Forest Science, Vol. 35, No. 2, pp. 364-387.

Development and Evaluation of an Expert System for Diagnosing Pest Damage of Red Pine in Wisconsin

DANIEL L. SCHMOLDT
GEORGE L. MARTIN

ABSTRACT. An expert system for diagnosing pest damage of red pine stands in Wisconsin, PREDICT, runs on IBM or compatible microcomputers and is designed to be useful for field foresters with no advanced training in forest pathology or entomology. PREDICT recognizes 28 damaging agents including species of mammals, insects, and pathogens, as well as two types of abiotic damage.

Two separate development tools (EXSYS and INSIGHT2 +) were used. Each employs a rule-based method for representing knowledge, which was obtained from the literature and from human experts in the fields of forest pathology and entomology. The pest-inference rule blocks, for each damaging agent, are based on the abduction model of diagnosis and include rules for eliminating a pest from further consideration, diagnosing a pest as certain, and compiling evidence in favor of a pest. Both development tools employ a backward-chaining control strategy; however, it was necessary to modify this approach by designing a special block of rules to approximate the mixed strategy used by the human experts. A logic and completeness rule block was also constructed to deduce facts omitted by the user and to minimize the need for questioning.

Input to PREDICT is obtained from pest damage reports containing specific information about stand/site conditions, tree symptoms, and signs. Diagnoses from PREDICT take the form of a list of one or more possible agents with corresponding confidence values. Actual and hypothetical test cases were used to refine the knowledge base, then a separate set of 20 actual cases was used as a basis for testing and evaluating the completed system. It was necessary to develop special procedures for refining and evaluating the system to accommodate the often vague and uncertain nature of pest damage information.

Two versions of PREDICT (developed with the EXSYS and INSIGHT2 + tools, respectively) were evaluated and compared with three recognized experts and two field foresters. No significant differences were found between the performances of PREDICT and the experts; however, PREDICT performed significantly better than the two for-

---

The authors are Research Forester, USDA Forest Service, Pacific Southwest Forest and Range Experiment Station, 4955 Canyon Crest Dr., Riverside, CA 92507, and Associate Professor, Department of Forestry, University of Wisconsin-Madison, Madison, WI 53706. This research was completed while D. L. Schmoldt was a Research Assistant, and candidate for the Ph.D. degree, at the Department of Forestry, University of Wisconsin-Madison. Research was supported by the College of Agricultural and Life Sciences, University of Wisconsin-Madison, and by the McIntire-Stennis Cooperative Forestry Research Program, Project 2817.

The authors gratefully acknowledge Stanley Banash, Ronald Giese, and Robert Patton for invaluable assistance in developing the PREDICT knowledge base; David Hall of the Wisconsin Department of Natural Resources, David Czysz and Fred Souba of Nekoosa Papers, Inc., and AlIan Munson of the USFS Northeastern Area State and Private Forestry for providing case records of pest problems in red pine stands; Robert Hess, William Krochmalski, Frank Morse, Allen Prey, and Shane Weber of the Wisconsin Department of Natural Resources for diagnosing the test cases; Robert Ford, Laura Merrill, Manfred Mielke, and AlIan Munson of the USFS Northeastern Area State and Private Forestry for evaluating the test diagnoses; and Robert Muetzelfeldt and an anonymous reviewer for their valuable comments and suggestions for improving the manuscript. Manuscript received April 18, 1988.

esters, even though they both have training and experience in forest pest diagnosis. It was concluded that PREDICT is able to improve the diagnoses of field foresters to a level comparable with recognized experts. FOR SCI. 35(2):364-387.

ADDITIONAL KEY WORDS. Artificial intelligence, personal computers, forest pathology, forest entomology, pest management.

FORESTRY APPLICATIONS OF EXPERT SYSTEMS are still relatively new, and little information has been compiled on methods for developing and evaluating expert systems in forestry, or on their effectiveness as tools for solving forestry problems. This paper addresses the need for such information, and describes the design, construction, evaluation, and test results of an expert system for diagnosing insect, disease, and other types of damage in red pine *(Pinus resinosa* Ait.) stands in Wisconsin.

Expert systems are computer programs capable of solving difficult problems at a level comparable to human experts. Their expertise is usually restricted to a very specific problem area, e.g., apple orchard management problems (Roach et al. 1985), or diagnosis and treatment recommendations for bacterial infections (Davis et al. 1977). Expert systems have been applied to problems in many fields, including electronics, engineering, law, manufacturing, mathematics, medicine, meterorology, physics, and space science (Hayes-Roth 1984). Forestry applications have also been reported (Rauscher and Cooney 1986, Schmoldt and Martin 1986, Kourtz 1987, Rauscher 1987, Reinhardt 1987, and White and Morse 1987), and potential applications in wood products manufacturing are described by Mendoza and Gertner (1988).

Expert systems evolved from techniques developed in the field of artificial intelligence, a branch of computer science. A variety of names have been used to identify expert systems, e.g., knowledge-based systems, inference systems, rule-based systems, and consultation systems. The particular name used is influenced by the preference of the developer and/or the structure of the knowledge in the system. General references on the structure and functioning of expert systems are provided by Hayes-Roth et al. (1983), Nau (1983), and Waterman (1986), to name a few. Discussions about the expert system approach and its problems and potentials for forestry applications are given by Schmoldt and Martin (1986), Mills (1987), and Schmoldt (1987a, 1987c).

The authors were motivated to explore the use of expert systems for pest diagnosis in forestry because of the success reported by Davis et al. (1977) with similar applications of expert systems for medical diagnosis. It was determined that the time and effort required to develop an expert system would necessarily y restrict this project to pest problems of a single species. Red pine was selected because of its economic importance in Wisconsin. Also, the number of different pests threatening the health of red pine stands is greater than that of any other commercial species in Wisconsin (Lindberg and Hovind 1983).

The objectives of this project were threefold. First, it was desired to develop an expert system for pest diagnosis of red pine that would be useful for field foresters with little formal training in forest pathology or entomology. Required observations of tree and stand conditions, signs, and symptoms should be easily obtained by the average field forester.

The second objective was to determine if the expert system can perform on a par with recognized experts, i.e., forest pathologists and entomologists with extensive experience diagnosing pest problems of red pine. And, finally, it was desired to determine if the expert system can significantly improve the diagnoses of field foresters (nonexperts).

DESIGN AND CONSTRUCTION OF THE EXPERT SYSTEM

Expert systems consist of a knowledge base containing the knowledge and experience required for the particular area of expertise. They also consist of programs to create and modify the knowledge base and to apply the knowledge to specific problems. In addition to the time and effort required to develop the knowledge base, earlier expert systems also required a substantial programming effort to develop the necessary algorithms and user interfaces. Since then, a number of expert system development tools (called "shells") have appeared on the market (Cooney 1986). These shells incorporate all the necessary programs in a ready-to-run software package.

The authors elected to make use of a shell so that this project's development effort could focus on design and construction of the knowledge base. In fact, two separate shell systems were used, EXSYS[1] and INSIGHT2+.[2] The authors decided to implement the pest diagnosis system under two different shells to see if the performance of the final system would be influenced in any way by the particular shell used. These two shells were selected because they seemed to be sufficiently powerful for the application, they are inexpensive (less than $500 at the time of the study), and there are some differences in their knowledge representation and control strategies.

The pest diagnosis system PREDICT *(Pinus resinosa* Expert Diagnostic Consultation Tool), was developed on an IBM XT personal computer using each of the two shells. Hence, there are currently two versions of PREDICT: PREDICT/EXSYS and PREDICT/INSIGHT2+.[3] The final systems operate on most IBM and compatible computers.

Both EXSYS and INSIGHT2+ use a rule-based method of knowledge representation, i.e., knowledge is represented in the form of IF-THEN rules, and decisions are made by drawing inferences from these rules and the characteristics of a specific problem. However, differences exist in the structure of the rules and the manner in which facts are represented. Brief descriptions of EXSYS and INSIGHT2+ follow.

EXSYS (VERSION 3.1)

EXSYS runs on an IBM PC/XT/AT or compatible computer with at least 256K of memory and one double-sided disk drive. However, all the rules and any external programs must be co-resident in memory, so more than 256K of memory maybe required (PREDICT/EXSYS requires 340K). Also, much better performance is obtained when two disk drives or a hard disk are available.

The EXSYS development package contains several programs for creating

[1] EXSYS is a product of Exsys, Inc., P.O. Box 75158, Albuquerque, NM 87194.

[2] INSIGHT2+ is a product of Level Five Research, Inc., 503 Fifth Ave., Indialantic, FL 32903. Level Five Research has recently been purchased by Information Builders, Inc., and INSIGHT2+ has been replaced by LEVEL FIVE Version 1, a higher performance product that was previously only available on VAX computers.

[3] Copies of PREDICT are available for distribution to interested parties. For prices and ordering information write to George L. Martin, Dept. of Forestry, University of Wisconsin-Madison, Madison, WI 53706.

and running a knowledge base. These programs include a rule editor and a run-time inference engine. The rule editor is used in the development process to create and modify decision rules in the knowledge base. Each decision rule contains up to five parts: IF and THEN parts, and optional ELSE, REFERENCE, and NOTE parts.

The IF part (also called the "antecedent") is a set of statements concerning facts about a particular problem. Facts are either deduced by the expert system or provided by the user and are represented in EXSYS as numbers, text, or qualifiers. A qualifier is a description of some object or condition that may have several different values at the same time. For example, a qualifier might be "the affected trees are" and its possible values might be

present in patches or groups,

scattered,

widespread,

dead or dying,

leaning or loose in the ground,

deformed.

For a particular pest diagnosis problem, "the affected trees are" may simultaneously have the values "present in patches or groups" and "dead or dying," or some other combination of values. It is assumed in EXSYS that all facts are known with certainty. There is no provision for uncertain information about a problem.

If all of the IF statements (antecedent conditions) are found to be true, the actions listed in the THEN part (also called the "consequent") of the rule are performed, otherwise the actions listed in the ELSE part, if any, are performed. Possible actions include assigning values to numeric, text, or qualifier variables, or calling an external program. Another type of action is the assignment of a probability or confidence value to one of the possible conclusions the expert system is trying to reach. In the PREDICT knowledge base, possible conclusions correspond to the damaging agents (Table 1). The probability or confidence value is not really a probability yin the strict sense of the word, but is a number (0-1) indicating the confidence that the conclusion is correct.

More than one rule may assign probabilities to the same conclusion, so it is necessary to select an appropriate formula for combining probabilities from different rules. From the several alternatives offered by EXSYS, the authors selected an incremental formula for independent probabilities as the most appropriate for the PREDICT knowledge base. With this formula, each time a new rule assigns a probability y, P, to a conclusion, the new probability, P_N, is calculated as

$$P_N = P_O + P(1 - P_O) \qquad (1)$$

where P_O is the old probability y assigned to the conclusion by previous rules. In this way, the probability score for a particular conclusion increases incrementally and approaches a value of 1 as evidence accumulates for that conclusion. The reader should note that this method of combining probabilities is not affected by the order in which probabilities are assigned by different rules. Shortliffe and Buchanan (1975) refer to these probabilities as *measures of belief,* and they accumulate their belief values using mathematics identical to Equation (l).

Some optional information may be appended to each rule in the form of a NOTE and REFERENCE. A NOTE provides information about the rule, its facts, and reasoning. The NOTE is displayed anytime the rule is displayed. A REFERENCE contains information about the sources of facts and reasoning in the rule. The REFERENCE is displayed only if the user specifically requests it.

The run-time inference engine, an EXSYS program that allows the knowledge base to be applied to specific problems, uses a control strategy known as "backward chaining." Under this strategy, each of the possible conclusions, in turn, is set as a goal, and all of the rules that infer this conclusion are evaluated in the order in which they appear in the knowledge base. If a rule cannot be evaluated because certain facts are unknown, the rule is temporarily suspended, and the program examines rules that might be able to deduce the unknown facts. If the program fails to deduce the necessary information, the user is asked to provide it. An example of backward chaining is presented in Figure 1.

It is important to note that backward chaining is a goal-oriented strategy, i.e., it selects rules that might be able to satisfy the current goal under investigation. Whether or not a selected rule actually succeeds (satisfies or contributes to the current goal) is determined by its antecedent conditions, consequent actions, and available facts.

The inference engine allows the user, at any time, to ask why some item of information is requested. When this is done, EXSYS displays the rule(s) currently being evaluated. The display provides some explanation of why

TABLE 1. Damaging agents recognized by the PREDICT expert system.

Mammals
Porcupines *(Erithizon dorsatum* L.)
Pocket gophers *(Geomys bursarius* [Shaw])
Meadow mice *(Microtus* spp.)
Insects
Red pine sawfly *(Neodiprion nanulus nanulus* Schedl)
European pine sawfly *(Neodiprion sertifer* [Geoff.])
Red-headed pine sawfly *(Neodiprion lecontei* [Fitch])
Pine tussock moth *(Dasychira pinicola* [Dyar])
Red pine needle midge *(Thecodiplosis piniresinosae* Kearby)
Saratoga spittlebug *(Aphrophora saratogensis* [Fitch])
Red pine shoot moth *(Dioryctria resinosella* Mut.)
European pine shoot moth *(Rhyacionia buoliana* [D.&S.])
White pine weevil *(Pissodes strobi* [Peck])
Allegheny mound ants *(Formica exsectoides* Forel)
Northern pine weevil *(Pissodes approximates* Hopk.)
Pales weevil *(Hylobius pales* [Hbst.])
Red turpentine beetle *(Dendroctonus valens* Lee.)
Root collar weevil *(Hylobius radicis* Buch.)
Root tip weevil *(Hylobius rhizophagus* Millers, Benjamin, & Warner)
White grubs *(Phyllophagus* spp.)
Bark beetles *(Ips pini* [Say])
Pathogens
Scleroderris canker *(Gremmeniella abietina* [Lagerb.] Morelet)
Tip blight *(Diplodia pinea* [Desm.] Kickx)
Needle cast *(Lophodermium pinastri* [Schrad. ex Hook.] Chev.)
Red pine shoot blight *(Sirococcus strobilinus* Preuss)
Armillaria root rot *(Armillaria mellea* [Vahl ex Fr.])
Red pine needle rust *(Coleosporium* spp.)
Abiotic agents
Winter injury
Needle droop

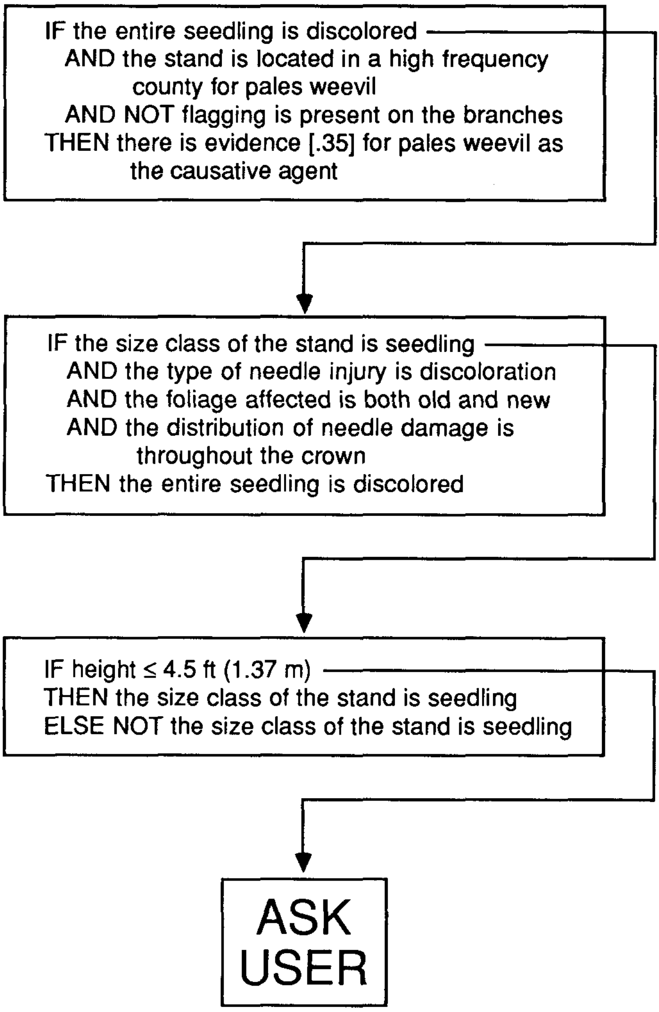

FIGURE 1. The interaction between rules and user input is illustrated in this example of backward chaining. The first rule is trying to establish evidence for pales weevil, but it has been temporarily suspended because it is not known if the entire seedling is discolored. The second rule in the chain is trying to deduce that the entire seedling is discolored, but it has also been suspended because it is not known if the size class of the stand is seedling. The third rule in the chain is invoked to determine if the size class is seedling or not, but it is suspended because the height of the stand is not known. Finally, because there are no rules to deduce the height of the stand, the user is asked to supply this information.

the requested information is needed. If several rules are listed, the information that has just been requested will appear in the IF part of the first rule displayed. One of the consequent actions of the first rule will provide facts for the IF part of the second rule, and so on. A chain of rules is displayed in which each rule's evaluation, except the first, is dependent on the evaluation of the preceding rule in the chain. This display sequence reflects the backward-chaining strategy described above.

After the list of possible conclusions has been exhausted and all inferences have been made, the inference engine displays a list of all conclusions that received a probability score greater than zero. The user can then request a display of all the rules used to infer a particular conclusion. This is helpful for understanding how the conclusion was reached.

INSIGHT2+ (Version 1.2)

Hardware and software requirements for INSIGHT2+ are similar to EXSYS. An IBM PC/XT/AT or compatible computer with at least 256K of memory is required. However, to use the full capabilities of this development package, 512K is recommended. Similar to EXSYS, any external programs that are called by a knowledge base must be co-resident in memory with the inference engine, but the entire knowledge base need not be memory resident. A minimum of two double-sided disk drives is required, but a hard disk is recommended.

The INSIGHT2+ development package contains a text editor, a knowledge base compiler, a run-time inference engine, and software to create and access database files. The text editor is used to create and edit the knowledge base using a special language called PRL (Production Rule Language). Before the knowledge base can be used by the inference engine, it must be compiled by the knowledge base compiler into an executable code. This process is similar to the edit-compile-run cycle typical of many programming languages.

All facts, hypotheses, final conclusions, and goals are treated identical y by INSIGHT2+. They are all considered as statements which can be described by one of the fact types. Any fact type may, optionally, have an associated confidence value in the range 0–100. Confidence values less than 50 (this default value may be reset by the system developer) are treated as measures of belief in a fact's falsehood. Values greater than or equal to 50 are measures of belief in a fact's truth. Four fact types are provided for describing problems. Numeric and text facts are identical to their counterparts in EXSYS. Object facts are very similar to EXSYS qualifiers. The fourth fact type, simplefact, is a simple statement that is either true, false, or has some confidence value in between. An example of a simplefact is "sandy soil is present in the stand. "

Rules in INSIGHT2+ are similar to those in EXSYS, and include IF, THEN, and ELSE parts. A backward-chaining control strategy similar to that of EXSYS is employed, and an incremental formula [Equation (l)] was selected for combining probabilities.

Extensive facilities are available for explaining to the user how a particular problem is being solved. At any time, the user can ask why some item of information is requested. This provides access to the menus that produce explanations. By default, the current rule and all rules dependent on it in the current chain can be viewed. This is identical to the display produced by asking "why" in EXSYS. Another menu allows the user to select from two displays of the line of reasoning. One display shows, in chronological order, all facts provided by the user and the inferred hypotheses. The other display shows the facts and hypotheses required for the determination of each goal pursued. A facts menu shows all facts in the current session, their values, and the origin of their values. Any of the facts supplied by the user can be changed any time during a session, and the problem can then be reevaluated by the inference engine (EXSYS provides a similar feature, but it is only available at the end of a session).

Damaging Agents

Damaging agents were selected for inclusion in the PREDICT expert system on the basis of three criteria: (1) the agent is economical] y important in red pine, (2) the agent is easily confused with economically important agents, or (3) the symptoms of the agent are so overt that it appears economically important even though it is not. Table 1 lists the agents selected in each of four categories: mammals, insects, pathogens, and abiotic agents.

Many mammals, including man, can severely damage red pine, causing significant mortality in small trees and greatly reducing the growth of larger ones (Wilson 1977). Of these, porcupines, pocket gophers, and meadow mice have had the greatest impact. Injury results from gnawing and chewing of the bark on the stem and major roots. The size and location of teeth marks help discriminate between these three mammals. Unlike many of the insects and diseases, visual identification of mammals provides little difficulty for the casual observer.

Insects comprise the largest group of damaging agents of Wisconsin red pine. Some insects injure red pine during both their adult and larval stages. Therefore, diagnosis must occasionally consider several variants of the same insect. Included in this category are defoliators, shoot and bud feeders, bark feeders, and root feeders.

Diseases of red pine are often more difficult to diagnose than insect injury. Symptoms are usually limited to changes in the appearance of the needles. Because the pathogens are fungi, signs are often small fruiting bodies that are difficult to distinguish outside of the laboratory. Knowledge of the specific needles affected is very important. For example, some diseases affect only the current year's needles, others only the previous year's needles, and still others affect both. The location of the affected needles on the tree is also important. Dark and moist conditions on lower portions of the tree are important for development of some diseases.

Trees can exhibit injury from their surrounding environment as well as from living organisms. Damaging environmental conditions may be present in the air, the soil, or a combination of the two. Temperature extremes, pollution, nutrient deficiencies, drought, and flooding are some of the conditions that may be injurious. Two physiological responses to environmental conditions are included in the abiotic category: winter injury and needle droop. Both are related to transpiration loss in the needles.

Conceptual Models of Pest Diagnosis

Although the transition from medical diagnosis to forest pest diagnosis is not a great conceptual leap, there are aspects of the forestry application that require special attention. In the medical setting, clinical data are usually obtained by trained specialists, and the remaining information is provided by reliable laboratory tests. In forestry, pest damage reports are usually completed by field foresters who are not highly trained in pest diagnosis, so important diagnostic information may be incomplete. Most often, only gross symptoms are identified, and laboratory data are rarely available. Also, foresters usually observe the pest problem at a single point in time and are unable to obtain a verbal history of the "patient's" illness. Ambiguities in forestry terminology further exacerbate the difficulty of obtaining reliable reports (Schmoldt 1987b). Because of the uncertainty and incompleteness associated with forest pest observations, thorough diagnoses often include more than one possible cause of stand damage, rather than produce a single, and possibly erroneous, conclusion.

To help provide the most consistent and complete data possible for use with PREDICT, a Red Pine Damage Report form was prepared. This form is a checklist of important diagnostic information ranging from general stand information and gross symptoms to more specific symptoms and signs.

Bearing in mind the problems associated with forest pest diagnosis, two

TABLE 2. Examples of different types of rules based on the abduction model of pest diagnosis. Numbers in brackets indicate the confidence values of the rules.

Rule 1:	Cause $\Rightarrow$ Effect
	IF root tip weevil is the damaging agent
	THEN the trees have branch flagging
Rule 2:	Evidence-Accumulation (Abduction)
	IF the trees have branch flagging
	THEN root tip weevil is the damaging agent [.3]
Rule 3:	Elimination
	IF the trees do not have branch flagging
	THEN root tip weevil is not the damaging agent [1]
Rule 4:	Certainty
	IF small roots have been chewed off
	THEN root tip weevil is the damaging agent [1]

conceptual models were used in the development of PREDICT. An abduction model was used to design and structure the decision rules, and a classification model was used for refinement of the knowledge base.

Abduction is the reverse of a cause-and-effect relationship. In the usual cause-and-effect implication, some pest *causes* stand damage, and its *effects* are observed as symptoms (Rule 1 in Table 2). This is the way pest information is presented in the literature and perceived by experts. In the course of describing a particular pest, its manifested symptoms are also described. However, in a diagnostic situation it is necessary to make inferences in the other direction, i.e., given certain effects, a cause must be inferred (Rule 2 in Table 2). This use of a cause-and-effect relationship in the reverse direction is referred to as *abduction* (Reggia et al. 1985). After many years of experience, human experts do this implicitly.

From elementary logic it is known that if a cause, p, implies an effect q ($p \Rightarrow q$), it is not necessarily true that the effect, q, implies p, but this is the type of inference that must be made in a diagnosis. It can be stated, however, that q provides evidence for p. Similarly, if r is another effect of p ($p \Rightarrow r$), it can be said that r provides additional evidence for p. The measure of evidence provided by the presence of an effect, q, and the implication $p \Rightarrow q$, is the confidence value assigned to an inference rule relating p and q in the knowledge base. This provides the basis for the evidence-accumulation rules used in PREDICT (Rule 2 in Table 2). In general, the relationship $p \Rightarrow q$ will not be a strict implication, i.e., q may not always be an effect of p. To accommodate this situation, it is necessary to reduce the confidence assigned to the decision rule relating p and q.

In Rule 2 (Table 2), the confidence value is determined by the amount of information conveyed by the presence of branch flagging, and also the strength of the cause-effect implication (Rule 1). Because branch flagging results from man y other agents, its information content is low and, despite the fact that the cause-effect implication is strong (i.e., root tip weevil attack always causes branch flagging), the overall result is a relatively small confidence value.

If $p \Rightarrow q$ is a strict implication, the logical contra-positive (NOT $q \Rightarrow$ NOT p) can be used to eliminate the cause, p, in the absence of its effect, q. This provides the basis for the elimination rules used in PREDICT (Rule 3 in Table 2). However, when using an elimination rule, a distinction must be made between a condition *not present* and a condition *not observed*. The observer must have given a thorough search for the symptom and failed to observe it before it can be stated that the symptom is not present. This stipulation reduces the possibility of erroneous elimination of an agent on the basis of questionable observations.

Another important use of the abduction model occurs when $p \Rightarrow q$ and q p, i.e., the implication is both ways $(p\ q)$. This means that the effect, q, occurs exclusively in the presence of p. Examples of this "if and only if" implication include symptoms and signs that are specific to a particular pest. In Rule 4 (Table 2), "small roots chewed off" carries a lot of information because it occurs exclusively during root tip weevil attack. Relationships of this type permit a definitive diagnosis and provide the basis for the certainty rules used in PREDICT.

The classification model of pest diagnosis can be viewed as the assignment of all possible observations of signs, symptoms, and stand conditions to particular agents (pests). Hence, associated with each pest is a set of possible damage reports that could arise from the occurrence of that pest (Figure 2). Depending on its diagnostic strength, a particular report maybe assigned to several pests, so there are varying degrees of overlap in the sets

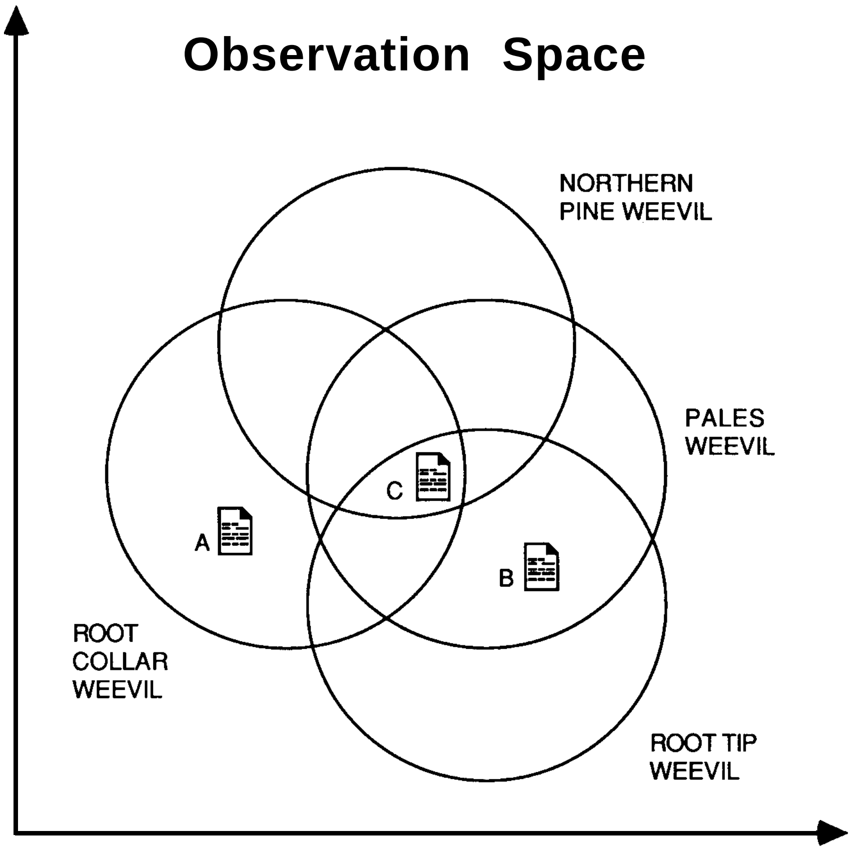

FIGURE 2. According to the classification model of pest diagnosis, all damage reports that might arise from the occurrence of a particular pest can be thought of as a circle or sphere in an "observation space" consisting of all possible signs, symptoms, and stand conditions. Circles from different pests overlap to varying degrees, so many damage reports are associated with more than one pest. Damage report A is highly diagnostic since it can arise from only one pest. Damage report B could arise from either of two pests, so it is less diagnostic. Damage report C has the least diagnostic value of the three reports, since it is associated with four different pests.

belonging to each pest. In practice, it is not possible to enumerate all of the sets; however, the concept of the classification model proved useful for refining the PREDICT knowledge base (see "Refinement" below).

Initial testing revealed groups of pests that were very similar in their effects, so initial refinement was aimed at discriminating between groups by identifying distinct sets of observations common to each group. Later refinement involved discriminating within groups by identifying sets of observations unique to individual pests.

Knowledge Acquisition and Rule Construction

Knowledge acquisition involves collecting and organizing the information and expertise necessary for problem solving and encoding it in a set of inference rules. Both literature and human experts were consulted for knowledge relevant to pest diagnosis of Wisconsin red pine.

Factors important to making a diagnosis were classified as stand/site conditions, tree symptoms, and signs. Specific stand/site conditions may be necessary for the occurrence of a particular pest, or they may predispose the stand to population buildup and subsequent attack. Included in the stand/site factors are stand location, soil and topographic descriptions, tree diameter and height, stand characteristics (e.g., crown closure), and presence of certain other tree species. Symptoms describe the injury resulting from the unknown agent. Different parts of the tree maybe affected, and the effect on a particular part of the tree (e.g., shoots) may differ for each pest. Signs are specific evidence that indicate the presence of a particular pest. Signs, such as larvae or cocoons, are highly diagnostic.

Because the literature is very descriptive and exhaustive, it proved a good source to begin enumerating the factors for each of the 28 pests. The above three categories were used to structure these lists. Several experts were then presented with the lists, and their comments were solicited. None of the experts was familiar with every pest; but among them, coverage of the pests was complete. The experts provided additional factors that were not identified in the literature, indicated which factors were more important, and eliminated some that were erroneous.

Stand location was identified by the experts as a very important factor for determining which pests to consider and which to rule out for a particular diagnosis. To best utilize this factor, data from an extensive search of pest damage records in Wisconsin were obtained from Ronald L. Giese.[4] Records spanning up to 30 years were obtained, and counties were designated, for most of the insects, as nonoccurrence (insect never observed), presence (observed at least once), low frequency, and high frequency (according to the frequency of observation). For some of the insects, the mammals, and the pathogens, limited records allowed counties to be designated only as nonoccurrence and presence. To minimize the potential for over-zealous elimination of a pest due to stand location, counties with no record of occurrence, but adjacent to a presence county, were also designated as presence counties.

Prior to constructing the rules, it was important to formulate a strategy for diagnosis, so several experts were asked to describe the procedures they employ. The y all use essential y the same strategy. Before arriving at a site, the expert already has a good idea of the pests most likely responsible for damage, because of familiarity with the pest problems common to a local

[4] Personal communication.

area (stand location) and, possibly, advance information from the reporting source. Location of the stand indicates which pests should be considered first and effectively narrows the search. With these pests in mind, trees are examined closely for symptoms. Observation of a particular symptom results in a focus on the pests related to this symptom, and a search is made for other symptoms associated with these pests. However, observation of certain other symptoms may change the focus of attention to different pests. This approach is described in the artificial intelligence literature as a mixed strategy (Klahr 1978), where intermediate conclusions are formed by forward chaining, i.e., applying rules in the knowledge base whose IF part is satisfied by known facts. These intermediate conclusions are then set as goals to be further investigated by backward chaining.

Neither EXSYS nor INSIGHT2+ employ a mixed strategy (both use backward chaining with only a limited forward-chaining capability); however, it was possible to incorporate the main aspects of the human experts' strategy in PREDICT. The strategy is implemented in three consecutive steps:

1. Identify potential agents by bark, roots, root collar, or soil examination, and pursue these agents in Step 2, else consider all agents.
2. Eliminate agents under consideration according to stand location, other stand/site conditions, and the description of injury.
3. Compile evidence for agents under consideration and not previously eliminated.

The first step utilizes a few preliminary symptoms to focus on particular pests for further investigation. All pests are considered if this fails. Focusing on a few likely agents at the outset can greatly increase the efficiency of a diagnosis. It was discovered that, using stand location alone, at most only half of the pests could be eliminated. This meant that at least 14 pests needed to be considered in every case. Preliminary testing indicated that PREDICT could spend a large amount of time investigating unlikely pests before the user was given an opportunity to provide highly diagnostic information. Because diagnoses can be performed quite quickly in situations where some careful observations have been made (e.g., examination of bark, roots, root collar, or soil), a better strategy is to inquire about such information at the beginning of a session. Only one question is needed to determine if valuable information is available. The alternative may be a lengthy, poorly directed consultation session.

Elimination similar to the human experts' approach occurs in Step 2. Step 3 attempts to accumulate favoring evidence for pests not eliminated. Generall y, these two steps occur sequentially for each pest. First, rules search for disfavoring evidence for a particular pest. If the pest cannot be eliminated on the basis of these rules, then favoring evidence is sought.

Construction of the initial knowledge base involved writing rules to infer each of the pests. Diagnostic factors in each rule were combined according to the four criteria described below.

The antecedent conditions in a rule must all occur within the same *time frame*. An observer in a stand will only see current symptoms, not symptoms that present themselves at different times or appear differently over time. Rules must exist to address the possibility y that observations may not be made at the time of injury and to consider different opportunities for observation.

Diagnostic content should vary between rules for each pest. Due to the incompleteness of information, strong diagnostic characters are not always reported, but it is still desirable to infer an agent to a lesser extent.

Observational skill must be given consideration. Observers vary in their ability to distinguish symptoms. If rules contain only conditions that require an experienced observer, they may never be applied when someone less experienced uses the system. Some rules were written only with very basic facts and others with symptoms and signs that require a trained eye for observation. Often, observational skill is related to diagnostic content; conditions that are more difficult to notice tend to provide stronger evidence for one pest over another.

The *level of predisposition* of a stand will determine how strongly other factors infer an agent. For example, the presence of sandy soil and jack pine (predisposing conditions) in a stand located in a high frequency county for root tip weevil is a stronger diagnostic statement than if sandy soil and jack pine are not present.

After the initial set of rules was constructed, experts were asked to assign confidence values to each rule to indicate how strongly they felt the premises imply the conclusion. Each rule was considered independently, and the weight assigned was dependent only on that rule's antecedent conditions. Rules that the experts did not consider important were revised or removed. Some new rules were also added at this time. Rules with identical premises are not independent, so the confidence values assigned to these rules must not total greater than one (Shortliffe and Buchanan 1975).

In addition to the pest-inference rules, it was necessary to construct special rules for implementing the diagnostic strategy, and rules for logic and completeness. Given certain conditions, the strategy rules select specific pests to be investigated, eliminate pests under consideration, and change the order in which pests are investigated. Logic and completeness rules were added to minimize questioning and supplement information omitted by the user. For example, the antecedent conditions in the following rule are asked in almost every session, so they represent information that is usually known to the system. From this, the system can deduce that "the affected trees are dead" without expecting or requiring the user to specifically state this fact:

IF—

the condition of the shoots is dead,

AND the condition of the buds is dead,

AND the type of needle injury is discoloration,

AND the foliage affected is both old and new,

AND the distribution of needle damage is throughout the crown,

THEN—

the affected trees are dead.

Some types of facts have a number of values, several of which may be true simultaneously. Instructions to the user indicate that all applicable values of a fact should be specified when requested. However, users do not always follow instructions, so rules were added to ensure that information obtained by the system is as complete as possible. For example,

IF—

the type of needle injury is death,

THEN—

the type of needle injury is discoloration.

Because of the large size of the PREDICT knowledge base (over 700 rules for the INSIGHT2+ version), it was necessary to impose a structure for organizing the rules and facilitating comprehension. The rules are grouped into blocks (Figure 3) according to their function, i.e., strategy rules, pest-inference rules for each pest, and logic and completeness rules. Each of the pest-inference rule blocks are further structured by subblocks for elimination rules, certainty rules, and evidence-accumulation rules. Within each elimination subblock, rules with very basic antecedent conditions appear and are evaluated first; more specific conditions of latter rules may not need to be asked. The conditions within each rule are ordered from basic facts to more specific facts, so the specific facts are investigated only if the preceding basic conditions are true. For example, most pest-inference rules

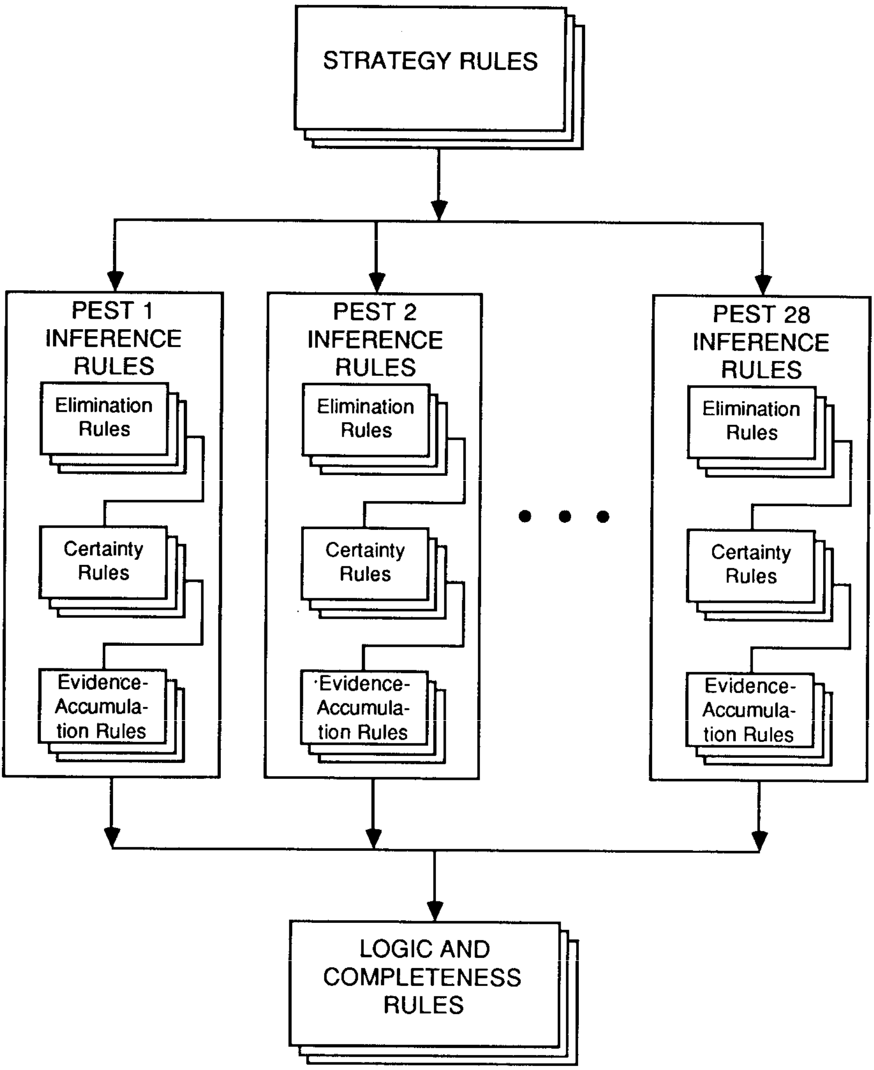

FIGURE 3. Structure of the PREDICT knowledge base. Arrows indicate the general order in which rules are evaluated. Strategy rules identify potential pests for further investigation, eliminate pests, and control the order in which pests are investigated. Pest-inference rules include elimination rules which compile disfavoring evidence, certainty rules which assign a confidence value of one for a pest, and evidence-accumulation rules which compile favoring evidence. Logic and completeness rules deduce facts omitted by the user and minimize the need for questioning.

have stand location (a condition which is always known) as the first antecedent condition. If it does not apply, then the remaining conditions are not checked. Again, this was done to minimize questioning.

REFINEMENT

Actual test cases provide realistic examples of problems and provide a basis for refining the knowledge base and evaluating the completed expert system. Case records of pest problems in red pine were obtained from the Wisconsin Department of Natural Resources, Nekoosa Papers, Inc., and the USFS Northeastern Area State and Private Forestry. Twenty of these test cases were set aside for use in evaluating the completed expert system. The remainder were used to refine the preliminary knowledge base.

The experts were asked to create some hypothetical cases to supplement the actual cases, but the hypothetical cases were used only for refining the knowledge base. They were not used for evaluating the completed system so as to avoid the possibility of biased results.

Weiss and Kulikowski (1984, pp. 151–155) describe an approach for refining rule-based inference systems using test cases for which a single, correct answer can be identified. Under this approach, a number of test cases are evaluated by the expert system, and the number of correct evaluations and false positives (conclusion reached, but incorrect) are tabulated for each of the possible conclusions. If a conclusion receives a low percentage of correct evaluations, the rules inferring this conclusion are generalized. Generalization is accomplished by removing some conditions from the IF part of a rule and/or increasing the confidence value assigned by the rule. If a conclusion receives a large number of false positives, rules are specialized. Specialization is accomplished by adding conditions and/or decreasing confidence values.

It was necessary to modify this approach for use with the PREDICT knowledge base, because pest diagnosis problems do not always have a single, correct answer. Given a particular problem, a single answer, multiple answers, or no answer may be appropriate. Hence, it is not always easy to distinguish correct evaluations and false positives.

The modified approach used to refine PREDICT is based on the classification model of pest diagnosis, discussed earlier in this paper. According to this model, as the diagnostic value of a damage report decreases, the number of pests that must be considered potentially responsible increases. With this in mind, several categories of final confidence scores for pests were defined: (1) 1.00 corresponds to absolute certainty, (2) 0.75 –.99 corresponds to very probable, (3) 0.50–.74 corresponds to possible, (4) 0.20–.49 corresponds to slightly possible. The following procedure was then employed for each of the 28 pests:

1. Run a test case that is highly diagnostic for the current pest. If the pest receives a confidence score of 1.00, or nearly so, go to Step 2; else, generalize the rules for this pest and repeat Step 1.
2. Reduce, slightly, the diagnostic value of the information provided by this test case, and run the case again.
   a. If the current pest receives the highest score (of all pests listed by PREDICT), go to Step 3; else, proceed to 2b.
   b. If all of the other pests are listed for good, diagnostic reasons, generalize rules for the current pest, specialize rules for pests that received a higher value, run the case again, and return to 2a; else, proceed to 2c.
   c. Perform the necessary specializations for each pest that should not have appeared in the list, run the case again, and return to 2a.
3. Further reduce the diagnostic value of the information provided, so that the

current pest receives a score in the next lower category (very probable, possible, or slightly possible). If other pests have substantially higher scores than the current pest, specialize rules for those pests and run the case again.
4. Repeat Step 3 for the remaining probability y categories.
5. Repeat Steps 1-4 using different test cases.

The reasoning behind this procedure is quite simple. At the outset, with highly diagnostic information for a particular pest, PREDICT is operating correctly if that pest receives a score of approximately 1.00. As the diagnostic information is reduced, other pests should also be listed as possible agents by PREDICT. However, none of the other pests should have a score much greater than the pest actually causing the damage. Diagnostic information is being reduced for the true pest, but no information is being added for other pests.

In order to implement this refinement procedure, it was necessary to provide some means of assessing the diagnostic value of information in Steps 2 and 3. As lists of diagnostic factors for each pest were identified during the early stages of knowledge acquisition, some measures of diagnostic importance were also specified by the experts. These measures provided a rough approximation of diagnostic value.

Refinement using this method of decreasing diagnostic value has some distinct advantages. Groups of pests with similar effects are identified early (Figure 4a), so the refinement effort can quickly focus on discriminating between pests within each group (Figure 4b). This within-group discrimination progresses with each iteration of the procedure. As other pests within a particular pest's group undergo refinement, overgeneralizations and overspecialization tend to be corrected. The procedure was found to be quite robust.

TESTING AND EVALUATION

After refinement of the knowledge base, an experiment was designed to compare the performance of the PREDICT expert system with that of recognized experts and field foresters. Although both versions of PREDICT use essentially the same knowledge base, there are some differences in operation that could affect performance. Therefore, both versions, PREDICT/EXSYS and PREDICT/INSIGHT2+, were tested.

Five human subjects were selected for the comparison, a forest pathologist, two forest entomologists, and two foresters. All of the subjects have work experience with both insect and disease problems. The two foresters attended training sessions in forest pest diagnosis, given by the Wisconsin Department of Natural Resources, but they do not have advanced training in entomology or pathology. None of the subjects was involved in any way with the development of PREDICT.

The 20 test cases selected for system evaluation ranged from easy to very difficult to diagnose. Also, the cases were selected to include many of the more common and economically important pests. None of the cases had been used during the refinement phase of the project. Each of the five human subjects was shown only the written damage report for each case, the same information that was input to the two versions of PREDICT. The subjects were instructed to work alone (consultations with other people were prohibited), assess each case with the aid of any reference materials they normally use, and identify all pests that could possibly be causing the stand damage. They were also asked to assign a numerical value (20–100%) to each pest indicating their confidence in that pest.

The diagnoses in this experiment could not be evaluated on the basis of a single, correct answer. Also, it is rare to find an individual skilled in both insect and disease diagnosis. So two evaluation teams, consisting of one entomologist and one pathologist each, were chosen to evaluate the diagnoses of the five human subjects and the two versions of PREDICT. None of the members of the evaluation teams was involved with the development of PREDICT.

The two evaluation teams worked independently and arrived at diagnoses

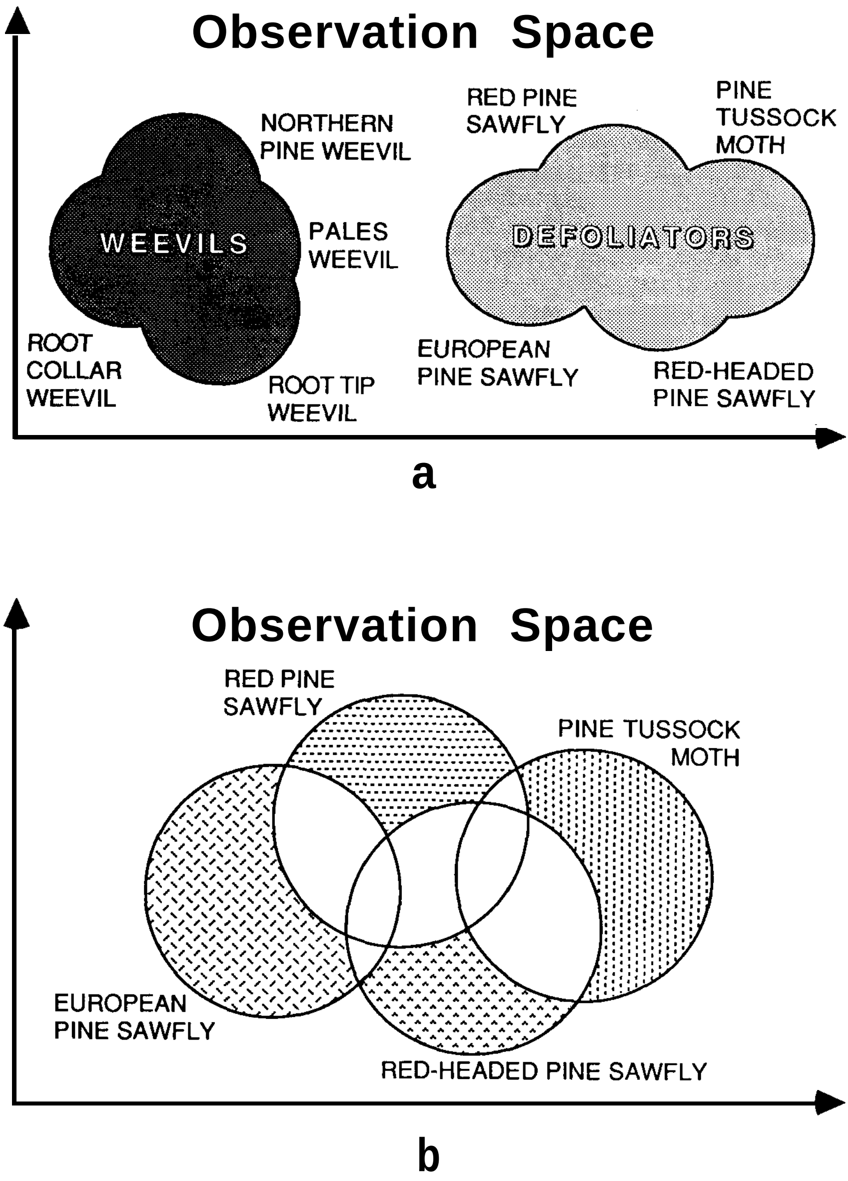

FIGURE 4. The classification model of pest diagnosis was very useful for refining the PREDICT knowledge base. Initial refinement (a) was aimed at identifying groups of pests with similar effects, e.g., the different weevils overlap considerably in their associated signs, symptoms, and stand conditions, as do the defoliators. Subsequent refinement (b) focused on discriminating within each group by identifying sets of observations (illustrated by the different shaded areas in b) unique to individual pests.

for each of the test cases. Each team was then shown the diagnoses from the human subjects and the two versions of PREDICT, but the diagnoses were scrambled so the evaluation teams did not know their sources. Also, there was no way the teams could determine if a diagnosis came from one of the human subjects or one of the versions of PREDICT.

Each evaluation team was asked to provide a comprehensive score, on a scale from zero (worst) to ten (best), for each diagnosis (noninteger values were permitted, but rarely used). This score is a subjective rating of the complete diagnosis, including an assessment of the appropriateness of the agents listed and their associated confidence values. Also, each evaluation team determined, for each diagnosis, if the assessment of the single most likely cause of stand damage was in agreement with that of the evaluation team.

RESULTS AND DISCUSSION

The results of the evaluation are summarized in Table 3. Disagreement between PREDICT and the evaluation teams, as to the most likely cause of stand damage, occurred in 6 of the 20 test cases; and in 4 of these cases the evaluation teams disagreed with each other. Most notably, both versions of PREDICT failed to diagnose Armillaria root rot and produced false positives for root collar weevil and bark beetles. Evaluation Team 2 disagreed with several of PREDICT’s sawfly diagnoses, but Team 1 did not. These particular test cases have since been analyzed by another entomologist, who concluded that valid diagnoses were made by PREDICT in all but one instance.

Chi-square tests were performed to identify significant differences between the numbers of diagnoses in agreement. The results are presented in Table 4. The differences between the versions of PREDICT and the experts are not significant, but PREDICT/EXSYS performed significant y better than both foresters, and PREDICT/INSIGHT2+ performed significantly better than Forester 2. There is no significant difference between the two versions of PREDICT.

With the exception of Team 2’s evaluation of PREDICT/INSIGHT2+, the standard deviations of the comprehensive scores (Table 3) are lower for the versions of PREDICT than for any of the human subjects, and in many instances they are significantly lower. This suggests that PREDICT performs at a more consistent level than the human subjects. Consistency was considered a desirable attribute during the construction of PREDICT and is one of the reasons for listing several potential pests when a single pest cannot be diagnosed with certainty.

Differences between mean comprehensive scores for each pair of subjects are presented in Table 5. No significant differences were found between the versions of PREDICT and the three experts. Both versions of PREDICT performed significantly better than Forester 2, but Forester 1's performance is not significant y different.

It is interesting to note that the mean difference in comprehensive scores between the two versions of PREDICT, although small (0.50), is statistically significant. The two versions of PREDICT are almost identical in their assessments of the single most likely cause of damage. However, PREDICT/INSIGHT2+ listed more potential pests (an average of 2.9 per case), and the evaluation teams tended to assign lower ratings for producing longer lists. While this seems valid from the view of experienced pest specialists,

TABLE 3. Total number of diagnoses in agreement, average comprehensive scores, and standard deviations from the 20 test cases.

	Evaluation Team 1			Evaluation Team 2		
		Comprehensive[b]			Comprehensive[b]	
Subject	Total no. in agreement[a]	Ave.	Std. dev.	Total no. in agreement[a]	Ave.	Std. dev.
Pathologist	16	7.20	3.22	14	6.40	3.08
Entomologist 1	16	7.75	2.73	14	6.35	2.68
Entomologist 2	16	7.40	2.98	13	7.00	2.73
Forester 1	12	5.60	3.56	13	6.50	3.07
Forester 2	9	5.05	3.43	8	5.05	2.98
PREDICT/EXSYS	18	7.45	2.50	15	7.05	2.56
PREDICT/INSIGHT2+	17	6.80	2.21	15	6.70	3.13

[a] The number of diagnoses in agreement is based on the subject's and evaluation team's assessment of the single most likely cause of stand damage in each case.

[b] The comprehensive score is a subjective rating (0, worst – 10, best) of the complete diagnosis, including an assessment of the appropriateness of the agents listed and their associated confidence values.

TABLE 4. Chi-square tests for significant differences between the numbers of diagnoses in agreement with the evaluation teams', for each pair of subjects.

	Pathologist	Entomologist 1	Entomologist 2	Forester 1	Forester 2	PREDICT/ EXSYS
Entomologist 1	0.00 NS					
Entomologist 2	0.06 NS	0.06 NS				
Forester 1	1.45 NS	1.45 NS	0.91 NS			
Forester 2	8.72 **	8.72 **	7.37 **	3.21 NS		
PREDICT/EXSYS	0.67 NS	0.67 NS	1.15 NS	4.01 *	13.65 **	
PREDICT/INSIGHT2+	0.29 NS	0.29 NS	0.62 NS	2.99 NS	11.85 **	0.08 NS

Note: NS = not significant, * = significant at 5% level, ** = significant at 1% level.

subsequent discussions with Wisconsin DNR pest management personnel[5] indicated that, in field applications, foresters would benefit more from the longer lists of potential agents. The two versions of PREDICT differ in their threshold values for eliminating pests from consideration. PREDICT/INSIGHT2+ is more conservative, and this results in longer lists. The elimination threshold can easily be changed to conform to the average list size of the human experts (1.7 pests per case).

Aside from differences in list size, the two versions of PREDICT produced almost identical diagnoses. The main difference in performance was found to be the amount of supplemental information available to users. INSIGHT2+ provides more extensive facilities for explaining how a particular diagnosis was deduced or why some item of information is being requested. Also, facts provided to the INSIGHT2+ version can be easily changed at any time during a diagnostic session. These differences in performance were not evaluated in the authors' experimental design (only the final diagnoses were evaluated), but the y are important considerations for developers and users of expert systems. For example, the authors found that it is much easier to analyze and refine the rules in PREDICT's knowledge base using the explanation facilities of INSIGHT2+.

A number of additional comparisons can be made from the perspective of system development; but, because of the rapid turnover of software releases, some of these comments may no longer be valid for the current versions of the two shells. The Production Rule Language (PRL) of INSIGHT2+ represents a rich programming environment with a relatively easy means for encoding knowledge bases. Because it is a very high level language, compilation time can be quite long for large knowledge bases. The PRL permits separation of a knowledge base into several different files, but there is no provision for separate compilation and subsequent linking; the entire knowledge base must be compiled each time changes are made.

EXSYS employs an interactive editor for knowledge base creation and execution, so editing and compiling are combined in a single step. Each rule is interpreted as it is entered or modified. This allows the knowledge base to be executed immediately, and the effects of a new rule can be observed almost instantly.

More detailed comparisons of the two shells are given by Schmoldt (1987a, 1988). It is worth noting that many of the new development tools now available include multiple knowledge representation schemes (Harmon et al. 1988). For larger and more complex applications, this additional flexibility is essential.

CONCLUSIONS AND SUGGESTIONS FOR FURTHER RESEARCH

Through use of the expert system approach, the authors were able to capture the knowledge and experience of a few human experts in the field of pest diagnosis of red pine. The system was designed to operate with the same kinds of information typically recorded by field foresters on pest damage reports. Hence, it can be used by foresters with no advanced training in forest pathology or entomology.

Test results indicate that the PREDICT expert system diagnoses pest damage as well as recognized experts with extensive experience, given the same information about stand damage. Further, PREDICT performed significantly better than the two field foresters, even though they both have some

[5]Personal communications with David H. Hall and Allen J. Prey.

TABLE 5. Mean difference in comprehensive scores between subjects (column subject minus row subject). The significance level of each mean difference, based on a paired t-test, is shown in parentheses. Comprehensive scores were combined from both evaluation teams.

	Pathologist	Entomologist 1	Entomologist 2	Forester 1	Forester 2	PREDICT/EXSYS
Entomologist 1	−0.25 NS (.70)					
Entomologist 2	−0.40 NS (.52)	−0.15 NS (.72)				
Forester 1	0.75 NS (.23)	1.00 NS (.054)	1.15 * (.044)			
Forester 2	1.75 ** (.0088)	2.00 ** (.0001)	2.15 ** (.0004)	1.00 NS (.058)		
PREDICT/EXSYS	−0.45 NS (.42)	−0.20 NS (.72)	−0.05 NS (.92)	−1.20 NS (.072)	−2.20 ** (.0005)	
PREDICT/INSIGHT2+	0.05 NS (.93)	0.30 NS (.61)	0.45 NS (.38)	−0.70 NS (.28)	−1.70 ** (.0065)	0.50 * (.018)

Note: NS = not significant, * = significant at 5% level, ** = significant at 1% level.

training and experience in forest pest diagnosis. It can be concluded, then, that PREDICT is able to improve the diagnoses of field foresters to a level comparable with recognized experts. Test results also indicate that PREDICT is very consistent and should prove reliable over a wide range of pest damage cases.

PREDICT, like all expert systems, should continue to undergo testing and refinement. Ease of use has already been enhanced by incorporating on-line definitions of most of the terms used in describing pest damage. Eventually, a complete description of each pest will be included. Other, less frequent types of damage (e.g., frost and drought), will also be added. Plans are currently underway to add a knowledge base for treatment recommendations. This component would take the output from the diagnostic system and then incorporate stand management objectives and economic considerations to produce one or more treatment prescriptions. With continuing advances in computer hardware and software, graphic displays of many of the symptoms could be presented during a diagnostic session. Enhancements like these would increase the educational value of the system as well.

Successful development of the PREDICT expert system suggests that similar systems can and should be developed for other tree species and other regions. An entire library of expert systems for forest pest diagnosis could be created. However, developers will find, as the authors did, that knowledge acquisition is the most important, as well as the most difficult and time-consuming, aspect of expert system development. New techniques must be found to make this process more efficient. Future efforts might explore the use of inductive learning systems. These have been used extensively in agricultural expert system development, where large numbers of test cases are available. An induction system creates discriminant rules on the basis of a large repertoire of examples. A preliminary set of rules constructed in this manner can serve as a good starting point for further revisions and refinement. In applications where test cases are abundant, a combination of inductive and direct knowledge acquisition could be a very effective approach.

LITERATURE CITED

COONEY, B. A. 1986. Representing knowledge through expert systems. Compiler 3(4):5-9.

DAVIS, R., B. G. BUCHANAN, and E. H. SHORTLIFFE. 1977. Production rules as a representation for a knowledge-based consultation program. Artif. Intell. 8: 15–45.

HARMON, P., R. MAUS, and W. MORRISSEY. 1988. Expert systems: Tools and applications. Wiley, New York. 289 p.

HAYES-ROTH, F. 1984. Knowledge-based expert systems. Computer 17(10):263-273.

HAYES-ROTH, F., D. A. WATERMAN, and D. B. LENAT. 1983. Building expert systems. Addison-Wesley, Reading, MA. 444 p.

KLAHR, P. 1978. Planning techniques for rule selection in deductive question-answering. P. 223-239 *in* pattern-directed inference systems, D. A. Waterman and F. Hayes-Roth (eds.), Academic Press, New York.

KOURTZ, P. 1987. Expert system dispatch of forest fire control resources. AI Applic. Natur. Resour. Manage. l(l):1–8.

LINDBERG, R. D., and H. J. HOVIND. 1983. A strategic plan for Wisconsin forests. Wisconsin Dep. Natur. Resour. Bur. of For. 121 p.

MENDOZA, G. A., and G. Z. GERTNER. 1988. Expert systems: A promising tool in wood products manufacturing. For. Prod. J. 38(2):51-54.

MILLS, W. L., JR. 1987. Expert systems in forest resources management. P. 349–362 *in* Proc. 1985 Symp. on Systems Analysis in Forest Resources, P. E. Dress and R. C. Field (eds.), Georgia Cent. for Cont. Educ., Athens, GA.

NAU, D. S. 1983. Expert computer systems. Computer 16:63-85.

RAUSCHER, H. M. 1987. Expert systems in natural resources management. Compiler 5(5):19-27.

RAUSCHER, H. M., and T. M. COONEY. 1986. Using expert-system technology in a forestry application: The CHAMPS experience. J. For. 84(3): 14– 17.

REGGIA, J. A., et al. 1985. Answer justification in diagnostic expert systems—Part I. Abductive inference and its justification. IEEE Trans. on Biomed. Eng. 32:263-267.

REINHARDT, E. 1987. An expert system for designing tire prescriptions. P. 223–225 *in* Proc. Symp. on Wildland Fire 2000, J. B. Davis and R. E. Martin (eds.), USDA For. Serv. Gen. Tech. Rep. PSW-101.

ROACH, J. W., et al. 1985. POMME: A computer-based consultation system for apple orchard management using Prolog. Expert Syst. 2:56–69.

SCHMOLDT, D. L. 1987a. Evaluation of an expert system approach to forest pest management of red pine *(Pinus resinosa).* Ph.D. diss., Univ. Wisconsin-Madison (Univ. Microfilms DA8708112). 211 p.

SCHMOLDT, D. L. 1987b. Implementing design considerations in a forest pest diagnosis expert system. *In* Proc. First Annu. Conf. on INSIGHT/LEVEL5 Expert System Applications, USERS.PRL, Ramsey, NJ. 15 p.

SCHMOLDT, D. L. 1987c. Using rule-based inference systems for forestry applications. P. 402–417 *in* Proc. 1985 Symp. on Systems Analysis in Forest Resources, P. E. Dress and R. C. Field (eds.), Georgia Cent. Cont. Educ., Athens, GA.

SCHMOLDT, D. L. 1988. PREDICT: An expert system for pest diagnosis. Compiler 5(5):29-41.

SCHMOLDT, D. L., and G. L. MARTIN. 1986. Expert systems in forestry: Utilizing information and expertise for decision making. Comput. Electron. Agric. 1 :233–250.

SHORTLIFFE, E. H., and B. G. BUCHANAN. 1975. A model of inexact reasoning in medicine. Math. Biosci. 23:351-379.

WATERMAN, D. A. 1986. A guide to expert systems. Addison-Wesley, Reading, MA. 419 p.

WEISS, S. M., and C. A. KULIKOWSKL 1984. A practical guide to designing expert systems. Roland and Allanheld, Totowa, NJ. 174p.

WHITE, W. B., and B. W. MORSE. 1987. ASPENEX: An expert system interface to a geographic information system for aspen management. AI Applic. Natur. Resour. Manage. 1(2):49-53.

WILSON, L. F. 1977. A guide to insect injury of conifers in the Lake States. USDA For. Serv. Agric. Handb. No. 501. 218p.