Predicate Calculus, Artificial Intelligence, and Workers' Compensation

Philip Harber, MD, MPH; and J. Michael McCoy, MD

Application of principles of predicate calculus (PC) and artificial intelligence (AI) search methods to occupational medicine can meet several goals. First, they can improve understanding of the diagnostic process and recognition of the sources of uncertainty in knowledge and in case specific information. Second, PC provides a rational means of resolving differences in conclusion based upon the same premises. Third, understanding of these principles allows separation of knowledge (facts) from the process by which they are used and therefore facilitates development of AI-based expert systems. Application of PC to recognizing causation of pulmonary fibrosis is demonstrated in this paper, providing a method that can be generalized to other problems in occupational medicine. Application of PC and understanding of AI search routines may be particularly applicable to workers' compensation where explicit statement of rational and inferential process is necessary. This approach is useful in the diagnosis of occupational lung disease and may be particularly valuable in workers' compensation considerations, wherein explicit statement of rationale is needed.

Two physicians may examine the same patient and reach different diagnostic conclusions. Occasionally, such differences in opinion occur if one physician has inadequate knowledge of either "textbook information" (generalized knowledge such as "asbestos causes restrictive lung disease") or case-specific information (eg, this particular patient's carbon monoxide diffusing capacity is abnormally low). However, differences in conclusions may more commonly be due to differences in utilization of information rather than inadequate information. Explicit understanding of how information is utilized in clinical inference is requisite for development of artificial intelligence (AI)-based expert systems in medicine. In developing an expert system to facilitate clinical recognition of occupational lung disease, it became evident that predicate calculus, fuzzy set theory, and AI principles could explain and provide a rational basis for resolving differences in diagnostic conclusions. Therefore, the means of clinically using information can be as explicit and objective as the research techniques employed for acquiring knowledge.

Predicate calculus is the branch of logic that provides a general method for determining the truth of statements about facts. Fuzzy set theory is a branch of mathematical set theory permitting quantitative analysis when inclusion in a set is not absolutely dichotomous (absolutely yes or absolutely no). Artificial intelligence methods include strategies for searching information to identify appropriate relationships relevant to the diagnostic process. In addition, the underlying search strategy must be specified.

In this application, these methods are applied to the diagnosis of occupational lung disease such as pulmonary fibrosis, but the method is generalized and easily applicable to many other areas.

Methods and Results

The first step in analyzing the utilization of information is its expression in symbolic format. This method is illustrated for diagnosing occupationally related pulmonary interstitial disease (eg, asbestosis). Consider the relatively simple diagnostic paradigm: if a worker has had significant asbestos exposure and has pulmonary fibrosis, then he has asbestos-related lung disease. First, the problem must be represented symbolically. Table 1 illustrates the symbolic approach. In general,

From the Occupational Medicine Branch/Pulmonary Division and General Internal Medicine Division/Department of Medicine, University of California, Los Angeles and the University of California Southern Occupational Health Center, Los Angeles, CA 90024 (Dr Harber, Associate Professor of Medicine and Chief, Occupational Medicine Branch; Dr McCoy, Assistant Dean for Medical Computing). Address correspondence to Dr Harber).

0096-1736/89/3105-0484$02.00/

Table 1 Rule Representations

1.1	Asbestos exposure + pulmonary fibrosis implies asbestosis disease	
1.2	AE + PF ➔ AD	
1.3	AE + PF ➔ AD (CF = 80)	
1.4	AE (CF=80) + PF (CF=60) ➔ AD (CF=?)	CF of AD = 60 (minimum), or = .60 x .80 = .48 (probability product), or = (.60 = 80)/2 = 70 (average)
1.5	FEE* + PF ➔ FED† (CF=70)	
1.6	FEE (CF=80) + PF (CF=60) ➔ FED (CF=?) (CF=70)	CF of FED = minimum (.60, .80) x .70

* FEE, fuller's earth exposure.
† FED, fuller's earth disease.

an if-then proposition of the form X implies Y ($X \rightarrow Y$) can express most diagnostic paradigms. The diagnostic rule described will be of the form AE + PF $\rightarrow$ AD, meaning asbestos exposure (AE) plus pulmonary fibrosis (PF) implies asbestosis disease (AD). With such a conjunction, both components of the predicate clause (the "if" clause) must be true in order for the conclusion (the "then" clause) to be true. This is analogous to common clinical practice in which neither asbestos exposure alone nor the presence of pulmonary fibrosis without known exposure should lead to the diagnosis of asbestosis disease.

Often, the presence or absence of AD in a particular worker cannot be determined with absolute certainty. Recognizing and dealing with clinical uncertainties is the second step in understanding clinical inference. Unfortunately, traditional propositional logic does not fully apply in this medical setting for both clinical and theoretical reasons. The "law of the excluded middle" of traditional logic requires that a statement either be absolutely true or false.[1], p. 223 However, reality often does not correspond to the rule of the excluded middle: clinicians frequently state, "I am *fairly* certain that he has asbestosis disease"; workers' compensation judges base decisions on the statement, "it is more probable than not that he has AD." Statisticians state that "a study demonstrated that PF implies AD with a $P < .05$" (indicating a 5% chance that the "true fact" is in error), and even logicians have derived mathematical models to deal with truth values in the continuum between 0 and 1. Therefore, the expanded symbolic model must permit conclusions about the certainty of AD which are neither 0 nor 1, but may be in the intermediate range (eg, there is a "confidence factor" (CF) of 0.80 that he has asbestosis disease). The conclusion need not be stated with absolute certainty but can be quantitatively expressed.

Certainty of a statement, expressed by the confidence factor (CF), has considerable overlap with probability but it is not synonymous. In the clinical sense, the statement that "I am 70% certain that he has AD" does not contradict the statement that, "if a thorough autopsy were performed, it would be shown that he certainly has AD or he certainly does not have AD."

There are several sources of clinical uncertainty:

(1) Competing Causes: How commonly do nonoccupational factors lead to the disease?
(2) Case-specific Data: How certain are the individual "clinical findings"?
(3) Entire Rule: How true is the relationship?
(4) How to combine certainties?

Competing causation is the first origin of uncertainty about conclusions. Pulmonary fibrosis occurs occasionally even in the absence of AE, although it is considerably more frequent with AE. For example, if the relative risk (determined epidemiologically) for interstitial disease with AE were 4.0, then 80% of cases of fibrosis in an asbestos-exposed population would be attributable to asbestos and 20% to other diseases. This concept, applied to a population, is the attributable risk and might be extrapolated on an individual case basis, perhaps as the "ascribable risk".[2] This concept is therefore subject to empiric determination at least on a population basis.

The next source of uncertainty is related to questions in the premise itself. Each of the predicate's clauses may in itself be uncertain. For example, AE may be intermediate between absolutely true and absolutely false (eg, "I am 60% certain that he had significant AE"). The CF related to this clause has two sources of uncertainty: First, what is the confidence that the individual actually had asbestos exposure? Second, what is the likelihood that it was "significant"? Similarly, the presence of PF may be uncertain. Radiologists frequently "hedge" their interpretations (eg, "80%" sure that PF is present). Differences in the interpretation of how uncertainty in the "if" phrase affects the conclusion may underlie many of the differences in physician opinions. This specific instance represents a conjunction, in which both components are necessary for the conclusion to be valid. Table 2 illustrates several means of combining the uncertainties of the predicate to determine the certainty of the conclusion of AD. From the probability theory standpoint, the conclusion (AD) has a probability of the product of the probability of AE and the probability of PF (ie, $.60 \times .80 = .48$). Alternative means of combining the uncertainties are, however, available and may be more appropriate. The probability interpretation assumes that CFs represent probabilities, but actually a CF represents certainty (including subjective estimates) and/or represents membership in a "fuzzy set" (discussed below). Hence, it is reasonable to follow the statement that, "the conclusion is only as strong as the weakest link in the chain of logic" (in this case, giving the conclusion of AD a CF of .60). A detailed discussion of the theoretical basis for this combining rule is provided by Winston,[3] Parsaye and Chingnell,[1, Ch 6] and Negoita.[4] Other alternatives, with their common sense justifications, are also shown in the table.

Fuzzy set theory provides a conceptual basis for dealing with these uncertainties. Unlike classic set theory, in which an element absolutely is or absolutely is not a member of a set, fuzzy set theory associates a degree of membership with inclusion in a set.[1, pp 223–228; 4] The CF may be considered to reflect the degree of membership in a set. A classic example illustrating the utility of fuzzy set theory has been ascribed to the Greeks[1, pp 223–228]: There is a pile of 100 stones, and one concludes that this pile certainly is a member of the set of "large piles." If one stone is removed, this bunch of stones would still be a member of the set of "large piles." However, as

Table 2 Combining Uncertainties

2A. Relationship between jobs and exposures

2A.1 Custodian → AE-C → AE
CF=40 CF=40

2.A.2 Shipyard worker → AE - SYW → AE
CF=99 CF=99

2A.3 Plumber → AE - P → AE
CF=60 CF=60

2A.4 Schoolteacher → AE - ST → AE
CF=4 CF=4

2A.5 Welder → nitrogen oxide exposure W → NOX
CF=90 CF=90

2B. Example of backward chain

2B.1 AE + PF → AD
? CF=40 ?

2B.2 Custodian → AE
CF=40

2B.3 AE = PF → AD
CF=40 CF=40 CF=40

2 A.1 to 2A.5 represent facts in the knowledge base. To determine whether asbestosis disease (AD) is present in example 2B, one must determine whether exposure (AE) is present; this is found by searching the knowledge base, finding rule 2A.1, and using it to reach a conclusion about AD (2B.3).

more stones are progressively removed, ultimately a time is reached when the pile is no longer a member of the set. However, this does not occur suddenly, and there is therefore a continuum of degree of membership of the pile in the set of "large piles." Using this analogy, one may derive that the membership of an element in the set composed of the union of two fuzzy sets is expressed by the minimum of its association with either of the sets. (Example: a particular pile of stones is associated with the "large pile" set with an association measure of 0.80, and its color, being slightly washed out, associates it with the set of "red stones" with an association of 0.60. Hence, its association with the set of "large piles of red stones" would be expressed as 0.60.)

A third source of uncertainty is the certainty of the entire rule (rather than of a particular clause thereof). Consider, for example, a situation if fuller's earth (FE) was the exposure of concern rather than asbestos. Symbolically, Table 1, line 1.5 summarizes these considerations. Unlike the situation for asbestos, there is some epidemiologic and other scientific uncertainty about whether FE causes significant interstitial lung disease. That is, there is a confidence factor associated with the entire statement that FE exposure (FEE) plus pulmonary fibrosis (PF) implies FE disease (FED). In Table 1 this is expressed by the parentheses bracketing the entire expression, implying that the CF applies to the entire rule. Here, as seen in line 1.5 of Table 1, a CF of 70 for the entire rule implies that, even if it is absolutely certain that an individual had FEE and that he has PF, it is only 70% certain (CF = 70) that he has FED.

Uncertainty associated with the overall rule interacts with uncertainty in the predicate clause as shown in line 1.6 of Table 1. In essence, a CF applying to the entire rule also applies to the certainty of the conclusion. Therefore, the CF of the conclusion is calculated and then multiplied by the overall rule CF. In the example shown in which there is uncertainty about FEE and about PF, the conclusion of FED has an assigned CF based upon the fuzzy set theory rule of taking the minimum CF of a conjunction and then using the rule of multiplying by the overall CF. Hence, the assigned calculated CF is minimum (.80, .60) × .70 = .60 × .70 = .42 ("less likely than not").

In addition to application of predicate calculus and fuzzy set theory, development of the expert system further suggests application of AI search procedures. In actuality, neither clinicians nor expert systems work with only a single rule, but rather rely upon a very large knowledge base composed of many rules, selecting the appropriate ones as needed. Furthermore, application ("firing" in AI jargon) of a particular rule depends upon establishing the truth of its predicate clause. Determination of whether AD is present, for example, depends upon knowing how true is AE and how true is PF. Lack of knowledge about the CF (CF = unknown) is different from knowing that it is untrue (CF = 0), because in the latter instance one can then conclude that AD is untrue (ie, if AE has a CF of 0, then application of the minimum combining rule shows that AD has a CF of 0). Clinicians and good AI expert systems do not rely upon a single rule, but can search their large knowledge bases of many rules to determine which are applicable. For example, if the "goal" is to determine whether AD is present, knowledge base search could identify all rules with AD in the conclusion. However, these rules cannot be directly utilized unless the premise clause components' CFs are known. Such a knowledge base search for the goal of identifying whether AD is present will of course locate the rule used in this example. The rule could be directly applied to reach the conclusion if a worker with definite PF said, "I am an asbestos-exposed worker with a CF of .70." However, workers generally do not present such information, but rather provide information such as, "I worked as a custodian . . . " Therefore, the clinician/computer must "backward chain" further. For example, consider the knowledge base partially illustrated in Table 2, composed of a series of rules about relationships between jobs and exposures, laboratory tests and PF, and clinical findings and diagnoses. The expert system/clinician backward chains from the goal of AD to the premise of AE. Recognizing that this is unknown, further backward chaining to rules having AE in their conclusion is necessary (2A.1–2A.4). Thereafter, when the premise of the rule (2A.1) for the custodian, a CF is assigned to AE in the conclusion of the former rule and then this CF (40) is applied to AE in the premise clause of the final rule, leading to a diagnosis of AD with 40% certainty. Of course, the process can extend back many more levels.

Combination of the AI search method and the predicate calculus method of determining CF can therefore permit an objective, fully explicable logical train and avoid conflicting conclusions as illustrated by the following example:

"He worked as a janitor, some janitors have asbestos exposure, and he probably has pulmonary fibrosis on this radiograph; I therefore conclude that he has asbestosis disease," v, "He worked as a janitor, but not all janitors have asbestos exposure; the radiologist is not completely certain that pulmonary fibrosis is present anyway; I therefore conclude that he clearly does not have asbestosis disease."

In this example, the same facts presented by a patient and his medical evaluation lead to apparently contradictory conclusions by two "competent" clinicians. These methods, however, permit objective resolution of the apparent conflict. Without clear thinking about symbolic representation, the two physicians would reach opposite conclusions even if they dealt with the same facts and the same confidence factors in each fact. The physicians' statements quoted above differ in how the facts are used not in the facts themselves.

A final consideration is the delineation of the underlying search strategy (whether by "clinician" or "computer").[1, Ch 7; 3] Two general methods of searching the knowledge base exist. The "backward chaining" method was illustrated above. In this method, a "goal" is defined and the knowledge base and inference system is used to determine whether this goal can be validated (ie, the system starts with the goal of determining whether asbestosis disease is present). It searches the knowledge base so that the premise of every rule with AD in its conclusion is assessed, backward chaining further if necessary. Hence, if at all possible, the system will find rules which ultimately lead to demonstrating the AD is present.

An alternative search strategy is to "forward chain," in which the system is data-driven, starting with the facts provided by the patient himself or herself rather than a goal to be reached. With a forward chain technique, one does not seek to determine whether AD is present, rather seeking to determine the implications of the specific facts presented by the patient. In expert system development (and by analogy, in clinical thinking) there is no a priori optimal choice between a forward or backward chain strategy. However, the search strategy employed by a clinician to select facts from his or her knowledge base might help explain differences in conclusions. One physician may start with the goal of finding an occupational disease and search through the available facts about the specific patient and general facts in the knowledge base to determine whether there is any way to support such a conclusion. Another, equally competent, might use the alternative "forward chain" search strategy, starting with the facts presented by the patient and following where they lead in a deductive manner. Particularly if the clinician or expert system does not employ a fully exhaustive search strategy (ie, is selective about which facts to include), different conclusions might be reached. A completely exhaustive search strategy is one which examines all possible links, but in reality this is practical neither in the clinical nor in the computer systems. Hence, choice between a backward and forward chaining search strategy can by analogy help explain differences in conclusions. This is further illustrated in the example below. (Actually, forward and backward chain searches are not mutually exclusive in either the clinical or computer sense. For example, the expert system we are developing uses a forward search to reach tentative conclusions and eliminate obviously extraneous facts and then backward chains to seek additional confirmatory information. Similarly, experienced clinicians actually "forward chain" from a small number of clinical observations to reach tentative hypotheses about diagnoses and then backward chain to seek confirmatory information to accept or reject their hypotheses.)

Comparison of Forward and Backward Chaining

Figures 1 and 2 illustrate a comparison of the two reasoning modes. Both employ the same patient-specific facts:

(1) Patient is a plumber (P = true).
(2) He has rheumatoid arthritis (RA = true).
(3) He has pulmonary fibrosis with 90% certainty (PF CF = 90).

Exposure rules are those shown in Table 2, and diagnostic rules are those of Table 1. In addition, a diagnostic rule is added: Pulmonary fibrosis in a patient with rheumatoid arthritis suggests that rheumatoid lung disease (RLD) is probable: (RA + PF → RLD) (CF = 95).

Figure 1 demonstrates backward chaining. In Panel 1A, the goal is to determine whether AD is present. The system identifies rules with AD in conclusion and evaluates their premise clause. Since AE and PF are not yet known, it backward chains to try to find their values. It therefore seeks rules with PF or AE in their conclusions. In Panel B, it finds 3 rules with AE in conclusion and

Backward Chain Reasoning

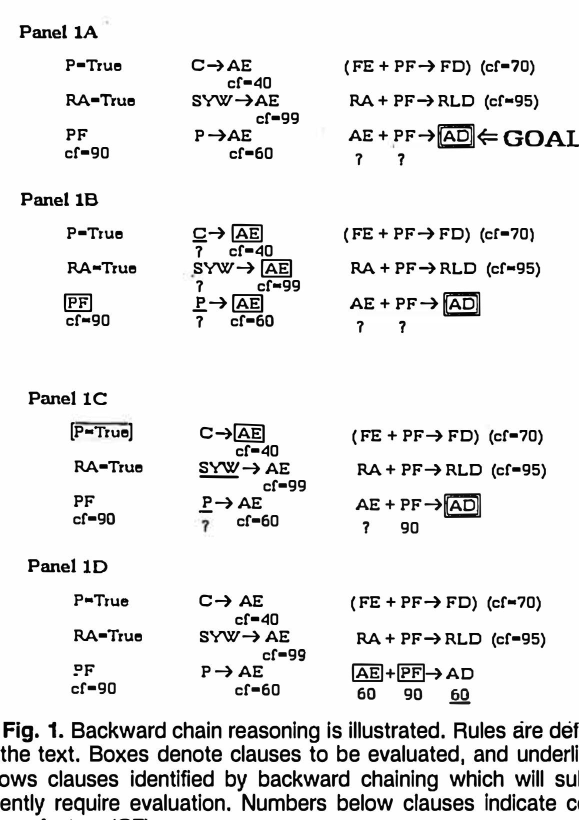

Fig. 1. Backward chain reasoning is illustrated. Rules are defined in the text. Boxes denote clauses to be evaluated, and underlining shows clauses identified by backward chaining which will subsequently require evaluation. Numbers below clauses indicate confidence factors (CF).

one with PF. Although the value of PF is immediately known, it must further backward chain to determine values of AE and therefore must check the three underlined premise clauses. Since their values are unknown, it backward chains further and (Panel 1C) finds the value for P. Hence, as shown in Panel 1D, it now knows values of both AE and PF and can therefore conclude that AD is is present with a CF = 60.

Figure 2 illustrates forward chain inference. In Panel 2A, the system starts with the patient-derived data (rather than the goal of AD). From the patient data, the values in the premises of other rules are derived, shown by the double boxes of CFs. The implications (conclusions) to which these lead are shown in Panel 2B. In particular, since the values of all clauses in the RLD diagnostic rule are known, the system concludes that RLD is present with a CF = 86. It may be programmed to proceed further or to stop. If it proceeds, Panel 2C shows that the clauses of the premise of the AD diagnostic rule are now known, and the system can also diagnose AD with CF = 60.

As illustrated, the two different methods may lead to different conclusions: In one instance, AD is found. The other method either finds only RLD as a diagnosis or, if it is designed to proceed further, finds both RLD and AD, but states that RLD is much more likely.

There are clinical analogies to these computer algorithms. Differences in diagnostic conclusions may depend on whether the clinician seeks to determine whether AD is present or seeks to see simply where the patient data lead.

Discussion

Rational, objective inference underlies much of Western thought and the scientific process. Furthermore, it should optimally be incorporated in the practice of medicine. Clear specification of the inference process is particularly important in two circumstances: the development of computer-based AI expert systems to emulate clinical thinking and, second, to facilitate the understanding by nonclinicians (eg, workers' compensation judges) of conflicting clinical reports. The application of predicate calculus, fuzzy set theory, and artificial intelligence search methods to the clinical problem illustrated herein may be generalized without difficulty to many other clinical questions.

Central to the application of these methods is the concept that truth is not absolute (ie, a fact need not be exclusively either absolutely true or absolutely false). Such a simple world view does not correspond to clinical reality, although the legal system occasionally forces occupational physicians to behave as if the world were composed of black and white with no shades of grey. Instead, facts may be stated with varying degrees of certainty.

Fuzzy set theory is similarly a means of dealing with uncertainties and is used in a very complementary fashion. Whereas in classical set theory an individual element is absolutely a member or a non-member of a particular set, fuzzy set theory provides a mathematical and conceptual framework for allowing varying degrees of membership in a particular set. This is directly analogous to truth: saying that, "I am 80% confident that it is true that he has asbestosis disease" is equivalent to saying, "his degree of membership in the set of persons with asbestosis disease is '80%'."

As discussed earlier in this paper, occupational medicine is replete with uncertainties. These include uncertainties in general knowledge (whether a particular agent causes disease, of occupational *v* nonoccupational causation, etc) and of facts about a specific worker (eg, how heavy was his or her exposure, is the radiograph truly abnormal, etc). The methods discussed herein provide a rational means of inference in view of real world uncertainty.

Certainty and confidence include but are not limited to probability considerations. Certainly, probability-based empiric observations are preferable to other forms of certainty, but these are not always available or fully accepted. Occupational health professionals are well versed in epidemiologic methods, but good epidemiologic data are often unavailable about a particular agent, are controversial (and hence not certain themselves), are possibly not generalizable to include the specifics of a particular case, or are contradictory. The recent popularization of "meta-analysis"[5,6] attests to the inability to accept epidemiologic study-based *P* values and estimates at face value.

Although these methods provide a scientifically valid, externally verifiable means of resolving differences in inference, they are not fully deterministic. For example, several means of combining confidence factors in the conjunction example (Table 1) were shown. Nevertheless, in most instances, these methods do permit resolution of conflicts.

These methods are not limited to individual clinical

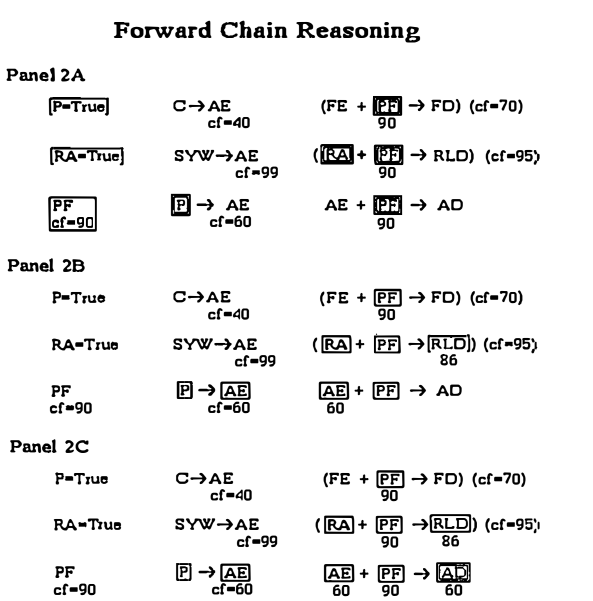

Fig. 2. Forward chain reasoning is illustrated. Terms are defined in the test. Double boxes denote clauses whose values have just been determined, and single boxes indicate clauses whose values are already known. Numbers represent confidence factors (CF).

cases as discussed in this paper. Rather, these methods may be applied more generally to members of populations and used in surveillance efforts and epidemiologic studies. The AI-based expert system being developed by the authors permits application in both the clinical and the population settings. However, the underlying principles are the same. Indeed, one advantage of applying formal logical methods to clinical, single-case-oriented situations is that it may provide a rational link between "objective scientific population-based studies" and clinical medicine.

In summary, predicate calculus, fuzzy set theory, and artificial intelligence search methods permit emulation and explanation of clinical inference. As such, they are important tools in development of computer-based diagnostic systems to assist physicians and other health care providers. They also can help improve diagnostic practices of clinicians by helping to focus upon the process by which conclusions are reached. Finally, they are particularly useful in those situations in which two competent clinicians present opposite conclusions about the same question even when starting with the same facts. Attention to the inference process as well as to the data employed holds great promise as a means for resolving conflicts in conclusions. Of course, when one physician is ignorant of the facts or when different facts are employed, attention to the inference process itself cannot resolve conflicts and conclusions. Nevertheless, in such instances, clear separation of the data from the means by which the data are used may still be useful by helping to clearly identify those data components which must be more carefully assessed. Although the techniques have a considerable mathematical underpinning, they are not hopelessly obscure and limited to application by a small number of mathematical theoreticians. As emphasized in this paper, these methods have a clear, common-sense basis and may therefore be of general utility.

Clinical medicine has shaped the early development of artificial intelligence and now methods underlying AI can be used to improve the understanding of clinical medicine.

Acknowledgments

This work was supported by grant RO1-OH-02288 from the National Institute for Occupational Safety and Health.

We thank Nancy Marshello for preparing the manuscript and Kevin Bailey, Katherine Howard, John Luo, Philip Coffman, and Ronald Kanter for their technical assistance.

References

1. Parsaye K, Chignell M: *Expert Systems for Experts.* New York, John Wiley, 1988, 223.
2. Harber P, Rothenberg LS: Controversial aspects of respiratory disability determination. *Semin Respir Med* 1985;7:257–269.
3. Winston PH: *Artificial Intelligence*, ed 2, Reading, MA, Addison-Wesley, 1981.
4. Negoita CV: *Expert Systems and Fuzzy Systems.* Menlo Park, CA, Benjamin-Cummings, 1985.
5. Dinman BD, Sussman N: Uncertainty, risk and the role of epidemiology in public policy development. *J Occup Med* 1983;25:511–516.
6. Sacks HS, Berrier J, Reitman D, et al: Meta-analyses of randomized clinical trials. *N Engl J Med* 1987;316:450–454.