

DTIC FILE COPY

4

AD-A197 241

RUTGERS UNIVERSITY
Center for Expert Systems Research

Quarterly Report:
*Empirical Analysis and Refinement of
Expert System Knowledge Bases*

Contract Number
N00014-87-K-0398
Office of Naval Research

May 31, 1988

Principal Investigators:
Sholom M. Weiss
Casimir A. Kulikowski

DTIC
ELECTE
JUL 12 1988
S D

DECLASSIFICATION STATEMENT A
This report is classified "Secret"
and is not to be released without
prior approval of the Office of Naval Research

1. Technical Project Summary

Knowledge base refinement is the modification of an existing expert system knowledge base with the goals of localizing specific weaknesses in a knowledge base and improving an expert system's performance. Systems that automate some aspects of knowledge base refinement can have a significant impact on the related problems of knowledge base acquisition, maintenance, verification, and learning from experience. The SEEK system was the first expert system framework to integrate large-scale performance information into all phases of knowledge base development and to provide automatic information about rule refinement. A recently developed successor system, SEEK2, significantly expands the scope of the original system in terms of generality and automated capabilities.

Based on promising results using the SEEK approach, we believe that significant progress can be made in expert system techniques for knowledge acquisition, knowledge base refinement, maintenance, and verification.

2. Principal Expected Innovations

We are proposing to demonstrate a rule refinement system in an application of the diagnosis of complex equipment failure. The expected candidate application is computer network troubleshooting. The expert system should demonstrate the following advanced capabilities:

- automatic localization of knowledge base weaknesses
- automatic repair (refinement) of poorly performing rules
- automatic verification of new knowledge base rules
- some automatic learning capabilities.

3. Objectives for FY88

1. functioning equipment diagnosis and repair knowledge base, suitable for refinement (expected in the area of computer networks).
2. initial demonstration of functioning equipment diagnostic system with capabilities of localization of weak rules, automatic refinement, automatic verification.
3. demonstration of initial rule learning capabilities.



by	per ltr
Date	
Approved	Order
Date	
A-1	

4. Summary of Progress

Here are the highlights of progress has been made in meeting our stated objectives for fiscal 1988:

- Last quarter, Dr. Peter Politakis of the Digital Equipment Co. transferred to us DEC's Network Troubleshooting Consultant program that we proposed to use in our system. Dr. Politakis directed the development of this software and serves as our expert in the refinement of the knowledge base. Previously, we circumscribed the knowledge base to the following problem types: line, circuit, or cable problems. During the last quarter, the subset of the knowledge base consisted of 287 observations, 138 hypotheses, and 324 rules. During the current quarter, we further revised the knowledge base. At the present time the KB consists of 215 observations, 148 hypotheses, and 390 rules. The purpose of this application is to serve as a vehicle for further experimentation. We expect the knowledge base to remain stable for the remainder of the contract while we develop systems with advanced refinement and learning capabilities.
- During the previous quarter, we noted that Politakis had obtained documented cases of network problems. He had supplied about a dozen, and we hoped to obtain others from DEC's stored records. During the present quarter, we were able to obtain an additional 60 cases. This brings the total to 72 cases. We will supplement a core group of documented cases with simulated cases derived from verified correct rules in the knowledge base.
- In our previous quarterly report, we noted that substantial progress was being made in our rule induction, i.e. learning system. Several experiments have been underway using data obtained from other researchers who have published results. These include data from Michalski and Quinlan. These efforts are extensions of the procedures we reported at the AAAI-87 conference [Weiss, Galen, and Tadepalli 87]. We note that unlike other fields, it is unusual for AI researchers to re-analyze other researchers data. Complete details of the experimental results have appeared in a technical report entitled *Minimizing Error Rates for Induced Production Rules*. The abstract of this technical report and some experimental results are reported in the next section.

In terms of the three objectives for fiscal year 1988, we have completed the first objective: producing a functioning computer network diagnostic and repair knowledge base suitable for refinement.

The second objective was for an initial demonstration of functioning equipment diagnostic system with capabilities of localization of weak rules, automatic refinement, automatic verification. We believe the current system has these capabilities. However, the knowledge base we have produced is already quite accurate and therefore has limited potential for further refinement. While additional topics could be covered by adding many new rules, this is a not a principal objective. We have embarked on a novel approach to testing the system. Because the current knowledge base is considered correct, we feel we can develop the following tools for experimentation:

- A case generator that randomly generates cases for given hypothesis from a correct knowledge base. This allows us to gather many more *simulated* cases than is otherwise possible.
- A rule modifier that randomly changes rules in a given knowledge base. In effect it introduces errors into the rules.

These tools will allow us to randomly modify a correct knowledge base and see whether the refinement system can recover from the errors. We expect that these tools will be completed during the next quarter, and that the second FY88 objective will be fully met.

The third FY88 objective is a demonstration of initial rule learning capabilities. The work reported in the next section and in our technical report [Weiss 88], further amplifies on a new approach to pure rule induction. For applications where a relatively short rule is required or can provide a good solution, our Predictive Value Optimization (PVO) procedure appears superior to other rule induction procedures reported in the literature.

PVO is an autonomous induction system that learns rules in restricted situations. During the contract period we expect to integrate this procedure into the overall knowledge base refinement system. However, during the next quarter we expect to produce heuristics and procedures that can immediately produce a learning capability within the context of the SEEK2 refinement system.

Progress in Rule Induction Techniques

During the current quarter, we completed our comparative experiments on rule induction. We have issued a technical report entitled *Minimizing Error Rates for Induced Production Rules*. We reproduce the abstract and a few of the key results below.

Abstract: Empirical techniques for induction of decision rules have evolved from procedures that cover all cases in a data base to more accurate procedures for estimating error by train and test sampling. Procedures that prune a set of decision rules and the components of these rules have been successful in increasing the performance of an induced rule set on new test cases. Recently, we reported on a technique for learning the single best decision rule of a fixed length. In this paper we show how resampling techniques for estimating error rates, can be integrated into this procedure for induction of decision rules. Superior results are reported on data sets previously analyzed in the AI literature.

In 1987, we reported on a technique for learning the *single* best decision rule of a fixed length [Weiss, Galen, and Tadepalli 87]. In contrast to other methods of rule induction, the PVO rule induction procedure does not generate and prune a complete set of decision rules. Instead, this method is an approximation to exhaustive generation of all possible rules of a fixed length. While a true exhaustive search is not feasible in most applications, a small number of heuristics reduce the search space to manageable proportions.

Experiments were performed on two sets of data for which published studies are available. The results are summarized in Figures 4-1 and 4-2.

Method	Variables	Rules	Error Rate
AQ15	7	2	32%
PVO	2	1	23%

Figure 4-1: Comparative Summary for AQ15 and PVO on [Michalski, Moztetic, Hong, and Lavrac 86] Data

Method	Variables	Rules	Errors (1985)	Errors (1986)
C4 pruned rules	8	2	31	43
PVO random resampling	8	2	17	30

Figure 4-2: Comparative Summary for C4 and PVO on [Quinlan 87] Data

In this paper, we re-analyzed data that had been analyzed using prominent machine learning techniques. We showed that superior rules could be induced from these data sets. In the case of the Michalski data, a simple two variable rule produces better results than the more complex rules cited in the literature. While Quinlan's original data analysis produced excellent results, we showed that somewhat better rules could be induced than those he cited in his reports on thyroid disease.

For our analysis, we used the classical resampling techniques of statistical pattern recognition to estimate error rates for nonparametric classifiers. These techniques can be time-consuming, but can lead to better induction results. Because PVO induces rules for a fixed, relatively short length, resampling procedures are a natural extension of the basic method. The major advantage is that error estimates can be derived, while essentially the complete data sample may be used for classifier design. While resampling is a natural fit to PVO, its use with other induction techniques is feasible.

We do not claim that PVO is universally superior to other empirical rule induction procedures. Unlike AQ15 or C4, in practice PVO is limited to the induction of single short rules. However, if a good solution exists in the form of a single short rule, PVO has a decided advantage. Unlike incremental empirical induction procedures that select one test at a time, PVO examines combinations of tests with varying constants. There are many applications, such as expensive instrument testing, where a short rule that limits the number of tests to be performed is a requirement.

References

- [Michalski, Mozetic, Hong, and Lavrac 86]
 Michalski, R., Mozetic, I., Hong, J., and Lavrac, N.
 The Multi-purpose Incremental Learning System AQ15 and its Testing
 Application to Three Medical Domains.
 In *Proceedings of the Fifth Annual National Conference on Artificial Intelligence*, pages
 1041-1045. Philadelphia, Pa., 1986.

- [Quinlan 87] Quinlan, J.
 Simplifying Decision Trees.
Int. Journal of Man-Machine Studies :in press, 1987.
 also Tech. Report 87.4, New South Wales Institute of Tecnology, School of
 Computing Sciences.

- [Weiss 88] Weiss, S.
Minimizing Error Rates for Induced Production Rules..
 Technical Report LCSR-TR-106, Rutgers University, Department of Computer
 Science, 1988.

- [Weiss, Galen, and Tadepalli 87]
 Weiss, S., Galen, R., and Tadepalli, P.
 Optimizing the Predictive Value of Diagnostic Decision Rules.
 In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*.
 Seattle, Washington, 1987.
 in press.

END

DATE

FILMED

DTIC

9-88