Integrating Multiple Knowledge Representations and Learning Capabilities in an Expert System: The ADVISE System

R. S. Michalski
A. B. Baskin
Department of Computer Science
University of Illinois
Urbana, Illinois 61801-2987

ABSTRACT

The ADVISE system is an integrated set of tools for the development of and experimentation with expert systems in various specific application domains. It functions as a multi-purpose inference system that employs three knowledge representations: a rule-base, a conceptual network and a relational data base. In addition, it includes learning capabilities by incorporating the inductive learning programs GEM (for learning from examples) and CLUSTER (for constructing classifications).

Three expert systems have been developed using ADVISE: PLANT/ds (version 2)--for diagnosing soybean diseases, PLANT/cd—for predicting cutworm damage to corn, and BABY—a consultant for the neo-natal intensive care unit.

I INTRODUCTION

At present, the technology of expert systems is undergoing very rapid growth and is being applied to a wide spectrum of practical problems. Current systems, however, suffer from a number of limitations that restrict their usefulness. They typically employ only one form of knowledge representation, usually a rule-base representation (Davis & King, 1976; Buchanan & Duda, 1982). They have no learning abilities, use only one type *of* inference procedure, and use only a single control strategy. Our research on the ADVISE "meta-expert" system has the goal of overcoming these limitations. Some of the relevant papers to our research are by Duda, Hart &. Sutherland (1978), Van Melle (1979) and Feigenbaum (1983).

II MAJOR FEATURES OF ADVISE

The tools which comprise ADVISE form a "knowledge engineer's workbench'' for developing expert-systems in particular domains, specifically, for building knowledge bases and implementing appropriate inference procedures on them. Main features of ADVISE are:

- allows use of three different forms of knowledge/data representation (a rule-base, a conceptual network, and a relational database),
- supports different schemes for uncertainty propagation in a knowledge base (probabilistic, approximate Bayesian, min/max logic, weighting evidence),
- separates control and planning during a consultation *(strategic information)* from more procedural information such as specific action rules *(tactical information),*
- supports different control strategies chosen when the knowledge base is built (utility optimization, probabilistic network traversal, forward rule chaining, backward rule chaining).
- incorporates a number of existing tools for inductively deriving decision rules and control information from examples, and
- is implemented in Pascal and operates in a standard UNIX environment.

ADVISE supports multiple knowledge representations because different types of information about a problem domain can best be expressed using different formal structures. For instance, situation-action pairs can be efficiently represented as rules; interrelationships between concepts, constraints, or clues used in planning can be more suitably described in a network; a large collection of facts (e.g. patient records) can be stored and manipulated as a relational database. For similar reasons, ADVISE supports a number of different schemes for defining aspects of the consultation process on a problem specific basis. Even within a single knowledge base, it is frequently useful to mix strategies.

Another major feature of the system is that knowledge is acquired not only through the standard technique of "learning by instruction," but also through "learning from examples" and "learning from observation."

The latter two modes of learning have been implemented by modifying and adapting the inductive learning programs GEM and CLUSTER. The GEM program generalizes examples of different concepts and creates formal logic rules for recognizing these concepts. The program is a modified version of the inductive program AQII (Michalski & Larson, 1978). The feasibility of using this program for inductive knowledge acquisition for an expert system has been proven by its successful application in the development of an expert system for soybean pathology (Michalski and Chilausky, 1980; Michalski, et al., 1982).

The second program, CLUSTER, automatically constructs a classification of given entities using an approach called "conceptual clustering." Each grouping in the obtained classification corresponds to a class of entities (objects, situations, etc.) closely described by a single conjunctive concept representing a logical product of relations on selected object attributes. Nodes with the same parent in the hierarchy are optimized according to some global "quality" criterion of classification (Michalski & Stepp, 1983).

The system also includes several inferential programs that support inductive learning and deal with the selection of most relevant attributes,, selection of most representative learning examples, etc. (Spackman, 1983).

III COMPONENTS OF ADVISE

A functional diagram of the ADVISE system is shown in Figure 1. The system consists of four major components: A) Control block and user interface, B) Knowledge base, C) Query block and D) Knowledge acquisition block.

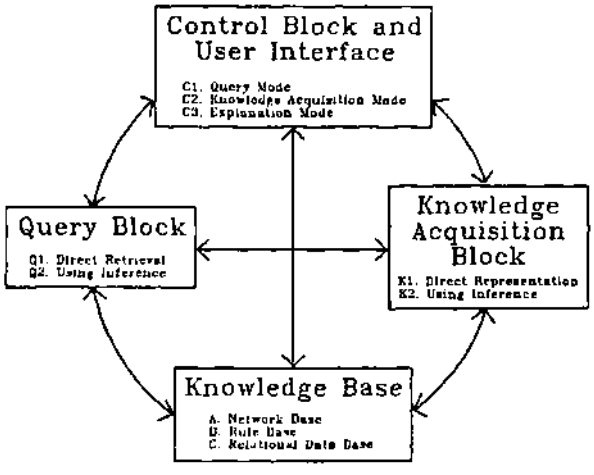

Figure 1. A block diagram of ADVISE showing the major functional components of the system

A. Control Block and User Interface

This component provides three modes of user interaction with other components:

- Query mode
- Knowledge acquisition mode
- Explanation mode.

The user interface portion of the system provides utility routines to manage the terminal screen and supports the explanation of internal data representations and operations in a textual and graphical form.

1. Query mode (Q-mode)

Query mode is used during consultation. In this mode, the system:

- selects questions to ask the user,
- accepts user answers and
- conducts an inference process involving the knowledge base and information provided by the user in order to compute advice with an associated strength of supporting evidence.

In one respect, the consulting portion of the system is radical in design. There is no single problem solving strategy. Rather, localized problem solving behavior is defined by the choice of an evaluation scheme (three of which are implemented) and global problem solving behavior is governed by a choice of control scheme (three of which are also implemented). Explicit provision has been made in the design for the inclusion of additional control schemes and additional evaluation schemes.

2. Knowledge acquisition mode (K-mode)

The knowledge acquisition mode coordinates both the encoding of expert derived rules in a knowledge base and the interactive invocation of the separate induction programs. This mode includes modules for defining expert rules, manual refinement of rules, induction of rules from examples, and automated correction and improvement of the rules. The system also provides facilities for testing rules in interactive mode on individual cases, as well as in batch mode on a collection of cases.

3. Explanation mode (Ivmode)

The explanation mode paraphrases decision rules, enables a user to understand the organization and functioning of the system in query and knowledge acquisition modes, allows simple interrogation of the contents of the knowledge base, and displays the steps in the process which led to a given advice.

B. Knowledge Base

The knowledge base integrates three types of representation through a unified "access protocol:"

- a network base (e.g., conceptual network),
- a rule base, and
- a relational data base.

The network base contains structures for representing domain knowledge about relationships among various conceptual units. The network organization is a form of the "Logic Net" formalism described in (Baskin 1980).

The rule base consists of inference rules in the form:

CTX CONDITION => CONCLUSION: a,B

where

CTX is a logical expression defining the context in which the rule is applicable,

CONDITION is a logical expression defining the conditions that have to be satisfied to assert the CONCLUSION,

α - the strength of forward implication (from CONDITION to CONCLUSION),

B - the strength of backward implication (from CONCLUSION to CONDITION).

The relational data base consists of a collection of relational tables. The tables can be created, modified and retrieved through ordinary relational database operators. In addition, a variety of inferential operators can be applied to the database through the facilities of QUIN (Query & inference) subsystem (Spackman, 1983).

C. Query Block

The Query block contains the programs supporting retrieval from the knowledge base. Direct retrieval is supported as a relational algebra for tables and a simple infix display of rules. Retrieval with inference will deductively and inductively extend the knowledge base in response to a user query'

D. Knowledge Acquisition Block

The system supports knowledge acquisition for networks, tables, and rules. Direct encoding is supported by input of tables or rules (''learning by being told") and knowledge acquisition through induction is supported by linkage to learning programs. Additional details about the Query and Knowledge Acquisition blocks can be found in (Michalski, et. al., 1983).

IV SUMMARY

The major modules of the ADVISE system have been implemented in the PASCAL language (approximately 45000 lines of code) for the VAX 780 computer under Berkeley-Unix operating system. The knowledge base management component of the system is well developed and has been used extensively. The query block is partially complete and is best developed for the relational database. Knowledge acquisition currently supports direct representation of information supplied by a human expert as well as induction of classifications and rules. The explanation mode is least developed and currently provides little more than an English paraphrase of rules and links within the network.

The system has already been applied for the development of three expert systems: PLANT/ds (v.2), PLANT/dc, and BABY. PLANT/ds (v.2) provides advice on the diagnosis of common soybean diseases in Illinois. PLANT/dc predicts damage to corn due to black cutworms. Its operation is based on a close cooperation "surface" model utilizing a rule base, and a "deep" model simulating the black cutworm corn development (Boulanger, 1983). The system called BABY provides expert advice on monitor and laboratory data for the neonatal intensive care unit. The system reasons from the available data toward speci6c pathologic states using a Bayesian update scheme (Rodewald, 1983).

ACKNOWLEDGEMENTS

The research described here has been supported in part by the Office of Naval Research, Grant No. N000I4-82-K-0186 and in part by the National Science Foundation, Grant No. MCS 82-05166. The authors thank A. Boulanger, M. Seyler, L. Rodewald, K. Spackman, C. Uhrik and R. Reinke for their contribution to the development of ADVISE, its first experimental applications and comments on this paper.

REFERENCES

Baskin, A. B., "LOGIC NETS: Variable-valued Logic plus Semantic Networks." in Policy Analysis and Information Systems, No. 3, 1980.

Boulanger, A., "The Expert System PLANT/cd: A case study in applying the general purpose inference system ADVISE to predicting black cutworm damage to corn," M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, Illinois, 1983.

Buchanan, B. and Duda, R., "Principles of Rule-Based Expert Systems," Hueristic Programming Project, Report No. HPP-82-14, August 1982.

Davis, R. and King, J., "An Overview of Production Systems" in Machine Intelligence 8, Elcock and Michie (eds.), 1976.

Duda, R. O., Hart, P. E. and Sutherland, G. L., "Semantic Network Representation in Rule-Based Inference Systems," in Pattern Directed Inference Systems, D. A. Waterman and F. Hayes-Roth (eds.), Academic Press, 1978.

Feigenbaum, E. A., "Knowledge Engineering for the 1980's," Computer Science Department, Stanford University, 1982.

Michalski, R. S. and Larson, J. B., "Selection of Most Representative Training Examples and Incremental Generation of VL, Hypothesis: The Underlying Methodology and the Descriptions of Programs ESEL and AQ11," Report No. 877, Department of Computer Science, University of Illinois, Urbana, Illinois, 1978.

Michalski, R. S. and Chilausky, R. L., "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, 1980.

Michalski, R S., Davis, J. H , Bisht, V. S. and Sinclair, J. B., "PLANT/ds: An Expert System for the Diagnosis of Soybean Diseases," European Conference on Artificial Intelligence, Orsay, France, July 12-14, 1982.

Michalski, R. S., "A Theory and Methodology of Inductive Learning," Chapter in the book, MACHINE LEARNING: An Artificial Intelligence Approach, TIOGA Publishing Co., R. S. Michalski, J. Carbonell and T. Mitchell (eds.) 1983.

Michalski, R. S. and Stepp, R., MACHINE LEARNING: An Artificial Intelligence Approach, TIOGA Publishing Co., R S. Michalski, J. Carbonell and T. Mitchell (eds.), 1983.

Michalski, R. S., Baskin, A. B., Boulanger, A. G., Seyler, M. R., "A Technical Description of the ADVISE Meta-expert system," in preparation.

Rodewald, L. E., "BABY: An expert system for interpretive reporting in a neo-natal intensive care unit," M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, Illinois, 1983.

Spackman, Kent, "Integration of Inferential Operators with a Relational Database in an Expert System," Masters Thesis, Department of Computer Science, University of Illinois, Urbana, Illinois, 1983.

Van Melle, W., "A Domain Independent Production Rule System for Consultation Programs," Proc. Sixth IJCAI, 1979.