* For the reference see the last page

A Cost Effective Question Asking Strategy

Evangelos Triantaphyllou[1] and Jinchang Wang[2]

[1]: Kansas State University
[2]: Missouri Western State College

Abstract

One of the main objectives of a successful expert system is to be able to reach a conclusion by asking a small number of questions. However, this may be a misleading approach when the cost of answering questions posed by the system varies. When different questions are associated with different costs then, it is possible a question asking strategy to be very costly when it focuses only on minimizing the total number of questions. This paper expands an existing effective question asking strategy and illustrates the development of a cost effective question asking strategy for Horn clause systems.

Introduction

The great proliferation of the use of rule based systems in a vast area of applications has defined the need for new research problems to be solved. Most of the new problems have emerged from the requirement to handle larger rule bases. When the size of a rule base increases, the number of questions posed to the user by the inference engine may increase as well. A good question asking strategy may lead the inference process to the final goal efficiently.

A question asking strategy is **efficient** if it can generate the pertinent questions quickly. Furthermore, such a strategy is **effective** if it can reach the final goal by asking as few questions as possible. A second way to define effectiveness is by minimizing the total cost needed in answering the pertinent questions posed by the system. This paper presents the development of a question asking strategy for Horn clause systems which is both efficient and effective.

Some of the early applications of experts systems consider the issue of question asking strategies. The system EXPERT [Hayes-Roth, Waterman and Lenat, 1983] uses pre- ordered lists of rules and questions. Also the PROSPECT system [Duda, Gasching and Hart, 1979] uses a scoring function in implementing a question asking strategy. Another question asking strategy, called "Alpha-Beta pruning", was introduced by Mellish [1985] for acyclic inference nets. In that strategy irrelevant questions are dropped from consideration.

de Kleer and William [1987] developed a general entropy technique for diagnosing multiple faults in an electrical circuit. The idea of that technique can be applied to the problem of question selection in rule bases. This is possible because a rule base may be viewed as an electrical circuit composed of AND-gates and OR-gates. Wang and Vande Vate [1990] proved that an effective strategy, which optimally selects the next question, is never efficient. That is, it is NP-hard to design such a strategy. This is also the case for Horn clause systems. They also developed the concept of the Minimum Usage Set (or MUS) strategy which attempts to approximate effective strategies in O(nlogn) time.

In general, a process for choosing a question is composed of two steps [Hayes-Roth, Waterman, and Lenat, 1983]. These two steps are goal selection and question selection. In the goal selection step, a question asking strategy chooses a potential conclusion to pursue. In the question selection step, the strategy selects a question which may help to reach the goal selected in the first step.

The present paper presents some theoretical results regarding the problem of selecting the next question given a set of candidate questions. Each question is an inquiry of the value (true or false) of the corresponding variable. It is assumed that the system knows the likelihood (or probability) that a given variable has true value. It is also assumed that for each such question the system knows the cost in answering that question.

In this paper it is first assumed that all these costs are equal. Therefore, the objective in this case is to try to **minimize the total number of questions** required to prove a given goal. Next, it is assumed that questions are associated with different costs. In this case the objective is to try to **minimize the total cost** in proving a given goal. Some experimental results are also provided. Since the proposed strategies are based on the original MUS strategy [Wang and Vande Vate, 1990], the next section briefly describes this strategy.

The MUS Strategy

A <u>Horn clause</u> is a rule in which a finite set of positive assertions (i.e., binary variables) implies one positive conclusion. Consider a collection of Horn clauses. Let the set N index the assertions in this collection. These assertions are denoted as A_j, where $j \in N$. The set of natural numbers D(i) indexes the Horn clauses which have conclusion A_i. Let the ordered pair (i,d) denote the d-th clause with conclusion A_i. Furthermore, let I(i,d) index the assertions in the antecedent part of the clause (i,d). This notation is adopted from [Wang and Vande Vate, 1990]. Based on this notation a <u>Horn clause system</u> can be expressed as:

IF (A_j is true for all $j \in I(i,d)$) THEN (A_i is true)
for any $i \in N$ and $d \in D(i)$ (1)

Horn clause systems are a special class of a logical system which is widely used in practice (see, for instance, [Hooker, 1988a and 1988b]). PROLOG, a popular logic programming language, relies on Horn clauses. The most distinguished characteristic of Horn clause systems is that inference in a Horn system can be carried out quickly [Dowling and Gallier, 1984], although inference is hard in ordinary logical systems. Jeroslow and Wang [1989] observed that inference in a Horn clause system can be represented by a Leontief flow problem which provides information on the proof structure.

If the value of an assertion (variable) can be provided by the user we call it an <u>observable assertion</u>. An assertion is called an <u>unconfirmed observable variable (or UOV)</u> if it is observable but its value is not given yet. Let A_i be a potential goal to be proved. A set of unconfirmed assertions is an <u>unconfirmed observable set (or UOV set) of A_i</u>, if the assertion (goal) A_i can be proved true when all the assertions (variables) in that set have "true" value. But if any one of these assertions were false, then A_i could not be concluded from the others. It is possible an assertion A_i to have many such sets. The smallest such set is called the <u>minimum inquiry set of A_i</u>. A smallest UOV set over all the UOV sets of the potential goal A_i is called a <u>global minimum inquiry set of A_i</u>.

Wang and Vande Vate [1990] showed that to determine the global minimum **inquiry** set is an **NP-hard problem.** They proved it by using the concept of the so-called sub-effective strategy. Given the problem of proving whether a goal A_j (where A_j is a non-observable assertion) is true or false, a sub-effective strategy selects that UOV set S of A_j which is **most likely** to have **all** its members affirmed. In other words, the product

$$\prod_{A_i \in S} P_i ,$$

is maximum (where P_i denotes the likelihood or probability that an unconfirmed observable variable A_i is true).

However, the problem of finding the previous set S is **NP-hard.** The **usage of A_i on A_k** is defined as the number of times A_i is used in the process of proving a final goal A_k. The <u>usage of an unconfirmed observable set S_j of A_i</u> is the sum of the usages of the assertions in S_j on A_i. The <u>minimum usage set of A_i</u> is an unconfirmed observable set of A_i with the minimal usage. The global minimum usage set has the smallest usage among all UOV sets. In fact, the following is true for this set:

$$\prod_{A_i \in S_j} P_i^{u_i} \quad \text{is maximum,}$$

where u_i is the usage of A_i in proving the goal corresponding to the set S_j.

A global minimum usage set can be found in <u>log-linear</u> time by using the labeling algorithm developed in [Wang and Vande Vate, 1990]. The original MUS strategy was organized in terms of the following steps:

Step 1: *Select an MUS set (it takes O(nlogn) time).*
Step 2: *Select to ask a question about an unconfirmed variable in the MUS set.*
Step 3: *If the response is true, go back to step 2. If the response is false, then select a new MUS set and go back to step 1.*
Stopping rules:
Stop: *If all of the variables in the MUS set are true. In this case the goal is proved to be true.*
Stop: *If no new MUS set can be determined. In this case the goal cannot be proved.*

The following theorem states an optimal question selection rule <u>only for step 2</u>, above. Therefore, by using this rule it is not guaranteed that the entire strategy will be optimal. This theorem is proved in [Triantaphyllou and Wang, 1991].

Table 1. The Seven Strategies.

UOV Set Selection Scenarios		Variable Selection Scenarios		
		Random	MLT	LLT
Random	(RND)	RND-RND	RND-MLT	RND-LLT
Sub-effective	(SUB)	not considered	SUB-MLT	SUB-LLT
MUS		not considered	MUS-MLT	MUS-LLT

THEOREM 1:
If the next question in proving a goal A_k is to be selected from a given UOV set of A_k in a Horn clause system then, in order to minimize the expected number of questions, the next question must be on the value of the assertion which is least likely to be true compared to the other assertions in the UOV set.

In other words, the best strategy is to ask first on the value of the variable which is least likely to be true. As it will be shown in the next section, this modification in step 2, makes the original MUS strategy to be almost optimal.

Computational Experiments

Computer experiments compared the performance of seven question asking strategies on 64,000 randomly generated problems. For each problem it was assumed, without loss of generality, that there was only one potential goal to be proved. The test problems were in the form of a group of UOV sets. Since the objective of these experiments was to determine the number of questions required by each question asking strategy, deduction was not necessary and so the problems did not need to be in the form of "IF ... THEN ..." rules.

The test problems were randomly generated with two parameters; the number of observable variables and the number of UOV sets. For each parameter the eight values, 10, 20, 30, ..., 80 were considered. Furthermore, each observable variable was assumed to have probability 0.50 of belonging to a UOV set. For each combination of a number of variables and a number of UOV sets, 1,000 random problems were tested. Since there are 8x8 = 64 different cases of test problems, 64,000 random problems were solved by applying the seven strategies.

The probability P_i, that an observable assertion (variable) A_i is true, was a random number from the interval [0, 1]. Furthermore, the usages, u_i, were random integers uniformly distributed in the interval [1, 11] (this interval was considered arbitrarily). Each problem was tested according to the following steps:

Step 1: *Select a UOV set.*
Step 2: *Select to ask a question about an unconfirmed variable in the UOV set.*
Step 3: *If the response is true, go back to step 2. If the response is false, drop all the UOV sets that include this false variable and go back to step 1.*
Stopping rules:
Stop: *If all of the variables in an UOV set are true. In this case the goal is proved to be true.*
Stop: *If no UOV set is left. In this case the goal cannot be proved.*

All the strategies that were examined were comprised of two major phases. In the first phase a UOV set was selected (step 1). In the second phase a variable was selected from the previous UOV set (step 2). For the case of the UOV set two major scenarios were examined. In the fist scenario the UOV set was the set needed by a sub-effective strategy. That is, the set S which maximizes the product $\prod_{A_i \in S} P_i$. That set was called the SUB set.

In the second scenario the UOV set was the MUS set of the goal to be proved.
The next question was based on a variable of the UOV set selected in the first step. Three scenarios were considered in selecting the next variable (step 2). In the first scenario the variable was randomly selected. The second scenario was to select the most-likely-true (or MLT) variable. That is, first asking the variable which is most likely to be true. The third scenario was to select the least-likely-true (or LLT) variable. That is, first asking the variable which is least likely to be true. Besides the previous six strategies, the strategy of selecting a random UOV set with a random variable (or RND-RND) was also considered. These seven strategies are depicted in Table 1.
A UOV set was determined as follows (see also Table 1):

Under the random scenario, by choosing randomly a UOV set.

Under the sub-effective scenario, by choosing a UOV set S such that the $\prod_{A_i \in S} P_i$, product was maximum.

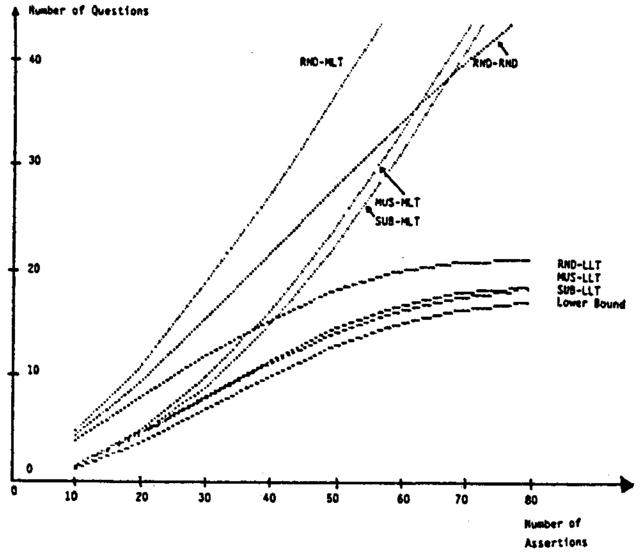

Figure 1.
Number of Questions Under Different Strategies

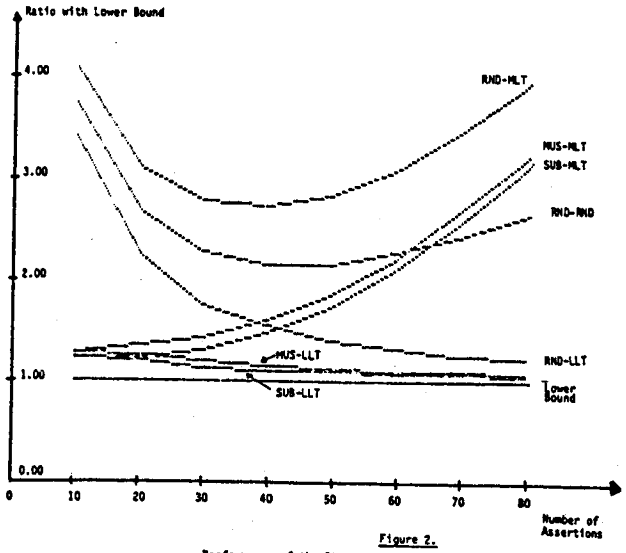

Figure 2.
Performance of the Strategies Relative to the Lower Bound

Under the MUS scenario, by choosing a UOV set S such that the product $\prod_{A_i \in S} P_i^{u_i}$, was maximum. That is, by selecting the MUS set.

As it can be seen from the way these test problems were considered in that investigation, the sub-effective strategy (SUB) and MUS sets can be determined by simply examining the values of the corresponding products. In this way the performance of a sub-effective strategy can be studied **without having to solve the NP-hard problem** needed in finding the required maximum product. In other words, the test problems **assumed** that all the related UOV sets are known a priory. However, in a real situation the UOV sets are **not known** a priory and finding the sub-effective strategy (SUB) set **is an NP-hard problem,** while finding the MUS set takes **only log-linear** time [Wang and Vande Vate, 1990].

Figures 1 and 2 illustrate the results of these experiments. Figure 1 illustrates the average number of questions asked under each strategy. The horizontal axis depicts the number of variables. The vertical axis depicts the average number of questions asked under a strategy.

These figures also present the lower bound of questions for each case. Since the actual values of the variables in the UOV variables were assumed to be known for the purpose of these simulated experiments, a lower bound of the number of questions asked by any strategy was calculated as a set covering problem [Triantaphyllou and Wang, 1991].

Figure 2 presents the performance of each strategy relative to the lower bound described above. For this reason, the lower bound is represented by a horizontal line with value on the vertical axis equal to 1.00. All the other lines were normalized subject to this line.

It can be easily seen from the experimental results that the MUS strategy is a good approximation of the sub-effective strategy on these random test problems. Taking under consideration the fact that the MUS strategy takes only O(nlogn) time, it is concluded that this strategy is also a good practical question asking approach. The same results also demonstrate that given a UOV set then, by starting with the least-likely-true variable is much better than starting with the worst-likely-true or just with a random variable.

Questioning Rules for AND-gates and OR-gates

In this section the original MUS strategy is extended to the situation where each question is associated with a different answering cost. Therefore, the objective now is to reach the final conclusion by trying to minimize the total questioning cost.

In this section a Horn clause system is viewed as a set of AND-gates and OR-gates. Then, two theorems are stated regarding these two types of gates. The first theorem describes the conditions for optimally asking questions **given** a clause. The second theorem describes what should be the optimal clause to be selected first. A heuristic approach is also developed in this section. This heuristic approximates the optimal UOV set, to be selected in step 1, for the case of having different answering costs.

A Horn clause (i,d): (IF (A_j is true for all $j \in I(i,d)$) THEN (A_i is true)), can be viewed as an AND-gate. This is possible because the assertion A_i is true exactly if every assertion A_j in the antecedent part of the clause (i,d) is true. Furthermore, it can be observed here that the assertion A_i can be proved by using any clause (i,d), where $d \in D(i)$. Therefore, such an assertion can be viewed as an OR-gate.

Let AND[(A_1, A_2, A_3, ..., A_k), A_t] denote an AND-gate. This gate represents the Horn clause with antecedent part the assertions (A_1, A_2, A_3, ..., A_k) and consequent part the assertion A_t. Let OR[(i1, i2, i3, ..., il), A_t] denote an

OR-gate. This gate indicates that assertion A_t can be proved by using any one of the following clauses: (t,i1), (t,i2), (t,i3),..., (t,il).

Let C_i be the cost in answering the question whether assertion A_i is true or false. Then, it can be easily proved that theorem 2 gives the necessary and sufficient conditions for asking questions **within a given** AND-gate so that, C(t,d), the expected cost (within that AND-gate) will be minimal.

THEOREM 2:
For an AND-gate, defined as AND[(A_1, A_2, A_3, ..., A_k), A_t], the questioning sequence (A_1, A_2, A_3, ..., A_k) yields the minimum expected cost if and only if the following is true

$$\frac{C_i}{1-P_i} \le \frac{C_{i+1}}{1-P_{i+1}} \quad \text{for } i=1,2,3,\ldots,k-1.$$

Furthermore, the expected minimum cost is:

$$C(i,d) = C_1 + C_2P_1 + C_3P_1P_2 + \ldots + C_k\prod_{i=1}^{k-1} P_i .$$

In the previous theorem (t,d) denotes the d-th Horn clause which has as consequent the assertion A_t. It is also assumed that this clause has in the antecedent part the assertions {A_1, A_2, A_3, ..., A_k}. Let P(t,d) denote the probability that all the assertions in the antecedent part of clause (t,d) are true. Then, the following relation is true:

$$P(t,d) = \prod_{j\in I(t,d)} P_j .$$

At this point it is assumed that the clauses involved in an OR-gate do not have assertions in common. Then, it can be easily proved that the following theorem states the necessary and sufficient conditions for selecting the next clause, so that the total cost will be minimal.

THEOREM 3:
For an OR-gate, denoted as OR[(1, 2, 3, ..., k), A_t], the questioning sequence (1, 2, 3, ..., k) yields the minimum expected cost if and only if the following is true

$$\frac{C(i,t)}{P(i,t)} \le \frac{C(i+1,t)}{P(i+1,t)} \quad \text{for } i=1,2,3,\ldots,t-1.$$

Furthermore, C_{At}, the expected minimum cost is:

$$C_{At} = C(t,1) + C(t,2)\,(1-P(t,1)) + C(t,3)\,(1-P(t,1))\,(1-P(t,2)) + \ldots + C(t,k)\prod_{i=1}^{k-1}(1-P(t,i)) .$$

If all the costs C_i are equal then, theorem 2 states that given an AND-gate the best strategy is to start with the least- likely-true assertion first. This is in agreement with the approach and experimental results presented in section 3. The next section uses the last two theorem to design an efficient heuristic when the total questioning cost is under consideration.

A Heuristic Approach for Selecting a Cost-Effective UOV Set

When different assertions bear different question costs, then it is desired to find in step 1 that UOV set (i.e., set of candidate questions) which has the minimum expected cost. Once this UOV set is determined, then theorem 2 can be used to select the questions.

Determining the optimal UOV set in step 1 is an NP-hard problem. This can be proved easily as in the UOV set selection problem discussed in [Wang and Vande Vate, 1990]. However, here an efficient heuristic approach is presented to solve the UOV set selection problem. The proposed approach takes only log-linear time. This approach is based on the labeling algorithm presented in [Wang and Vande Vate, 1990] and on theorems 2 and 3.

Consider a Horn clause system given as (1) in section 2. Let C_i be the cost to determine the value of assertion (variable) A_i. As it was noted earlier, let P_i be the probability that assertion A_i has true value. If A_i is a non-observable variable, then it is assumed that C_i equals to the infinity.

In the following developments LC_i and LC(i,d) denote the **labels of cost** for variable A_i and clause (i,d), respectively. Similarly, LP_i and LP(i,d) denote the **labels of probability** for variable A_i and clause (i,d), respectively. Given the previous terminology the labeling procedure is as follows:

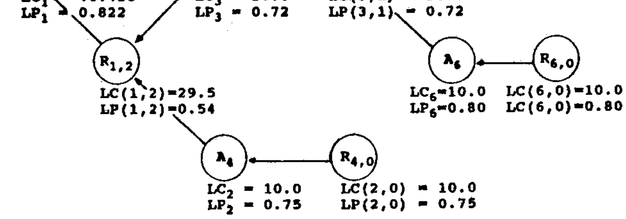

Figure 3. Diagram and Labels for Example.

LABELING PROCEDURE

Begin { Labeling Procedure }

For each unconfirmed observable variable A_i form labels

LC_i and LP_i as follows: $LC_i = C_i$ and $LP_i = P_i$

WHILE there is an unlabeled element t, all of whose immediate nodes have been labeled, label t according to the following two rules:

__Rule 1:__ If the element t is a clause (i,d), then the labels LC(i,d) and LP(i,d) are:

$$LC(i,d) = C_{j_1} + C_{j_2}P_{j_1} + C_{j_3}P_{j_1}P_{j_2} + \dots + C_{j_m}\prod_{k=1}^{m-1} P_{j_k}, \qquad (1)$$

where: $j_1, j_2, j_3, \dots, j_m \in I(i,d)$

and m is the size of the set $I(i,d)$

$$\text{and} \quad \frac{C_{j_k}}{1-P_{j_k}} \le \frac{C_{j_{k+1}}}{1-P_{j_{k+1}}} \quad \text{for } k=1,2,3,\dots,m-1,$$

$$LP(i,d) = \prod_{j \in I(i,d)} P_j. \qquad (2)$$

__Rule 2:__ If the element t is an assertion A_i, then the labels LC_i and LP_i are:

$$LC_i = LC(i,d_1) + LC(i,d_2)(1-LP(i,d_1)) + LC(i,d_3)(1-LP(i,d_1)(1-LP(i,d_2)) + \dots + LC(i,d_m)\prod_{k=1}^{m-1}(1-LP(i,d_k)), \qquad (3)$$

where assertion A_i is the conclusion of the clauses (i,d_1), (i,d_2), (i,d_3),..., (i,d_m) and only of these clauses.

$$\text{and} \quad \frac{LC(i,d_k)}{LP(i,d_k)} \le \frac{LC(i,d_{k+1})}{LP(i,d_{k+1})} \quad \text{for } k=1,2,3,\dots,m-1,$$

$$LP_i = 1 - \prod_{k=1}^{m}(1-LP(i,d_k)). \qquad (4)$$

End { Labeling Procedure }

Relations (1) and (3), above, are derived by applying theorems 2 and 3, respectively. If the final goal is to determine the value of the non-observable variable A_i, then the previous labeling algorithm leads to a cost-effective UOV set for step 1. Suppose that in a Horn system there are n variables and totally m occurrences of these variables in the system. Then it can be proved, as with the MUS set selection algorithm in [Wang and Vande Vate, 1990], that this labeling procedure takes only O(mlogn) time.

The previous labeling procedure assumes that there are **no cycles** in the Horn system. However, it can be extended to the general situation where cycles are allowed by using the ASSIGN subroutine presented in [Wang and Vande Vate, 1990]. The following section presents an example of how the proposed cost-effective question asking strategy can be used.

An Example

Suppose that a Horn system is comprised of the following three clauses (rules):

IF (A_2 AND A_3) THEN A_1	(clause (1,1) or: $R_{1,1}$)
IF (A_3 AND A_4) THEN A_1	(clause (1,2) or: $R_{1,2}$)
IF (A_5 AND A_6) THEN A_3	(clause (3,1) or: $R_{3,1}$).

Therefore, the variables (assertions) A_2, A_4, A_5, and A_6 are unconfirmed observable variables and A_1, A_3 are non-observable variables. Suppose that the final goal is to determine the value of variable A_1. Let the following table represent the costs and probabilities associated with these unconfirmed observable variables.

Variable A_i	Cost C_i	Probability P_i
A_2	15	0.85
A_4	10	[illegible] 0.75
A_5	20	0.90
A_6	10	0.80

Given all the previous data, the diagram which represents this Horn system and the labels derived by applying the proposed labeling procedure are as in figure 3. From this figure and theorem 3 it can be seen that the cost-effective UOV set to be selected in step 1 is: $\{A_4, A_5, A_6\}$. Its estimated expected cost for this UOV set is: LC(1,2) = 29.5 and its probability label is: LP(1,2) = 0.54.

From theorem 2 it follows that the questions (in step 2) to be asked should be in the order: A_4, A_6, A_5. If the user answers one of these questions with "false", then the system will update the diagram and the labeling procedure will be applied again to derive a new cost-effective UOV set (step 1). This process will continue until no UOV set is left (in which case the final goal cannot be proved) or the user reaches a UOV set with only "true" answers. In the second case the final goal has "true" value.

Concluding Remarks

A successful question asking strategy guides the inference process in proving the final goal by asking as few questions as possible. When the total questioning cost is considered, then such a strategy should reach the final conclusion by minimizing the total cost. Two main steps can be identified in a question asking strategy; form a set of candidate questions (step 1) and select a question to ask next from the previous set (step 2).

In this paper we provided efficient rules for selecting the next question to ask from a set of candidate questions. This is done for two cases. In the first case the objective is to minimize the total number of questions. In the second case the objective is to minimize the total cost in answering these questions.

Furthermore, this paper also presents ways for determining the set of candidate questions (step 1). This is also done for the two cases mentioned above. The approaches proposed in this paper are based on two labeling algorithms and take log-linear time. These approaches are heuristics which attempt to approximate the optimal strategies. In the case of minimizing the total number of questions, experimental results indicate that the proposed approach is very effective. Since Horn clause systems are very popular in many practical applications, more research in this area is needed.

REFERENCES

de Kleer, J. and B.C. Williams, 1987. *"Diagnosing Multiple Faults,"* Artificial Intelligence, 32, pp. 97-130.

de Kleer, J., 1990. *"Using Crude Probability Estimates to Guide Diagnosis,"* Artificial Intelligence, 45, pp. 381-391.

Dowling, W.F. and J.H. Gallier, 1984. *"Linear Time Algorithms for Testing the Satisfiability of Horn Formulae"*, Journal of Logic Programming 3, pp. 267-284.

Duda, R., J. Gasching and P. Hart, 1979. *"Model Design in the PROSPECT Consultant System for Mineral Exploration"*, Expert Systems in the Micro-electronic Age, Edinburgh Univ. Press, Great Britain.

Hayes-Roth, F., D.A. Waterman and D.B. Lenat, 1983. *"Building Expert Systems"*, Addison-Wesley Publishing Company, Inc., Reading, MA.

Hooker, J.N., 1988a, *"Generalized Resolution and Cutting Planes"*, Annals of Operations Research, R.G. Jeroslow (Ed.), Vol. 12, No. 1-4, pp. 217-239.

Hooker, J.N., 1988b, *"Resolution vs. Cutting Plane Solution of Inference Problems: Some Computational Experience"*, Operations Research Letters, Vol. 7, No 1, pp. 1-7.

Jeroslow, R.G., 1988, *"Computation-Oriented Reductions of Predicate to Propositional Logic"*, Decision Support Systems, North-Holland, No. 4, pp. 183-197.

Jeroslow, R.G. and J. Wang, 1989. *"Dynamic Programming, Integral Polyhedra and Horn Clause Knowledge Bases"*, ORSA's Journal on Computing, Vol. 1, Winter 1989, pp. 7-19.

Mellish, C.S., 1985. *"Generalized Alpha-Beta Pruning as a Guide to Expert System Question Selection"*, Expert System 85, Proceedings of the Fifth Technical Conference of the British Computer Society Specialist Group on Expert Systems, Univ. of Warwick, 17-19 Dec. 1985. Cambridge Univ. Press, Cambridge CB2 1RP, pp. 31-41.

Triantaphyllou, E. and J. Wang, 1991. *"The Problem of Asking the Minimum Number of Questions in Expert Systems with Horn Clause Knowledge Bases"*, Working Paper, Kansas State University, Department of Industrial Engineering.

Wang, J., 1990. *"Rule-Based Expert Systems and Discrete Optimizations"*, Doctoral Dissertation. College of Management, Georgia Institute of Technology.

Wang, J. and J. Vande Vate, 1990. *"Question-Asking Strategies for Horn Clause Systems"*, Annals of Mathematics and Artificial Intelligence, 1, 1990, pp. 359-370.

APPLIED COMPUTING: TECHNOLOGICAL CHALLENGES OF THE 1990'S

PROCEEDINGS OF THE 1992 ACM/SIGAPP SYMPOSIUM ON APPLIED COMPUTING

VOLUME I

Kansas City Convention Center
March 1-3, 1992

editors:

Hal Berghel
Ed Deaton
George Hedrick
David Roach
Roger Wainwright

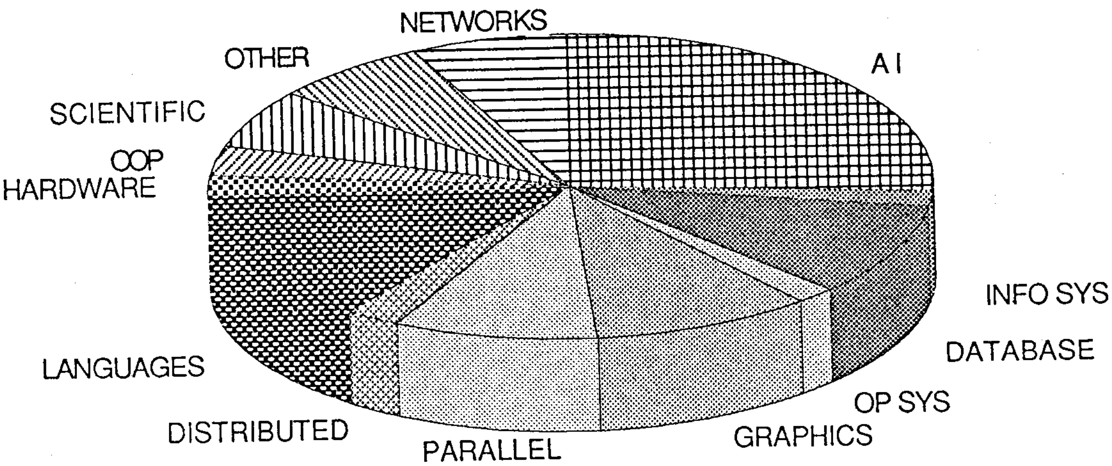

acm PRESS