activity in the observed games of baseball. These behavior patterns are described in terms of four attributes: *actor*, *action*, *location*, and *time*. The goal of the system is to acquire a hierarchical network of schemata and concepts that represent an understanding of the observed activity at various levels of abstraction. The generalized schemata and concepts capture the relationships between the actions of the players and the goals intended by those actions. A key objective of our research is to allow the acquired schemata to aid in the further understanding of the observed patterns of behavior. This learning process requires both general knowledge of the goals and causal relationships in competitive action-oriented games as well as knowledge about particular physical actions.

1. This work was supported by ARI grant DAHC19-76-G-0013.