

# Computational Complexity of Hypothesis Assembly

Dean Allemang, Michael C. Tanner, Tom Bylander, John Josephson

Laboratory for Artificial Intelligence Research  
Department of Computer Science  
The Ohio State University

## ABSTRACT

The problem of finding a best explanation of a set of data has been a topic of much interest in Artificial Intelligence. In this paper we present an approach to this problem by hypothesis assembly. We present this approach formally so that we can examine the time complexity and correctness of the algorithms. We then examine a system implemented using this approach, which performs red blood antibody identification. We use this domain to examine the ramifications of the assumptions of the formal model in a real world situation. We also briefly compare this approach to other assembly approaches in terms of time complexity and reliance on assumptions.

## I. Introduction

The problem of abductive reasoning (as proposed by the philosopher C.S. Peirce) has been a topic of much recent interest in Artificial Intelligence (Miller 1982, Reggia 1983, Charniak 1985). The general task faced by an abductive reasoning system is to find the best explanation of a set of data or observations, i.e., the best way to account for a set of data. Most of the work in Artificial Intelligence in this area has focused on a specific kind of abduction, which we call hypothesis assembly. The hypothesis assembly task assumes as given a set of hypotheses with some knowledge about what sorts of data each can account for, and finds the subset of these hypotheses that best accounts for the problem data. At the Ohio State Laboratory for Artificial Intelligence, Josephson et. al. (1985) have been developing an approach for abduction based upon hypothesis assembly.

In this paper we will begin by presenting a mathematical idealization of this approach. From this we will analyze the complexity and correctness of the algorithm. Then we will examine how well this idealization matches with real

This work has been supported by the National Library of Medicine under grant LM-04298, the National Science Foundation through a Graduate Fellowship, and the Defense Advanced Research Projects Agency, RADC contract F30602-85-C-0010. Computer facilities were enhanced through gifts from Xerox Corporation.

world concepts of abduction. In particular, we will examine a system called RED, based upon this approach, which performs antibody identification in the domain of red blood cell typing, as described in (Josephson 1984), (Josephson 1985) and (Smith 1985), and show how the general mathematical results respond to questions that have been raised (Mostow 1985) about its complexity.

## II. Mathematical Idealization of Abduction

### Definitions

In order to motivate the following definitions, we will examine briefly the domain of the current implementation of RED, the domain of blood bank antibody analysis. The primary data consists of results of several lab tests on blood samples. The blood bank technologist knows how *antibodies* can account for various reactions. The lab tests have the property that if antibody A accounts for some reaction  $r$ , and antibody B accounts for reaction  $q$ , then the presence of both antibodies A and B accounts for both reactions  $r$  and  $q$ . This property of a domain will be referred to as *independence* of hypotheses.

More formally, we define a domain for hypothesis assembly as a the triple  $(H, M, e)$ , where  $H$  is a finite set of *hypotheses*,  $M$  is a finite set of *manifestations*, and  $e$  is a map from subsets of  $H$  to subsets of  $M$ .  $e(S)$  is interpreted as the *explanatory power* of a set of hypotheses, and is the set of manifestations for which those hypotheses can account. An assembly problem is specified by a subset  $M_0 \subseteq M$ .  $M_0$  is interpreted as the set of observed manifestations<sup>1</sup>. In these terms, we have

*The Independence Assumption:*

If  $S$  and  $T$  are subsets of  $H$ , then

$$e(S \cup T) = e(S) \cup e(T)$$

Although many domains satisfy the independence assumption, we wish to strengthen our result by replacing

---

<sup>1</sup> In what follows, we will use the notation  $e(S)$  where we should, strictly speaking, write  $e(S) \cap M_0$ , for the restriction of  $e(S)$  to the observed manifestations.

the independence assumption with two assumptions, which when taken together are weaker than the independence assumption.

*The Monotonicity Assumption:*

If  $S$  and  $T$  are subsets of  $H$ ,

$$S \subseteq T \Rightarrow e(S) \subseteq e(T)$$

Clearly any domain that satisfies the independence assumption also satisfies the monotonicity assumption.

*The Accountability Assumption*

For a hypothesis  $h \in H$ , it is possible to know what manifestations  $m$  can account for. Formally, the function  $\alpha : H \rightarrow M^*$

$$\alpha(h) = \left\{ m \in M \mid \exists S \subseteq H \text{ with } h \in S \text{ such that } \begin{array}{l} m \in e(S), m \notin e(S \setminus h) \end{array} \right\}$$

can be computed as easily as  $e$  (polynomial calls to  $e$ ).

Notice that if the independence assumption holds, then  $\alpha(h) = e(h)$ , so the accountability assumption holds.

Central to the success of any hypothesis assembly algorithm is the following assumption:

*The Computability Assumption:*

For any subset  $S$  of  $H$ ,  $e(S)$  can be computed.

Clearly the complexity of this computation is central to the complexity of the assembly task, since it is difficult to ask for a set of hypotheses that accounts for some set of manifestations if it is difficult to compute what is accounted for by a set of hypotheses. In cases in which independence holds, this computation can be done in linear time by simply having a table that tabulates  $e$  for each individual hypothesis, so that  $e(S)$  can be computed by taking the union of  $e(s)$  for  $s \in S$ . This has allowed the current RED mechanism, the set covering algorithm of Reggia (Reggia 1983), and Internist (Miller 1982) to safely ignore the complexity of the computation of  $e$ . In fact, the incremental nature of all these algorithms allows them to compute  $e(S)$  in constant time, since the usual case is that some set  $T$  is to be extended with another hypothesis  $s$ , with  $e(T)$  already known, so that the computation of  $e(T \cup \{s\})$  can be done by simply looking up  $e(\{s\})$ . For the complexity analysis which follows, we will assume an oracle for the function  $e$ , that is, we make the calculations counting one step for each evaluation of  $e$ .

B. Description of the hypothesis assembly algorithm

We now describe the algorithm for abduction by hypothesis assembly. The algorithm has four phases; screen-

ing, hypothesis collection, parsimony and critique.

The input to the algorithm is the set of observed manifestations,  $M_0$ . The output is a set of hypotheses that constitute a 'best' explanation of the data.

The screening phase determines *plausibilities* (i.e., *prima facie* likelihoods) for the veracity of the hypotheses. Hypotheses with very low plausibilities are ruled out for further consideration. In this discussion we wish to examine and analyze the assembly phases of the algorithm, so we will omit a description of the screening phase. When we discuss applications, we will see one way to accomplish screening.

The collection phase works in the following way:

Let  $\langle H, M, e \rangle$  be a domain for hypothesis assembly; let  $M_0 \subseteq M$  be a problem in that domain.

$M'$ ,  $H'$  and  $m$  are variables in the following algorithm.

Let  $M' = M_0$  be the observed manifestations (the manifestations to be explained).

Let  $H'$  be the empty set (working set of hypotheses).

Until  $M'$  is empty, do the following:

Let  $m$  be the most salient hypothesis in  $M'$ .

Find the most plausible hypothesis  $h$  that offers to account for  $m$ , i.e.,  $m \in \alpha(h)$ . Adjoin  $h$  to  $H'$ .

Compute the reactions accounted for by  $H'$ .

Let  $M'$  be  $M' \setminus e(H')$ , i.e.,  $M'$  is the set of manifestations yet to be explained.

The set  $H'$  is the assembled explanation of the observed manifestations.

The parsimony phase works in the following way:

For each hypothesis  $h$  in  $H'$  (the working hypothesis set), starting from the least plausible to the most plausible, do the following:

If  $e(H' \setminus \{h\}) \supseteq M_0$  then  $h$  is superfluous, so set  $H' := H' \setminus \{h\}$ .

$H'$  is now a parsimonious set of hypotheses, that is, it has no redundant parts. Note that this does not guarantee that  $H'$  is a complete explanation of least cardinality.

The critique phase determines whether a hypothesis is *essential* to the explanation, that is, there is no way to account for the observed data without using that hypothesis. This phase works as follows:

For each hypothesis  $h$  in  $H'$  do the following:

Set  $G := H' \setminus \{h\}$ . Repeat the collection algorithm (above) using  $G$  as working hypothesis, with hypothesis  $h$  marked as unusable.

If collection completes (that is, there is a composite hypothesis that does not use  $h$  which explains  $A/O$ ), then  $h$  is not essential.

If collection fails to complete, then hypothesis  $h$  is essential.

### C, Complexity of the Assembly mechanism

We begin by analyzing the time complexity of the assembly phase.

The collection algorithm can be analyzed as follows:

Let  $h$  be the number of hypotheses that were not ruled out; let  $m$  be the number of manifestations observed.

For each manifestation, we select a hypothesis that offers to explain that manifestation, and evaluate the resulting composite. There are  $m$  such manifestations, and selection of a hypothesis takes at most  $h$  calls to the oracle  $a$  (for a linear search). The computation of  $M^*e(H')$  can be done in  $O(m \log m)$  steps (sort both sets, and merge the results). Thus, the complexity of this step is  $O(m(h + m \log m))$ .

Each time a new hypothesis is added to the working composite hypothesis  $H'$ , it is necessary to evaluate  $e(H')$ , so there are  $O(m)$  calls to the oracle  $e$ .

Next we analyze the parsimony algorithm:

At most  $h$  hypotheses are considered for redundancy, so there are  $O(h)$  calls to the oracle  $e$ . At each such step, we must compare the explained set of manifestations to the observed set, that can be done in time  $O(m \log m)$  (first sort both sets, then compare elementwise), so the entire complexity is  $O(hm \log m)$ .

Notice that this algorithm does not guarantee that we have found a compound hypothesis of smallest size that accounts for the observed data. This problem, which is the task of the Reggia et. al. set covering program (Reggia 1983), can be shown to be NP-complete (reduction to vertex cover). This algorithm simply guarantees that the parsimonious solution has no proper subset that is also a solution.

Now we analyze the critic for essential hypotheses:

The critic takes each of the hypotheses in the working hypothesis and marks it as unusable, then repeats the assembly algorithm. There are at most  $h$  hypotheses in the composite hypothesis, so this is at most  $h$  times the complexity of assembly, that is,  $O(hm(h + m \log m))$ .

### D. Correctness of the algorithm

That the collection mechanism produces a composite hypothesis that accounts for all of the observed findings (provided, of course, that such a collection exists) is obvious from the loop condition. It is also clear that the critique

algorithm will find all indispensable hypotheses, and all hypotheses deemed to be essential are in fact indispensable. The parsimony algorithm needs some justification.

We say that a composite hypothesis is a *parsimonious* explanation of some findings if no proper subset of it can account for those findings. The parsimony mechanism ensures that there is no single hypothesis that can be excluded from the composite without losing the capability of explaining some finding. It is possible to construct a domain in which a composite hypothesis has a proper subset that can account for all findings (that is, the composite hypothesis is not parsimonious), yet no single hypothesis is by itself redundant (that is, there is no hypothesis such that it alone can be removed without losing explanatory power). If the monotonicity assumption is true, this can not happen.

Suppose that  $S$  is a non-parsimonious composite hypothesis. Then for some subset  $R$  of  $S$ ,

$$e(S) = e(S \setminus R)$$

Now, for any  $r \in R$ , we have

$$S \setminus R \subseteq S \setminus \{r\} \subseteq S$$

So by monotonicity, we have

$$e(S) = e(S \setminus R) \subseteq e(S \setminus \{r\}) \subseteq e(S)$$

i.e.,

$$e(S) = e(S \setminus \{r\})$$

So for all  $r \in R$ ,  $r$  is itself redundant, that is, the RED mechanism will find a redundant hypothesis in  $S$ . Hence we conclude that if  $S$  is not parsimonious, then RED will eliminate some hypothesis from it.

## III. Adapting the formal model to real-world abduction

### A. The Abductive Answer

An abductive problem solver applied to actual problems tries to produce a best explanation for the observed data. Unfortunately, there is no hard and fast rule for demonstrating that a particular explanation is best. Human problem solvers argue about whether one explanation is better than another. Hence the best we can do to argue that the answer given by this algorithm is best is to list the features which recommend it.

The answer given by this approach is guaranteed to cover all the data. Every observation will be explained.

The answer will be parsimonious in the sense that it will have no superfluous parts. It is not necessarily the smallest answer, i.e., the answer with the least parts.

The answer will be plausible, since more plausible hypotheses are preferred over less plausible ones. The algo-

rithm prefers a large number of plausible hypotheses over a smaller number of less plausible ones. The algorithm will identify essential hypotheses, so when part of the answer is provably correct, the algorithm determines this.

Again, we cannot prove that such an answer is a 'best' explanation in any given situation, these are just the points by which one could argue for this answer. In the long run, the best answer is the one that is in fact the true situation. So now let us examine an application of this method to an actual domain.

## B. A Real-World application -

### Red blood antibody identification

#### 1. Description of the domain

One of the jobs done by a blood-bank technologist is to identify antibodies in a patient's serum to antigens that might appear on red blood cells. This is typically done by combining samples of patient serum with red blood cells that contain certain antigens. Some of these combinations will show reactions, others will not. The presence of certain antibodies in the patient serum will account for certain reactions. The reactions are additive in the sense that if the presence of one antibody explains one reaction, and presence of another antibody explains another, then the presence of both antibodies explains both reactions. If both antibodies can account for a weak result in some reaction, then the presence of both can account for a strong result in that reaction. Also, some pairs of antibodies cannot occur together. The task of the pathologist is to decide what antibodies are present, given a certain reaction pattern.

#### 2. Satisfaction of Assumptions

We now examine how well this task meets the assumptions of the formal model presented in the first part of this paper.

*Computability of e* - The pattern of antigens appearing on the test cells is available in table form. This information determines which antibodies account for any of the reactions. More detailed knowledge of the chemistry of the reactions allows us to build a table to tell the strengths of the reactions.

*Independence* - The additive nature of the reactions means that for separate reactions, the independence assumption is met. Since the model does not allow for parts of manifestations to be explained, we cannot really say that the additivity of reaction strengths is included in the independence assumption.

*Mono tonicity* - If we view a weak result for some reaction as a separate result from a strong result for the same reaction, then we can say that the phenomenon of additive reaction strengths satisfies the monotonicity assumption. That is, each antibody alone explains the manifestation of 'weak reaction'. Together, they can explain either a 'weak reaction', or a 'strong reaction'.

*Accountability* - Since the domain nearly satisfies the

independence assumption, it is not surprising that we can calculate a nearly as easily as we can in the independent situation. That is, if an antibody can account for any strength result for some reaction, then it can potentially contribute to an explanation of any other strength result for that reaction.

*Incompatibility of Hypotheses* - Implicit in the formal model is the assumption that any collection of hypotheses is possible. This domain has a restriction that invalidates this assumption. It can be shown that hypothesis collection in general under such a constraint is an NP-complete problem (see appendix). There are two reasons why this is not a problem in this domain. One is that the collection problem is not exponential in the number of hypotheses, or even the number of hypotheses left after screening, but in the number of incompatible pairs remaining after screening. This tends to be quite a small number. Also, since the screening rates hypotheses by plausibility, the algorithm is likely to stumble upon the correct answer early in the collection process.

#### 3. Implementation

The algorithm for this domain is presently implemented as follows:

*Screening* - this phase is done by a hierarchical classification system similar to the MDX system (Chandra 1983).

*Collection* - this phase is done as in the text, with the exception that when a hypothesis is to be entered that is incompatible with some hypothesis already in the collection, the conflict is resolved by throwing out the hypothesis which is already in the set. A check is made for loops each time a hypothesis conflict arises. This is, in the worst case, a most plausible-first, brute force search. The computations of  $e$  and  $a$  are done from a table of singleton values; since the independence assumption holds (almost), we can compute the value of  $e$  of a set from the value of  $c$  at each element.

*Parsimony and Critique* These phases are done as described in the text.

The collection and parsimony phases are repeated, using the essential hypotheses as a starting point.

This algorithm matches the one in the text, except for the problem of incompatible hypotheses. So, with that exception (that was discussed above), the time complexity of the RED algorithm is polynomial in the number of antibodies and reactions.

#### 4. Availability of Knowledge

We have seen that in the blood typing domain that the assumptions required by this model are easily satisfied. We now examine these assumptions to estimate how likely it is that they will extend to yet other domains.

Without the *computability of e*, it would be impossible to evaluate whether an answer constituted a complete

explanation, much less to construct such an answer. Many existing abduction systems (e.g. Set-covering model (Reggia 1983) and Internist (Miller 1982)) depend strongly upon this assumption.

The *independence* assumption makes computations easy, since we only need to know values of  $e$  for singleton sets. Diagnostic domains are likely to satisfy this assumption, since devices are often decomposable into independent parts, and hypotheses about these parts are often independent. The set covering model (Reggia 1983) relies strongly on this assumption. Internist (Miller 1982) relies on this, but it would only take a simple change to Internist to cope with more unruly domains, provided that they satisfy the monotonicity and accountability assumptions.

The *monotonicity* assumption is weaker than the independence assumption, hence is more likely to hold. It allows for some interaction of hypotheses, so long as no manifestations are canceled.

The *accountability* assumption is also weaker than independence, so is more likely to hold. This simply requires that we know to what manifestations a hypothesis can contribute.

In a more unruly domain, in which the monotonicity assumption fails to hold, that is, when several hypotheses together fail to account for data that some of them do separately (e.g., due to subtractive interactions between the hypotheses), all of these assembly techniques will have difficulty. It can be shown that certain types of subtractive interaction problems are NP-complete. In such cases, it may be fruitful to find different knowledge so that the problem is no longer one of assembly. One approach along this line is ABEL, a system described in (PATIL, 1982), which attempts to make this problem tractable by making use of elaborate causal structures.

#### Conclusion

An important characteristic of problems that are dealt with in Artificial Intelligence is that domain knowledge is necessary to make the computations feasible. If an approach is to be widely applicable, it is necessary to know what knowledge is necessary to ensure that the approach will yield correct answers in a reasonable time.

In this paper, we have constructed a formal model of the approach to abductive reasoning illustrated by the RED system. From this model, we have identified when the available knowledge will allow application of this approach. Under the assumption that the knowledge is of this form (that is, it satisfies the assumptions of monotonicity, accountability and computability), we have demonstrated the tractability and correctness of the computation.

## Appendix I Proofs of theorems

### NP-Completeness of the Reggia algorithm

After Reggia (1983) (using the notational conventions of this paper), we have the following definitions:

a *diagnostic problem* is a 4-tuple  $(H, A, C, M_0)$ .

$H, A$  are finite sets of hypotheses and manifestations respectively,

$C$  is a relation,  $C \subseteq H \times M$ ,

$M_0 \subseteq M$ , the 'observed manifestations'.

For a subset  $S \subseteq H$ , we write  $e(S)$  for the set  $\{m \in M \mid \exists h \in S \text{ s.t. } (h, m) \in C\}$

An *Explanation* for a diagnostic problem is a set  $S \subseteq H$  such that  $M_0 \subseteq e(S)$ ,  $|S| \leq |T|$ , for any  $T \subseteq H$  s.t.  $M_0 \subseteq e(T)$ .

We wish to show that given a diagnostic problem, finding an explanation is NP-complete. We give a reduction from minimal node cover.

Minimum node cover: Given a graph  $G = \langle V, E \rangle$ , find a set  $V' \subseteq V$  of smallest cardinality s.t.  $\forall e \in E, E$  has an endpoint in  $V'$ .

We form a diagnostic problem from the graph  $G$  as follows:

Set

$H := V$

$M := E$

$M_0 := E$

$C := \{(v, e) \mid v \text{ is an endpoint of edge } e\}$

Clearly an explanation of this diagnostic problem is also a cover of the graph ( $V'$  of minimum cardinality). Also clearly this transformation is polynomial in the size of  $G$ .

### NP-completeness of hypothesis assembly with Incompatible Hypotheses

An hypothesis assembly problem under constraints of incompatible hypotheses can be specified as a 5-tuple  $(H, A, C, I, A_f)$ , where

$H$  and  $A$  are finite sets of hypotheses and manifestations respectively,

$C$  is a relation,  $C \subseteq H \times M$ , a table for computing  $e$ ,  $I$  is a set of unordered pairs of elements of  $H$ , and  $M_0 \subseteq M$ .

For a subset  $S \subseteq H$ , we write  $e(S)$  for the set  $\{m \in M \mid \exists g \in S \text{ s.t. } (g, m) \in C\}$

The interpretation of these sets is that  $H$  is a set of hypotheses that can account for certain manifestations in  $A$ ,  $(h, m) \in C$  (i.e.,  $m \in e(h)$ ) means that  $h$  can account for  $m$ .

$\{h_1, h_2\} \in I$  means that hypotheses  $h_1$  and  $h_2$  are incompatible. An explanation is a set  $S, S \subseteq H$  such that  $M_0 \subseteq e(S)$ , and such that there is no pair  $h_1, h_2 \in S$  with  $\{h_1, h_2\} \in I$ .

We prove NP-completeness of this problem by reduction from 3SAT (Garey 1979).

3SAT: Given a statement in propositional calculus in disjunctive normal form, in which each term has at most three factors, find an assignment of the variables that makes the statement have value 'True'.

We form an hypothesis assembly problem from a 3SAT problem by mapping factors in a 3SAT problem to hypotheses and mapping terms to manifestations as follows:

Let  $P$  be a statement in propositional calculus in disjunctive normal form, where each term has at most three factors. Let  $\{u_1, u_2, \dots, u_n\}$  be the variables used in  $P$ , let  $t$  be the number of terms in  $P$ .

We now construct an instance of the hypothesis assembly problem under incompatibility constraints with  $2n$  hypotheses and  $t$  manifestations.

$$\text{Set } H := \{h_1, h'_1, h_2, h'_2, \dots, h_n, h'_n\}$$

$$M := \{m_1, m_2, \dots, m_t\}$$

$$C := \{ \{h_i, m_j\} \mid u_i \text{ is a factor in the } j^{\text{th}} \text{ term} \}$$

$$\cup \{ \{h'_i, m_j\} \mid \neg u_i \text{ is a factor in the } j^{\text{th}} \text{ term} \}$$

$$I := \{ \{h_i, h'_i\} \mid 1 \leq i \leq n \}$$

$$M_0 := M$$

If we find a complete explanation  $S$  for this problem, then we claim that the assignment

$$u_i \text{ is True iff } h_i \in S$$

will be a solution to  $P$ .

For any  $j$ ,  $1 \leq j \leq t$ ,  $m_j \in M_0$ , so there is some  $h \in S$  with  $m_j \in e(h)$ .

Case I. If  $h = h_i$ , then  $h_i \in S$ ,  $u_i$  is True, and since  $m_j \in e(h_i)$ , we have, from the construction of  $C$  (and hence,  $e$ ), that  $u_i$  is a factor in the  $j^{\text{th}}$  term of  $P$ , hence the  $j^{\text{th}}$  term of  $P$  has value True.

Case II. If  $h = h'_i$ , then  $h'_i \in S$ .  $\{h_i, h'_i\} \in I$ , so  $h_i \notin S$ , and hence  $u_i$  is False, and since  $m_j \in e(h'_i)$ , we have that  $\neg u_i$  is a factor in the  $j^{\text{th}}$  term of  $P$ , hence again the  $j^{\text{th}}$  term of  $P$  has value True.

Since every term of  $P$  has value True, so does  $P$ .

A similar argument shows that there exists a valid assembly corresponding to each satisfying assignment.

Clearly the construction of this problem from  $P$  is a deterministic polynomial process. Also it is clear that complete explanation construction under incompatibility constraints is NP, so hypothesis assembly with incompatible hypotheses is NP-complete.

## References

1. B. Chandrasekaran and S. Mittal, "conceptual Representation of Medical Knowledge for Diagnosis by Computer: MDX and Related Systems", In *Advances in Computers*, M. Yovits, Ed. Academic Press, 1983, pp. 217-293.
2. E. Charniak and D. McDermott, *Introduction to Artificial Intelligence*, Addison Wesley, 1985.
3. M. Garey and D. Johnson, *Computers and Intractability*, W. H. Freeman and Company, New York, 1979.
4. Josephson, J. R.; Tanner, M. C, Smith, J. W.; Svirbely, J. R.; Strohm, P., "Red: Integrating Generic Tasks to Identify Red-Cell Antibodies." In *Proc. The Expert Systems in Government Symposium, McLean, Virginia, October 24-25, 1985*.
5. J.R. Josephson, B. Chandrasekaran and J. W. Smith, "Assembling the Best Explanation." In *Proceedings of the IEEE Workshop on Principles of Knowledge-Based Systems*. Denver, Colorado, December 1984, pp.185-190.
6. Miller, R.A., Pople, H.E. Jr. and Myers, J.D., "INTERNIST-I, an experimental computer-based diagnostic consultant for general internal medicine." *New England Journal of Medicine* Vol. 307(1982), pp. 468-476.
7. J. Mostow, "IEEE Workshop on Principles of Knowledge Based Systems: A Personal Review", *SIGART Newsletter* 92, April 1985, p. 25.
8. R. S. Patil, P. Szolovits, and W. B. Schwartz, "Modeling Knowledge of the Patient in Acid-Base and Electrolyte disorders", in *Artificial Intelligence in Medicine*, P. Szolovits, ed., Westview Press, 1982, pp. 191-226.
9. H. Pople, "On the Mechanization of Abductive Logic." In *Proceedings of the International Joint Conference on Artificial Intelligence*. 1973, pp. 147-152.
10. J. Reggia, "Diagnostic Expert Systems Based on a Set Covering Model." *International Journal of Man-Machine Studies*. November 1983, pp. 437-460.
11. J.W. Smith, "RED: A Classificatory and Abductive Expert System." PhD Thesis, Ohio State University, Laboratory for Artificial Intelligence Research, Department of Computer Science, Ohio State University, Columbus, Ohio, 1985.
12. J.W. Smith, J.R. Josephson, M.C Tanner, J.R. Svirbely, P. Strohm, "Problem Solving in Red Cell Antibody Identification: RED's Performance on 20 cases," Technical Report, Laboratory for Artificial Intelligence Research, Department of Computer and Information Science, Ohio State University, Columbus, Ohio, 1986.