AAAI-83

A NEW INFERENCE METHOD FOR FRAME-BASED EXPERT SYSTEMS

James A. Reggia, Dana S. Nau, and Pearl Y. Wang

Department of Computer Science
University of Maryland
College Park, MD 20742

ABSTRACT

This paper introduces a new frame-based model of diagnostic reasoning which is based on a generalization of the classic set covering problem in mathematics. The model directly handles multiple simultaneous disorders, it can be formalized, it is intuitively plausible, it provides an approach to partial matching, and it is justifiable in terms of past empirical studies of human diagnostic reasoning. We are using this model as an inference method in diagnostic expert systems, and contrast it with the inference methods used in previous similar systems.

DIAGNOSTIC PROBLEM SOLVING

A diagnostic problem is a problem where one is given a set of abnormal findings (manifestations) for some system, and must explain why those findings are present. Diagnostic problems are common, occurring in medicine, software debugging, automotive repair, electronic circuit fault localization, etc. Search methods, statistical pattern classification, and rule-based deduction face significant limitations when applied to such problems [Reggia, 1982].

Recently a variety of inference methods which model the hypothesize-and-test process involved in human diagnostic reasoning have been proposed, especially in medicine (e.g., [Aikins, 1980; Mittal et al, 1979; Pauker, 1976; Miller et al, 1982; Pople, 1977; Patil et al, 1981]). While these models have produced impressive performance at times, they currently face a number of limitations when applied to real-world problems [Reggia, 1982]. For example, problems where multiple disorders are present simultaneously have proven very difficult to handle [Pople, 1977]. In addition, AI models of diagnostic reasoning are often criticized as being "ad hoc" by individuals outside of AI because of the absence of a formal, domain-independent theoretical foundation (e.g., [Ben-Bassat et al, 1980]).

Acknowledgement: Supported by NINCDS through grants 5 K07 NS 00348 and 1 P01 NS 16332. Dr. Nau was supported in part by NSF grant MCS81-17391. Computer time was provided in part by the Computer Science Center of the University of Maryland.

This paper introduces a new description-based (frame-based) model of diagnostic reasoning which is founded on a generalization of the set covering problem. This model, which we call the "generalized set covering" or GSC model, is of interest for several reasons. It directly addresses the problem of multiple simultaneous disorders, it provides a basis for a formal theory of diagnostic inference, and it provides an approach to such issues as partial match and inference in the context of incomplete problem data. The GSC model is summarized here informally, and further details and example applications are available in [Reggia, 1981; Reggia et al, 1983]. We have already used this model to implement both medical and non-medical expert systems. We view our work as an effort to bring mathematical rigor to an area of AI where it has previously been relatively lacking, and as an attempt to create an abstraction of expert system implementations in the sense that Nilsson has recommended [Nilsson, 1980].

```
BASILAR MIGRAINE
  [DESCRIPTION:
       AGE = FROM 20 THRU 30 <H>, 30 THRU 50 <L>,
             50 THRU 110 <N>;
       DIZZINESS
          [TYPE = VERTIGO <H>, RE$T <L>;
           COURSE = EPISODIC
                [EPISODE DURATION = MINUTES <L>,
                   HOURS <H>, DAYS <L>],
             ACUTE AND PERSISTENT];
       HEAD PAIN <A>
          [LOCATION = OCCIPITAL <H>, RE$T <L>];
       NEUROLOGICAL SYMPTOMS =
          TINNITUS <M>,
          DIPLOPIA [DURATION =
                TRANSIENT DURING DIZZINESS <A>],
          . . .
          SYNCOPE;
       NEUROLOGICAL EXAM FINDINGS =
          HOMONYMOUS FIELD CUT
             [DURATION = TRANSIENT DURING
             DIZZINESS],
          . . .
          CNS FINDINGS
             [TYPE = NON-SPECIFIC <H>, RE$T <L>;
             DURATION = TRANSIENT DURING
             DIZZINESS]                          ]
```

Figure 1: A DESCRIPTION for BASILAR MIGRAINE.

KNOWLEDGE REPRESENTATION

The basic unit of associative knowledge used by the GSC model is the frame-like DESCRIPTION. For each possible causative disorder in the domain of a knowledge base there is a corresponding DESCRIPTION. Figure 1 illustrates a DESCRIPTION for the disorder BASILAR MIGRAINE from the knowledge base of a diagnostic expert system dealing with the problem of dizziness [Reggia, 1982]. Letters in angular brackets represent subjective indications of frequency (A = always, H = high, M = medium, L = low, N = never).

Figure 1 means:
"Basilar migraine usually occurs in individuals from 20 to 30 years old, but many occur up to age 50. If a person is over 50, basilar migraine can be categorically discarded as a possible etiological factor. Basilar migraine causes dizziness which is usually of a vertiginous nature and occurs either in an episodic or an acute and persistent fashion. When episodic, the dizziness usually lasts for hours but may last for minutes or days. Headache, usually in an occipital location, is always present. Neurological symptoms caused by basilar migraine are . . .".

In the current dizziness knowledge base there are 50 disorders like basilar migraine. The key point is that each disorder has an associated DESCRIPTION that specifies, among other things, all manifestations caused by the disorder.

GENERALIZED SET COVERING AS A MODEL OF DIAGNOSTIC INFERENCE

The GSC model provides a useful method for making diagnostic inferences from DESCRIPTIONs without the use of production rules. In the GSC model the underlying knowledge for a diagnostic problem is viewed as pictured in Figure 2a. There are two disjoint finite sets which define the scope of diagnostic problems: **D**, representing all possible disorders d_i that can occur, and **M**, representing all possible manifestations m_j that may occur when one or more disorders are present. For example, in medicine, **D** might represent all known diseases (or some relevant subset of all diseases), and **M** would then represent all possible symptoms, examination findings, and abnormal laboratory results that can be caused by diseases in **D**.

To capture the intuitive notion of causation, we assume knowledge of a relation $\mathbf{C} \subseteq \mathbf{D} \times \mathbf{M}$, where $\langle d_i, m_j \rangle \in \mathbf{C}$ represents "d_i can cause m_j." Note that $\langle d_i, m_j \rangle \in \mathbf{C}$ does not imply that m_j necessarily occurs when d_i is present, but only that m_j may be caused by d_i. Given **D**, **M**, and **C**, the following sets can be defined:

$$man(d_i) = \{m_j | \langle d_i, m_j \rangle \in C\} \; \forall d_i \in D, \text{ and}$$

$$causes(m_j) = \{d_i | \langle d_i, m_j \rangle \in C\} \; \forall m_j \in M.$$

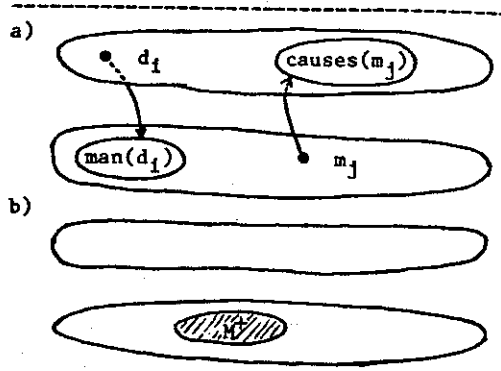

Figure 2: Organization of diagnostic knowledge (a) and problems (b).

These sets are depicted in Figure 2a, and represent all possible manifestations caused by d_i, and all possible disorders that cause m_j, respectively. These concepts are intuitively familiar to the human diagnostician. For example, medical textbooks frequently have descriptions of diseases which include, among other facts, the set $man(d_i)$ for each disease d_i. As noted earlier, the DESCRIPTION of BASILAR MIGRAINE in Figure 1 explicitly defined man(BASILAR MIGRAINE). In addition, physicians often refer to the "differential diagnosis" of a symptom, which corresponds to the set $causes(m_j)$. Clearly, if $man(d_i)$ is known for every disorder d_i, then the causal relation **C** is completely determined. We will use $man(D) = \bigcup_{d_i \in D} man(d_i)$ to indicate all possible manifestations of a set of disorders D, and $causes(M) = \bigcup_{m_j \in M} causes(m_j)$ to indicate all possible causes of any manifestation in M.

Finally, there is a distinguished set $M^+ \subseteq M$ which represents those manifestations which are known to be present (see Figure 2b). Whereas **D**, **M**, and **C** are general knowledge about a class of diagnostic problems, M^+ represents the manifestations occurring in a specific case.

Using this terminology, we define a diagnostic problem P to be a 4-tuple $\langle D, M, C, M^+ \rangle$ where these components are as described above. We assume that $man(d_i)$ and $causes(m_j)$ are always non-empty sets. We now turn to defining a solution to a diagnostic problem by first introducing the concept of explanation.

Definition: For any diagnostic problem P, $E \subseteq D$ is an explanation for M^+ if

(i) $M^+ \subseteq man(E)$, or in words: E covers M^+; and

(ii) $|E| \leq |D|$ for any other cover D of M^+, i.e., E is minimal.

This definition captures what one intuitively means by "explaining" the presence of a set of manifestations. Part (i) specifies the reasonable constraint that a set of disorders E must be able to cause all known manifestations M^+ in order to

be considered an explanation for those manifestations. Part (ii) specifies that E must also be one of the smallest sets to do so, reflecting the Principle of Parsimony or Ockham's Razor: the simplest explanation is the preferable one. This principle is generally accepted as valid by human diagnosticians. Here, we have equated "simplicity" with minimal cardinality, reflecting an underlying assumption that the occurrence of one disorder d_i is independent of the occurrence of another.

An explanation is a generalization of the concept of a minimal set cover [Edwards, 1962]. One difference from the traditional set cover problem in mathematics is that when $M^+ \neq M$, man(E) may be a superset of M^+. This difference, reflects the fact that sometimes when a disorder is present not all of its manifestations occur.

With these concepts in mind, we can now define the solution to a diagnostic problem P, designated Sol(P), to be the set of all explanations for M^+. Thus, solving a diagnostic problem in the GSC model involves a second generalization of the traditional set covering problem: we are interested in finding all explanations rather than a single minimal cover.

Example: Let P = <D,M,C,M^+> where D = $\{d_1,d_2,\ldots,d_9\}$, M = $\{m_1,\ldots m_6\}$, and man(d_i) are as specified in Table 1. Note that Table 1 implicitly defines the relation C, because C = $\{ \langle d_i,m_j\rangle \mid m_j \in man(d_i)$ for some $d_i\}$. Let M^+ = $\{m_1,m_4,m_5\}$. No single disorder can cover (account for) all of M^+, but some pairs of disorders do cover M^+. For instance, if D = $\{d_1,d_7\}$ then $M^+ \subseteq man(D)$. Since there are no covers for M^+ of smaller cardinality than D, it follows that D is an explanation for M^+. Careful examination of Table 1 should convince the reader that

Sol(P) = { $\{d_1\ d_7\}$ $\{d_1\ d_8\}$ $\{d_1\ d_9\}$ $\{d_2\ d_7\}$ $\{d_2\ d_8\}$ $\{d_2\ d_9\}$ $\{d_3\ d_8\}$, $\{d_4\ d_8\}$ }

is the set of all explanations for M^+.

d_i	man(d_i)	d_i	man(d_i)
d_1	m_1 m_4	d_6	m_2 m_3
d_2	m_1 m_3 m_4	d_7	m_2 m_5
d_3	m_1 m_3	d_8	m_4 m_5 m_6
d_4	m_1 m_6	d_9	m_2 m_5
d_5	m_2 m_3 m_4		

Table 1: Knowledge about a class of diagnostic problems (C is implicitly defined by this table).

Rather than representing the solution to a diagnostic problem as an explicit list of all possible explanations for M^+, it is advantageous to represent it as a collection of explanation generators. A generator is analogous to a Cartesian set product, the difference being that the generator produces unordered sets rather than ordered tuples. To illustrate this idea, consider the example diagnostic problem above. Two generators are sufficient to represent the solution to that problem: $\{d_1\ d_2\}$ x $\{d_7\ d_8\ d_9\}$ and $\{d_3\ d_4\}$ x $\{d_8\}$. The second generator represents two explanations $\{d_3\ d_8\}$ and $\{d_4\ d_8\}$, while the first generator represents the other six explanations in the solution. Generators are usually a more compact form of the explanations present in the solution, they are a convenient representation for developing algorithms to process explanations sequentially (see below), and they are closer to the way the human diagnostician organizes the possibilities during problem solving (i.e., the "differential diagnosis").

In adapting the GSC model for use in a real-world expert system several issues were addressed and resolved. One of these issues is the fact that diagnostic problem solving is inherently sequential in nature. The human diagnostician usually begins knowing only that one or a few manifestations are present, and must actively seek further information about others.

This sequential diagnostic process can be captured in terms of the GSC model, and represents a third generalization of the traditional set covering problem. The tentative hypothesis at any point during problem solving is defined to be the solution for those manifestations already known to be present, assuming, perhaps falsely, that no additional manifestations will be subsequently discovered. To construct and maintain a tentative hypothesis like this, three data structures prove useful:

MANIFS: the set of manifestations known to be present so far;

SCOPE: causes(MANIFS), the set of all diseases d_i for which at least one manifestation is already known to be present; and

FOCUS: the tentative solution for just those manifestations already in MANIFS; FOCUS is represented as a collection of generators.

These data structures are manipulated as follows:

(1) Get the next manifestation m_j.
(2) Retrieve causes(m_j) from the knowledge base.

(3) MANIFS ← MANIFS ∪ $\{m_j\}$.
(4) SCOPE ← SCOPE ∪ causes(m_j).
(5) Adjust FOCUS to accomodate m_j.
(6) Repeat this process until no further manifestations remain.

Thus, as each manifestation m_j that is present is discovered, MANIFS is updated simply by adding m_j to it. SCOPE is augmented to include any possible causes d_i of m_j which are not already contained in it. Finally, FOCUS is adjusted to accomodate m_j based partially on intersecting causes(m_j) with the sets of disorders in the existing generators [Reggia et al, 1983]. These latter operations are done such that any

explanation which can no longer account for the augmented MANIFS (which now includes m_4) are eliminated. Figure 3 illustrates this algorithm with a "trace" based on the earlier example.

DISCUSSION

This paper has proposed the construction and maintenance of generalized minimal set covers ("explanations") as a model of diagnostic reasoning and as a method for diagnostic expert systems. The GSC model is attractive for several reasons: it directly handles multiple simultaneous disorders, it can be formalized, it is intuitively plausible, it provides an approach to partial matching, and it is justifiable in terms of past empirical studies of diagnostic reasoning (e.g., [Elstein et al, 1978; Kassiner et al, 1978]). To our knowledge the analogy between the classic set covering problem and general diagnostic reasoning has not previously been examined in detail, although some related work has been done (e.g., assignment of HLA specificities to antisera, see Nau et al, 1978; Woodbury et al, 1979]). As noted earlier, other aspects of the GSC model relevant to expert systems, such as question generation, termination criteria, ranking of competing disorders, and problem decomposition are discussed elsewhere [Reggia et al, 1983 and 1984].

The GSC model provides a useful context in which to view past work on diagnostic expert systems. In contrast to the GSC model, most diagnostic expert systems that use hypothesize-and-test inference mechanisms or which might reasonably be considered as models of diagnostic reasoning depend heavily upon the use of production rules (e.g., [Aikins, 1980; Mittal et al, 1979; Pauker et al, 1976]). These systems use a hypothesis-driven approach to guide the invocation of rules which in turn modify the hypothesis. Rules have long been criticized as a representation of diagnostic knowledge [Reggia, 1982], and their invocation to make deductions or perform actions does not capture in a general sense such intuitively attractive concepts as coverage, minimality, or explanation.

Perhaps the previous diagnostic expert system whose inference method is closest to the GSC model is INTERNIST [Miller et al, 1982]. INTERNIST represents diagnostic knowledge in a DESCRIPTION-like fashion and does not rely on production rules to guide its hypothesize-and-test process. In contrast to the GSC model, however, it uses a heuristic scoring procedure to guide the construction and modification of its hypothesis. This process is essentially "depth first," unlike the "breadth first" approach implied in the GSC model. INTERNIST first tries to establish one disorder and then proceeds to establish others. This roughly corresponds to constructing and completing a single set in a generator in the GSC model, and then later returning to construct the additional sets for the generator. INTERNIST groups together competing disorders (i.e., a set of disorders in a generator) based on a simple but clever heuristic: "Two diseases are competitors if the items not explained by one disease are a subset of the items not explained by the other; otherwise, they are alternatives (and may possibly coexist in the patient)." [Miller et al, 1982]. In the terms of the GSC model, this corresponds to stating that d_1 and d_2 are competitors if M^+-man(d_1) contains or is contained in M^+-man(d_2). While this simple heuristic often works in constructing a differential diagnosis, we can produce examples in the context of the GSC model for which it will fail to correctly group competing disorders together.* It is also unclear that the INTERNIST inference mechanism is

*For example, suppose $M^+ = \{m_1 \ldots m_8\}$ and only d_1, d_2, and d_3 have been evoked where $M^+ \cap$ man(d_1) = $\{m_2\ m_4\ m_5\ m_6\ m_7\ m_8\}$, $M^+ \cap$ man(d_2) = $\{m_3\ m_4\ m_5\ m_6\ m_7\ m_8\}$, and $M^+ \cap$ man(d_3) = $\{m_1\ m_2\ m_3\}$. In the GSC model, Sol(P) = $\{\ \{d_1\ d_3\}\ \{d_2\ d_3\}\ \}$ which can be represented by the single generator $\{d_1\ d_2\}$ x $\{d_3\}$ where d_1 and d_2 are grouped together as competitors. Suppose that d_1 was ranked highest by the INTERNIST heuristic scoring procedure. Then M^+-man(d_1) = $\{m_1\ m_3\}$ and M^+-man(d_2) = $\{m_1\ m_2\}$, so INTERNIST would apparently fail to group d_1 and d_2 together as competitors.

Events in order of their discovery	MANIFS	SCOPE	FOCUS
Initially	$\emptyset$	$\emptyset$	$\emptyset$
m_1 present	$\{m_1\}$	$\{d_1\ d_2\ d_3\ d_4\}$	$\{d_1\ d_2\ d_3\ d_4\}$
m_2 absent	"	"	"
m_3 absent	"	"	"
m_4 present	$\{m_1\ m_4\}$	$\{d_1\ d_2\ d_3\ d_4\ d_5\ d_8\}$	$\{d_1\ d_2\}$
m_5 present	$\{m_1\ m_4\ m_5\}$	$\{d_1\ d_2\ d_3\ d_4\ d_5\ d_7\ d_8\ d_9\}$	$\{d_1\ d_2\}$ x $\{d_7\ d_8\ d_9\}$ and $\{d_8\}$ x $\{d_3\ d_4\}$
m_6 absent	"	"	"

Figure 3: Sequential problem solving using the set covering model.

guaranteed to always find all possible explanations for a set of manifestations. Reportedly, the "depth first" approach used in INTERNIST resulted in less than optimal performance [Miller et al, 1982]. Recent enhancements in INTERNIST's successor CADUCEUS attempt to overcome some of these limitations through the use of "constrictors" to delineate the top-level structure of a problem [Pople, 1977]. These changes are quite distinct from the approach taken in the GSC model, but they do add a "breadth first" component to hypothesis construction.

We are currently developing the GSC model in two ways: by studying its application in medical expert systems and by formally developing the mathematical theory. Currently, we have implemented two medical diagnostic expert systems based on the GSC model, one for dizziness (a difficult medical problem because of the many possible causes) and one for peroneal muscular atrophy [Reggia, 1981; Reggia et al, 1983]. While the GSC model forms the central mechanism of these expert systems, the basic model was augmented in a number of ways to make it more useful for real-world problem solving. For example, the "symbolic probabilities" illustrated in Figure 1 were introduced and are used to rank competing explanations after the final FOCUS is constructed. A heuristic approach to question generation and termination was adopted. When tested on prototype cases these expert systems functioned well, but modifications to the content of the knowledge bases (not the GSC model) would be necessary before more extensive evaluation in practice using a series of real patients could be done.

In parallel, we are developing the mathematical basis of the GSC model [Reggia et al, 1984]. This has involved defining a variety of operations on generators and expressing formal algorithms in terms of those operations. We are proving the correctness of the these algorithms and have established criteria for decomposing diagnostic problems into independent subproblems that are easier to solve. While the GSC model as it currently exists does not address all aspects of diagnostic problem solving, it does appear to provide a reasonable starting point from which to formalize the underlying abductive inference process that is involved.

REFERENCES

1. Aikins J: Prototypes and Production Rules - A Knowledge Representation for Computer Consultations, Memo HPP-80-17, Stanford Heuristic Programming Project, 1980.
2. Ben-Bassat M, et al: Pattern-Based Interactive Diagnosis of Multiple Disorders - the MEDAS System, IEEE Trans. Pat. Anal. Machine Intell., 2, 1980, 148-160.
3. Edwards J: Covers and Packings in a Family of Sets, Bull. Am. Math. Society, 68, 1962, 494-499.
4. Elstein A, Shulman L, and Sprafka S: Medical Problem Solving - An Analysis of Clinical Reasoning, Harvard University Press, 1978.
5. Karp R: Reducibility Among Combinatorial Problems, in R. Miller and J. Thatcher (eds.), Complexity of Computer Computations, Plenum Press, New York, 1972, 85-103.
6. Kassirer J and Gorry G: Clinical Problem Solving - A Behavioral Analysis, Ann. Int. Med., 89, 1978, 245-255.
7. Miller R, Pople H, Myers J: Internist-1: An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine, NEJM, 307, 1982, 468-476.
8. Mittal S, Chandrasekaran B, and Smith J: Overview of MDX - A System for Medical Diagnosis, Proc. Third Symposium on Computer Applications in Medical Care, IEEE, 1979, 34-46.
9. Nau D, Markowsky G, Woodbury M & Amos D: A mathematical analysis of human leukocyte antigen serology, Math Biosci., 40, 1978, 243-270.
10. Nilsson N: The interplay between experimental and theoretical methods in artificial intelligence, Cognition and Brain Theory, 4, 1980, 69-74.
11. Patil R, Szolovits P, and Schwartz W: Causal understanding of patient illness in medical diagnosis, Proc. Seventh IJCAI, Yale University, New Haven, CT, 1981.
12. Pauker S, et al: Towards the Simulation of Clinical Cognition, Am. J. Med., 60, 1976, 981-996.
13. Pople H: The Formation of Composite Hypotheses in Diagnostic Problem-Solving - An Exercise in Synthetic Reasoning, IJCAI, 5, 1977, 1030-1037.
14. Reggia J: Knowledge-Based Decision Support Systems - Development Through KMS, TR-1121, Department of Computer Science, University of Maryland, Oct., 1981.
15. Reggia J: Computer-Assisted Medical Decision Making, in Applications of Computers in Medicine, M. Schwartz, ed., IEEE Press, 1982, 198-213.
16. Reggia J, Nau D, and Wang P: Diagnostic Expert Systems Based on a Set Covering Model, Int. J. Man. Machine Studies, to appear late 1983.
17. Reggia J, Nau D, and Wang P: A formal model of diagnostic inference, 1984, manuscript in preparation.
18. Woodbury M, Ciftan E, & Amos D: HLA serum screening based on an heuristic solution to the set cover problem, Comp. Pgm. Biomed., 9, 1979, 263-273.