

REASONING WITH INCOMPLETE KNOWLEDGE IN A RESOURCE-LIMITED ENVIRONMENT:
INTEGRATING REASONING AND KNOWLEDGE ACQUISITION

Mark S. Fox

The Robotics Institute
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

ABSTRACT

This paper describes an approach to reasoning with incomplete information in a resource-limited environment. Approaches to date either assume infinite resources and proceed to enumerate a large inference space, or assume few resources and ignore the missing information. They do not reason about resource constraints and the inference methods admissible under them. A hearsay-II-like system is described where each knowledge source is a separate production system. During rule evaluation, a rule antecedent is evaluated using minimal - resource methods. A rule antecedent is evaluated to true, false, or an expected resource cost to acquire the information necessary to complete its evaluation. If conflict resolution chooses a partially evaluated rule, it posts a goal asking other knowledge sources to provide the missing information, suspends the knowledge source, and informs the knowledge source's manager about the suspension and accompanying goal. The manager decides whether the goal is worth pursuing now, the amount of resources to apply to the task, what knowledge source to apply, and when to give up. The knowledge sources that attempt the goal can implement a variety of inferential and knowledge acquisition techniques.

1. Introduction

The purpose of this paper is to describe a method of reasoning with incomplete information in rule-based systems. In particular it describes how an "expert program" can gracefully integrate the dynamic acquisition of knowledge with its reasoning process. The integration process allows the system to reason about the most appropriate methods for acquiring knowledge.

The primary representation of knowledge in rule-based systems is the production rule (Newell & Simon, 1972; Davis & King, 1975). A production rule is of the form:

Antecedent --> Consequent

The antecedent can be any expression which if true upon evaluation results in the evaluation (execution) of the consequent. The antecedent is classically a pattern to be matched against data in memory, but can be as complex as a first-order predicate expression or an arbitrary LISP expression. A basic assumption in rule-based systems is that the antecedent is a total-function which evaluates to true or false. In a rule system where the antecedent is a pattern to be matched, the

assumption holds; the pattern matches something in memory or it does not. But the failure of a match can arise from two sources, either the pattern does not exist in memory (is absolutely or temporally false), or memory lacks sufficient information to determine truth or falsity, i.e., the knowledge base is incomplete. In the latter case, an unsuccessful match should be interpreted as "I don't know" instead of false. While not new, this problem has many names such as partial-matching (Hayes-Roth, 1978), partial- information inferencing (Joshi, 1978), and backward chaining (Shortliffe, 1975). More generally, it is the problem of reasoning with incomplete information.

In a rule-based system the control system evaluates each rule antecedent and chooses one rule from those that are true. But it is often the case that pattern antecedents partially match memory, and for predicate antecedents a subset of predicates evaluate to true while others are not disprovable. The question is: "What rule should be chosen to execute?" Previous solutions to this problem fall into a continuum where one end uses maximal knowledge but requires large amounts of resources to learn, infer or acquire the missing information, while the other end uses little knowledge and minimal resources. In particular, the backward chaining of Mycin (Shortliffe, 1976) does a complete search of the rule space looking for rules which may provide the missing information required to evaluate an antecedent. Klahr (1978) reduces the search cost of Mycin's backward chaining by utilizing rule abstractions to plan the backward chaining of rules. Abstraction reduces the information available. In both these approaches, the rules determine what information is relevant and the method of acquisition. Joshi (1976) decreases the cost of choosing a rule even further by not searching for the missing information. Instead, a partial match metric based on specificity and certainty is derived and the highest valued rule is chosen. McDermott (1978) deals with partial-matches by creating a

separate rule for each sub-pattern of the antecedent which is of use to the problem solving. Rosenberg (1979) augments the Planner system with meta-theorems that propose alternate ways of supporting a goal when existing theorems fail. Mostow (1977) has proposed the specification of how, what, why and when information for each rule, to guide in the selection of rules and the gathering of information of discrimination amongst them. In all these approaches, the more resources expended, the more

complete the search, the more certain is the outcome. The important questions here are: what is the level of certainty required by the problem-solving, and what is the (resource) cost of reducing uncertainty?

2. Reasoning with Incomplete Knowledge in a Resource-Limited Environment

Reasoning, problem-solving, and learning have all been characterized as heuristic search. In the more interesting domains, the search space is so large that it cannot be completely searched, hence many of the problems center around reducing the size of the search space. Not being able to circumnavigate the search space is basically a resource limitation, e.g., time and space. Under such circumstances, reasoning programs must efficiently assign and use resources to solve a problem. Hence, rule-based systems when faced with the inability to evaluate an antecedent, must balance adequacy of performance against resource limitations. Lenat et al. (1979) have discussed the problem of reasoning in resource limited environments under the name of cognitive economy, and Lenat (1976) has used the schema of assigning time slices to rules to limit resource expenditures.

For example, if there exists a rule of the following form:

IF (X ∈ Vertebrate) THEN

and X in this case is a lion, then there are a number of ways of evaluating the antecedent "X is a member of Vertebrata":

1. Is X equal to Vertebrata?
2. Search the is a hierarchy of the knowledge base.
Lion \Rightarrow Panthera \Rightarrow Felidae \Rightarrow Carnivore \Rightarrow Mamalia \Rightarrow Vertebrata.
3. Look for other inference rules that may provide the answer.
4. Match the Lion schema against the Vertebrata schema.
5. Ask the human sitting at the terminal whether it is true.
6. Assume it's true and retract if it leads to inconsistent results.

This example raises a variety of problems that can be found in almost any domain. There are many algorithms (methods) available for answering a question including learning and knowledge acquisition methods. Some require only a few resources such as the equality check, some require potentially a lot of resources such as the search and the pattern match, while others require nonreplenishable resources. For example, if a system can ask the user questions, then it is in the best interests of the system not to ask the user too many questions because he may become annoyed and provide fewer and less accurate answers.

When faced with one or more rules whose antecedent cannot be evaluated, the reasoning system must address the following problem*:

- Should it continue to work at evaluating the antecedent?
- How many resources should be expended in evaluation?
- What methods should be chosen to evaluate the antecedent?

3. Meta-Evaluation: Resource Limited Evaluation

The KS system of Fox (1981a) is a discovery system which reasons with incomplete information. A simplified view of its program architecture is a distributed set of Hearsay-IMike systems (Erman et al., 1980). That is, there are *departments* where each department has *employee* knowledge sources and a *manager* knowledge source. Departments can form hierarchies or heterarchies with communication channels and shared data spaces. Each knowledge source is a rule-based system with access to a global knowledge base. When a knowledge source is executing, each of its *visible* rules are *meta evaluated*. Meta-evaluation attempts to evaluate an antecedent using minimal resources. If it cannot evaluate the antecedent successfully, it assigns to the rule an *expected cost* for successfully evaluating the rule using other methods. The conflict resolution algorithm for the knowledge source then chooses a rule based on specificity, relation to the knowledge source's goal, number of parts evaluated, and expected cost (if not successfully evaluated). If the antecedent of the rule chosen was evaluated successfully, then its consequent is evaluated and the process begins again. If the antecedent was assigned an expected cost, then the knowledge source is suspended and the portions of the rule not evaluated are placed as goals in the knowledge base for other knowledge sources to evaluate. Other knowledge sources' preconditions fire and instantiate them. Each knowledge source provides different methods at different costs. The manager knowledge source makes the decision as to which knowledge source to assign the task, and how many resources it will expend. A knowledge source may *run* out of resources before it finishes, leaving it to the manager to assign more resources or to assign the task (goal) to another employee. Once the antecedent is finally evaluated, the suspended knowledge source is awakened and passed the value by the manager. Processing of the knowledge source continues.

4. Meta-Evaluation: An Example

A rule antecedent in the KS system can be any arbitrary lisp expression but includes existential and universal quantification as a basic knowledge base search mechanism. To simplify matters, we will assume a single department with a manager and three employees:

1. CONSPEC: knowledge source that creates new schemata via specialization (figure 4-1)¹.
2. INTROSPEC: primary inference knowledge source.
3. OUESTASK: knowledge source that acquires knowledge by asking the user a question.

The schema depicted in figure 4-1 has the following form: A schema definition is surrounded by double braces "{{ <schema def> }}" . A schema is composed of multiple views, each which is surrounded by single braces "{ <view type> <viewed schema> <view slots> }" Each view has a set of slots defined (inherited) in that view Each SLOT, which is printed as SMALL CAPITALS, has a set of facets printed in *italics* One of the facets is a the *value* of the SLOT When only the *value* facet has a filler, the SLOT and the filler are printed without naming the facet A schema also has slots that are independent of its views They are printed separately. Each slot can act as a bidirectional inheritance relation. Slots that act as inheritance relations are printed in the view format. See (Fox, 1981 a. 1961b) for more information.

```

{{Conspec
  {is A Employee
    MANAGER: Example-Manager
    WINDOW:
      Restriction: (SET (TYPE isa rule))
      Value: RI R2
    STATE:
      Restriction: (OR rule block mail-block completed
        failed ready)
      Value: ready
    PRECONDITION:
      restriction: (SET (TYPE is a capability))
    BUDGET:
    SPACES:
    MAIL BOX:
    SUSPENDED RULE: }
  })

```

Figure 4-1: CONSPEC Schema

CONSPEC has two rules in its rule set (figure 4-2). Rule RVs antecedent tests whether X is a member of VERTEBRATA. Rule R2 is similar to R1 but has been specialized by the added condition "similar" which tests whether X's color is similar to the color of X's environment. For this example, X is bound to the schema LION.

```

{{R1
  { is a Rule
    ANTECEDENT: (X € VERTEBRATA)
    CONSEQUENT: ... }
  }}

{{R2
  {is A Rule
    ANTECEDENT: (X € VERTEBRATA) AND
      (X.Color IS SIMILARTO (X.Environment).Color)
    CONSEQUENT: ...}
  }}

```

Figure 4-2: CONSPEC Visible Rules

4.1. Low-Cost Evaluation

The first phase of rule metaevaluation is low-cost evaluation of functions in rule antecedents. For each function in the antecedent, there exists a schema in the knowledge base that defines how it is evaluated. These function schemata come in three flavors:

1. *Low-cost Function*: uses bounded, small set of resources during execution. Contains the code to be executed.
2. *Vnrtable>-cosi Function*: uses variable sot of resources (possibly unbounded) in evaluation. Search based functions (e.g., there-exists) are of this type. This schema contains the expected-cost of evaluation information and the goal name to be assigned to the function.
3. *Multifunction*: a function that can be evaluated in more than one way, some low-cost and others variable-cost. An example is set membership €. The schema contains the low-cost functions and the expectedcost and goal name of the variable-cost approach.

If the function is low-cost then the code in the schema is executed. If the function is variable-cost, then the expectedcost for evaluation and the name of the goal to attach to the function is returned. If the function is a multi function then the low-cost functions are evaluated until one succeeds. If none succeed, then the expected cost and goal name is returned.

We will assume that in this example, there is enough information to answer rule R2s color test, but none of the low-cost function tests for the multi function E succeed (figure 4-3). The first low-cost function €=equal uses equality to test membership, but LION = VERTEBRATA SO it fails. The second low-cost function €depth3, searches to depth 3 up the ISA hierarchy (Lion ↔ Panthera ↔ Felidae → Carnivora), and also fails. Hence, the expected cost of (seconds 30) is returned for € along

```

{is A Multi-Function
  LOW COST FUNCTION: C equal
    €-depth3
  EXPECTED COST: (seconds30)
  GOAL-NAME: € goal}
}}

```

Figure 4-3: € Multi-Function

with the CGOAL name to be assigned to the antecedent if the rule is chosen by conflict resolution.

4.2. Conflict Resolution

The rules R1 and R2 cannot be evaluated since lion and vertebrata are more than 3 levels apart; thus for each rule, the

expectedcost of (seconds 30) is returned. The conflict resolution system looks at the two rules and chooses R2 for three reasons: 1) R2 is more specific than R1, 2) answering R2 will also answer R1, and 3) R2 has more subparts successfully matched than R1. Coat, specificity, evaluability, and relation to employees goal are used to choose a rule.

4.3. Goal Posting and Assignment

Each department has associated with it one or more sections of the Knowledge base called spaces. Each employees precondition is keyed to one or more of these spaces. If the chosen rule has an expected cost associated with it. the knowledge source is suspended (figure 4 5) and the unevaluated functions are posted in the department's goal space as schemata with the goal-names assigned and a state attribute of POSTED (figure 4 6). All the other "interested" employees then test the goal against their precondition (apply the contents of the test slot

```

{{ Lion
  { ISA Panthers
    COLOR: tan
    ENVIRONMENT: Grassland}
  }}
{{ Grassland
  {is-A Place
    COLOR: brown, green, tan }
  }}

```

Figure 4-4: Lion and Grassland Schemata in Memory

in each of the capability schemata that (ill the procondition slot) and if true inform the manager of their appl cability In tins case, both the INTROSPEC and QUESTASK employees lire and inform the manager that they can (possibly) evaluate the € GOAL. The manager then compares the possible approaches offered by the employees, INTROSPEC (figure 4-7) would use more resources than OUESTASK (figure 4 8) but the factor of bothering the human user is considered, lowering the overall utility of OUESTASK. INTROSPEC is chosen, assigned resources based on availability and expected cost and what are available, and is initiated. The state of the E.GOAL schema in the goal space is changed from POSTED to ASSIGNED (figure 4-9).

```

{{ Conspec
  {is A Employee
    MANAGER. Example-Manager
    WINDOW: R1 R2
    STATE: suspended
    SUSPENDED-RULE: R2
    GOAL: €-goal-1 }
  }}

```

Figure 4-5: Suspended CONSPEC Schema

```

{{ C-goal-1
  { IS-A Goal
    GOAL: €-GOAL
    PARAMETERS: Lion Vertebrata
    SUSPENDED-KS: Conspec
    RULE: R2
    STATE: posted
    EXPECTED-COST: (seconds 30)
    RESULT:
    CONTRACTED-KS:
    BUDGET: (seconds 50)
    ATTEMPTS: }
  }}

```

Figure 4-6: Posted €-GOAL

4.4. Goal Satisfaction

The INTROSPEC employee tries a variety of methods available in its rule set to solve the goal. For each attempt that fails, it records on the goal schema the attempt type and the amount and type of resources consumed (figure 4 10). This information is used by the system to update function schemata expected-costs and by the manager to decide what to do next if the employee fails at its task. If INTROSPEC succeeds, it changes the state of the goal

```

{{Introspec
  {is A Employee
    MANAGER Example-Manager
    STATE: ready
    PRECONDITION; {is A Capability
      GOAL: € GOAL
      TEST; C-GOAL TEST INTROSPEC
      APPROACH; €- Match-Rule Set
      COST: (seconds 25)}
  }}

```

Figure 4*7: Introspec Schema

schema to EVALUATED, informs the manager, and ends. The manager activates (puts in the ready queue) the suspended employee. If INTROSPEC fails or runs out of resources, the manager is informed and decides whether to continue persuing the goal, what resources to expend, and to whom they should be assigned. In this case, if INTROSPEC fails, OUESTASK is initiated and the user is queried as to whether a lion is a vertebrata.

```

{{ Ouestauk
  {is A Employee
    MANAGKR: Example-Manager
    STATE: ready
    PRECONDITION: {is A Capability
      GOAL: C-GOAL
      TEST: CGOAL TESTOUESTASK
      APPROACH: €-question
      COST: (questions 5)}
  }
}

```

Figure 4-8: Questask Schema

```

-Attempt-1
{is A Goal-Attempt
  KS: Introspec
  METHOD: € Match-Rule Set
  COST, (seconds 20)
  REASON. <list of unmatched attributes> }

```

Figure 4-10: C-Goal Attempt by Introspec Knowledge Source

5. Observations and Conclusions

We described an approach for reasoning with incomplete knowledge in a rule-based system which integrates reasoning and knowledge acquisition. Simply, it attempts to evaluate a rule's antecedent using resource-miserly methods. If these fail, then it posts the antecedent as a goal for evaluation by other knowledge sources in the system. This approach has several strengths:

1. antecedents are quickly evaluated when the

```

{{ c-goal-l
  {is A Goal
    GOAL: E GOAL
    SUSPENDED KS: Conspec
    RULE: R2
    STATE: assigned
    EXPECTED COST: (seconds 30)
    RESULT:
    CONTRACTED KS: Introspec
    BUDGET: (seconds 50)
    ATTEMPTS: }
  }
}

```

Figure 4-9: Assigned C-GOAL

information to answer it exists. The cost of meta-evaluation at this point is the cost of interpreting the function's schema. Small for a LOW-COST FUNCTION, more expensive for a MULTIFUNCTION.

2. When information is incomplete a variety of methods can be brought to bear on evaluating the antecedent. These methods require various amounts of resources.
- 3 The choice of method to apply is left to the manager knowledge source which has a more global view of

the problem, hence can intelligently choose the best approach and the resources to assign².

4. Even if an antecedent cannot be evaluated, the manager or an employee can force it to be true or false. In essence, creating an hypothesis, requiring dependency analysis capabilities (London, 1978; Doyle, 1979).
5. The manager can postpone evaluating an antecedent to work on more important problems.

The problem of *attention focussing* (Hayes-Roth & Lesser. 1977) relies on good heuristics for choosing amongst alternatives, and a language for describing Knowledge source capabilities (Lesser, 1967; Fox, 1979b).

While this ;ippin.ich provides greater control of the application of inference and knowledge acquisition methods in filling in missing information required by reasoning, it incurs two types of expenses. The first, is the added cost of rule meta-evaluation. For each function in the antecedent, the corresponding schema must be accessed and interpreted according to its type: low-cost, variable, or multi. Hence, there is a per function constant overhead. In the KS system, the cost of interpretation has been removed. Each functional schema is both a declarative and procedural representation that can be executed directly without interpretive overhead (Fox, 1979a; 1981a). Hence, meta-evaluation costs are negligible.

The second expense is the increased complexity of the manager's rules to handle task selection, employee selection and resource assignment. Research in perceptual tasks such as speech vision (e.g., Hearsay-II) has shown the need for such capabilities when dealing with complex and uncertain data.

6. Acknowledgements

I would like to thank Herb Simon, John McDermott and Jaimie Carbonell for their criticisms of earlier versions of this paper.

7. References

- Davit R. and J.King, (1975), "An Overview of Production Systems," Technical Report STAN CS-75 524, Memo AIM-271, Stanford CA: Stanford University, Computer Science Department.
- Doyle J-, (1979), "A Truth Maintenance System", Artificial Intelligence, Vol. 12, No. 3
- Erman L.D., F, Hayes-Roth, V.R. Lesser, and D.R. Reddy, (1960), "The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty", Computing Surveys, Vol. 12, No. 2, June 1960, pp. 213-253.
- Fox M.S., (1976), "Knowledge Structuring: An Overview", Proc. of the 2nd Conf. of the Canadian Society for Computational Studies of Intelligence, Toronto, Ontario, pp. 146*155.
- Fox M.S., (1979a), "On Inheritance in Knowledge Representation", Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokyo Japan.
- Fox M.S., (1979b), "Organization Structuring: Designing Large, Complex Software", Technical Report CMU-CS-79-155, Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA.
- Fox M.S., (1981a), "Knowledge Structuring: Knowledge Accomodation and Discovery Using Specialization and Planning", (Ph.D. Thesis), Computer Science Department, Carnegie Mellon University, Pittsburgh PA, In preparation.
- Fox M.S., (1961b), "SRL: Schema Representation Language", Technical Report, The Robotics Institute, Carnegie-Mellon University, Pittsburgh. PA, in preparation.
- Hayes-Roth F., (1977), The role of partial and best matches in knowledge systems, In Waterman and Hayes-Roth (Eds.), Pattern directed infaranoa systems. N.Y.: Academic press.
- Hayes Roth F., and V. Lesser, (1977), "Focus of Attention in the HEARSAY II Speech Understand System," Fifth Int. Joint Conf. on Artificial Intelligence, Cambridge, MA, Aug. 1977.
- Joehi, Aravind K., (1978), "Some Extensions of a System for inference on Partial Information", in Pattern-Directed Infaranoa Systems, D. Waterman and F. Hayes-Roth (Ed.), NY: Academic Press.
- Klahr, PhiHp, (1978), "Planning Techniques for Rule Selection in Deductive Question-Answering", In Pattern-Directed Infaranoa Systems, D. Waterman and F. Heyee-Roth (Ed.), NY: Academic Press.
- Lenat, D.B., (1976), "AM: An Artificial Intelligence Approach to Discoverey in Mathematics us Heuristic Search," (Pn.D. Thesis) Computer Science Dapl., Stanford University
- Lenat D.B., F.Hayes-Roth, and P.Klahr, (1979), "Cognitive Economy in Artificial Intelligence Systems", Proceedings of the SixJh Annual Conference on Artificial Intelligence, pp. 531-6, Tokyo
- London P.E., (1976), "Dependency Networks as a Representation for Modelling in General Problem Solvers", Ph.D. Thesis, Technical Report TR-698, NSG-7253, , University of Maryland. College Park MD, Sept. 1978.
- McDermott J., (1979), "ANA: An Assimilating and Accomodating Production System", Technical Report, Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA.
- Moatow D.J., (1977), "A Revised Architecture for Rule-based Systems", Sigart Newsletter, No. 63, June 1977. pp. 856.
- Newell A., and H.A. Simon, (1972), Human Problem Solving, Englewood CMffe, N.J.: Prentice Halt.
- Rosenberg S., (1979), "Reasoning in Incomplete Domains", Proceedings of the Sixth International Joint Conicrenco on Aititirial Intvlligonce, Tokyo Japan, pp. 73t-737.
- Shorlliffo E.H.. (19/G), Computet Usutd Medical Consultations: MYCIN, New York: American Elsevier.

This research was sponsored by the Defense Advanced Research Protects Agency (DOD), ARPA Order No. 3697, monitored by the Air Force Avionica Laboratory under Contract F33615-76-C-1661.

The views snd conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.