PLANT/DS: AN EXPERT CONSULTING SYSTEM FOR THE DIAGNOSIS OF SOYBEAN DISEASES

by

R. S. Michalski
J. H. Davis
V. S. Bisht
J. B. Sinclair

Plant Diseases and Proceedings of the European Conference on Artificial Intelligence, Orsay, France, pp. 133-138, July 12-14,1982.

ECAI - 82

1982 EUROPEAN CONFERENCE ON ARTIFICIAL INTELLIGENCE

12 - 14 JULY 1982, ORSAY, FRANCE

CONFERENCE PROCEEDINGS

PLANT/ds: AN EXPERT CONSULTING SYSTEM FOR THE DIAGNOSIS OF SOYBEAN DISEASES

R. S. Michalski and J. H. Davis
Department of Computer Science

V. S. Bisht and J. B. Sinclair
Department of Plant Pathology

University of Illinois at Urbana-Champaign
Urbana, Illinois 61801-2987

ABSTRACT

PLANT/ds is an experimental consultation system for advising farmers and other users on the diagnosis of soybean diseases common in Illinois. It contains an encoded knowledge of the symptomatic properties of the diseases, and formulates advise in response to the information provided to it by answering questions appearing on the computer terminal. The computer terminal is sensitive to touch, so that questions can be answered by touching an appropriate place on the screen.

On request, the system can provide an explanation of its operation and justify any of its advice by tracing the sequence of steps which lead to the advice. A unique feature of the system is that it uses two types of decision rules: 1) the rules representing experts diagnostic knowledge, and 2) the rules obtained through inductive learning from several hundred cases of disease. Experimental testing of the system has indicated a high level of correctness of the system s advice (in an experiment involving a few hundred cases, approximately 98% of the diagnoses were correct).

1. INTRODUCTION

Computer databases that store facts and numerical data on a given subject have been developed in many fields. These systems allow a user to easily retrieve information stored in them, but leave all the decisions about its use and interpretation to the user. Except for trivial cases, such systems are unable to answer any question for which there is no stored answer. Although they will continue to be of great value for many applications, new types of information systems are now being developed, called knowledge-based systems or expert systems.

An expert system contains a "knowledge base" and an inference mechanism able to conduct formalized reasoning. By relating the contents of the knowledge base to the information supplied by a user's answers to system formulated questions, the system infers the most recommended action in any particular situation. The knowledge base includes factual data (as in a data base) and decision rules that represent the general knowledge of the given subject (e.g., the diagnostic rules linking symptoms with diseases).

A typical form of a decision rule is:

If CONDITION then DECISION with confidence α (1)

The CONDITION stands for a list of elementary conditions characterizing a situation or an object to which the rule is applied (e.g., a diseased plant). The DECISION stands for specific advice or action, which this rule indicates when the CONDITION is satisfied. The parameter α expresses the strength of confidence in the DECISION when the CONDITION is fully satisfied ($0 \leq \alpha \leq 1$). When $\alpha = 1$, the confidence is maximal. When $\alpha = 0.5$, the rule states that the DECISION may be correct only half of the time. Thus, the rule format permits one to express the conditional knowledge of experts, and also the expert's confidence or lack of confidence in this knowledge.

In general, the CONDITION part of the rule may be only partially satisfied. For example, the condition: precipitation = above normal is only partially true if the precipitation was just above average. The degree to which the condition is true is expressed by a parameter $\gamma(0 \leq \gamma \leq 1)$. If $\gamma < 1$, then the confidence in the DECISION is calculated as a function of both α and γ. The way this function is calculated depends on the so-called evaluation scheme, described in section 2.

The DECISION part of a rule in a knowledge base may be an assignment of the status 'TRUE' to some conditions which are in the CONDITON part of another rule. Consequently, a satisfaction of one rule may cause a satisfaction of another rule(s), etc., and in this fashion the system can perform a chain of inferences.

Most of the expert systems developed to date are still in the experimental phase, and typically address some relatively narrow but important practical problem. For example, MYCIN was developed to advise doctors on the antibacterial therapy (Shortliffe 1976, Davis 1976), INTERNIST for providing consultation on the diseases of internal medicine (Myers and Pople 1977) and CASNET - on the glaucomas (Weiss et al., 1978). In chemistry, the DENDRAL system determines the molecular structures of complex organic chemicals from mass spectrograms and related data (Buchanan & Feigenbaum 1978). Another system, PROSPECTOR, provides consultation about mineral deposits (Duda et al. 1979). Some current aspects of research on expert systems are described in Michie (1979 and 1980).

System PLANT/ds was developed at the University of Illinois at Urbana-Champaign to provide consultation on the diagnosis of soybean diseases. The system is a part of a more general system PLANT designed to advise users about the diagnosis and decision making regarding crop diseases and damages due to insects. The following sections describe the knowledge base containing diagnostic rules, the method for using the rules to determine diagnostic advice, and results of experimental testing of the system.

2. REPRESENTING DIAGNOSTIC KNOWLEDGE

The diagnostic knowledge of experts is represented in the form of decision rules which specify all conditions indicating each disease. Advantages of such rule representation are that it is relatively easy to comprehend all the conditions indicating a diagnosis and to correct or refine knowledge represented in this form. It is also easy to extend the knowledge base by adding new rules, and to explain the inference process leading to a given diagnosis.

2.1 Specification of Descriptors

The first step toward building the knowledge base for PLANT/ds was to select variables (called descriptors) which provide useful characterizations of plants and their environment for diagnosing the considered diseases. The choice of a descriptor depended upon the relevency of a descriptor to the problem (diagnosing soybean diseases), and also on the ease of reliably determining its value for any diseased plant. In this study we made the assumption that a typical grower with no special training in plant pathology and with no special tools (like a microscope), should be able to determine the value of any descriptor.

This work was supported in part by a grant from the United States Department of Agriculture No. 321512344 and in part by a grant from the National Science Foundation No. 79-06614.

Each descriptor was assigned a value set (domain) which specifies all possible values the descriptor may take for any diseased plant. In determining such value sets it is important to avoid excessive precision, i.e., to limit the set of values to only those which may make a difference in diagnostic decisions. For example, the values of the descriptor "Condition of Leaves" were chosen simply as "Normal" and "Abnormal". When more specific information is needed on a given subject, other descriptors are used. For example "Leaf Spots" may have the values "Absent", "Present", "With Yellow Halos" or "Without Yellow Halos". In total, 41 descriptors were selected to characterize the diseased soybeans and the relevant characteristics of their environment. The diseases considered in the pilot study included 19 of the most common soybean diseases in Illinois.

The value sets of some descriptors include the special value "does not apply", which indicates that the descriptor can be irrelevant in some situations. For example, when the "Condition of Leaves" is "Normal", then all descriptors describing leaf abnormalities have the value "Does Not Apply".

Relationships between particular descriptors uch as above, impose restrictions on the description space, defined as the set of all theoretically possible combinations of descriptor values. In the process of determining a diagnosis, the system does not allow these restrictions to be violated. For example, if "Condition of Leaves" had the value "Normal" then "Leaf Spots" could not have the value "Present with Yellow Halos," as this would be an illegitimate combination.

The relationships between the values of the same descriptor are also taken into consideration. Depending on the type of this relationship, a descriptor is either nominal, linear or structured. Nominal (categorical) descriptors assume that no relationship exists between their values. Condition of Leaves (Normal-Abnormal), Presence of Hail (Yes-No) and Seed Shriveling (Absent-Present) are examples of nominal descriptors. The linear descriptor has values which have a natural linear order (like numbers), i.e., in any pair of distinct values, there is a smaller and a larger value. "Time of Occurrence" and "Number of Years Crop Repeated" are examples of linear descriptors. When more complicated relationships have to be represented, structured descriptors are used. While values of a linear descriptor can be placed along a line, structured descriptors have values that are nodes of a hierarchy (a tree structure). "Leaf Spots" is such a descriptor. The configuration of its values (Fig. 1) specifies that the values "With Yellow Halos" and "Without Yellow Halos" are considered to be a special case of the value "Present". Thus, if a rule stated that a disease would have "Leaf Spots" present and the user of the system said the field had leaf spots with yellow halos, that part of the rule would be satisfied.

Figure 1: The Hierarchy of Values for Leaf Spots

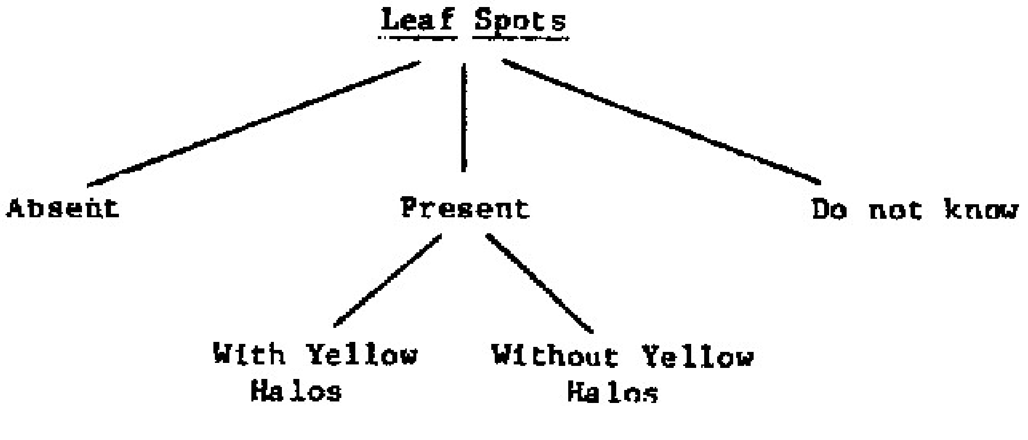

2.2 Relational statement: Basic Building Block of Decision Rules

A relational statement (or a selector) is an elementary condition, which specifies the scope of values which a descriptor is allowed to take on in order to satisfy a decision rule. In its simplest form a selector states that a given descriptor should take only one of its possible values, as in [Condition of Leaves = Abnormal] (square brackets always surround a selector). If a user indicates that the specimen is "Abnormal", then this selector is satisfied and assigned the evidence degree "1". If the user indicated that the leaves were "Normal", then the selector would be "not satisfied" and assigned the evidence degree "0".

In a more general case, a selector may allow a variable to take more than one value, for example, [Canker Lesion Color = Brown v Tan]. In this case the selector is satisfied if the user indicates either "Brown" or "Tan" as the "Canker Lesion Color", and is not satisfied otherwise. For linear descriptors a range of values can be specified, instead of listing the values individually, for example:

[Time of Occurrence = June..September]
or
[Precipitation = Normal]

Suppose that an expert wants to indicate that Bacterial Pustule is most likely in August, less likely to occur in July or September, unlikely in June, and quite unlikely in any other month. To express such information, a more general form of the selector is used, namely, a weighted selector, which allows one to specify the degree of evidence associated with each value of the descriptor. The evidence degree may range between 0 and 1, where "0" indicates no evidence and "1" - the maximum evidence. Suppose 0.8 represents evidence value "less likely", 0.6 - "unlikely", and 0.2 - "quite unlikely".

A weighted selector which represents the above stated information about Bacterial Pustule would be:

[Time of Occurrence = August: 1; July v September: 0.8; June: 0.6; Else 0.2]

There is another equivalent form for expressing a weighted selector. It uses a behavior indicator and a weight function. In this form, the above selector would be expressed: [Time of Occurrence : @ w], where

$$w = \begin{cases} 1, \text{ if August} \\ 0.8, \text{ if July or September} \\ 0.2, \text{ otherwise.} \end{cases}$$

The symbol @ is the behavior indicator, which tells that the evidence degree has the maximum value somewhere in the middle of the range "Time of Occurrence," and decreases for values on both sides of the middle point.

The weight function w is stored separately form the rules. It can be an arithmetic expression evaluated each time when it is needed. The concept of a weight function is similar to the concept of a set membership function in the fuzzy set theory (Zadeh, 1974).

Other possible symbols for the behavior indicator are v (minimum in the middle), ↑ (proportionality) and ↓ (inverse proportionality). The behavior indicator is used only with linear descriptors.

2.3 Conjunctive Statements: Complexes

A very common form of a decision rule is one which states that several conditions must be simultaneously satisfied in order to support a DECISION. The CONDITION part in such rules is the logical product (AND, conjunction) of selectors, called a conjunctive statement or a complex. The complex is expressed by either concatenating all the selectors involved (for implicit conjunction) or by joining them by the symbol "&". The latter way is useful when a complex takes more than one line of text. For example a complex describing Purple Seed Stain is:

[Time of Occurrence=September v October]&
[Condition of Seed=Abnormal][Seed Discoloration=Purple]

When all the selectors above are satisfied, then the evidence degree has the maximum value. Suppose, however, that only the last two selectors are satisfied and the first one is not, because the time of occurrence is August, rather than September. Using the classical interpretation of conjuction, if one condition is not satisfied, the whole complex would not be satisfied. Since August is 'close' to September, then it is clear that such an interpretation would be too rigid. A more flexible interpretation is needed. One way of handling this problem is to change the selector into a weighted selector specifying weights for each value. In this case, one would add August to the selector above with a weight slightly smaller than 1.

Another solution is to specify a default weighting function for selectors involving linear descriptors. In this case an (unweighted) selector would be evaluated as if it were a weighted selector, with a standard weight function of "bell" form, which assumes the maximum value for values specified in the selector, and continuously decreasing on both sides. Since August is "close" to September, the evidence degree would be slightly less than '1'.

A problem now arises of how to combine the degrees of evidence provided by each selector into the degree of evidence provided by the whole complex. PLANT/ds uses three techniques (evaluation schemes) for conjunction:

- PROD: The evidence degree of a complex is the arithmetic product of evidence degrees provided by its selectors,
- MIN: The evidence degree of a complex is the minimum of the evidence degrees provided by its selectors,
- AVE: The evidence degree of a complex is the average of the evidence degrees provided by its selectors.

These functions satisfy the relation PROD $\leq$ MIN $\leq$ AVE, i.e., given any two degrees of evidence, their product (PROD) will always be less than or equal to their minimum (MIN), and MIN less than or equal to their average (AVE). Theoretical and experimental studies show that the performance of each technique depends on the particular problem under consideration. The only reliable way to determine which technique to use seems to be experimentation. Therefore, we equipped the system with the ability to use any of the three schemes and the final interpretation scheme was chosen after complete testing of the system.

2.4 Disjunctive Statements

Combining selectors into complexes may not be sufficient for expressing complicated relationships between symptoms and diseases. Such a case occurs when two or more different combination of symptoms are associated with the same disease. For example, a soybean plant has brown spot when either small leaf spots without yellow halos or large spots with yellow halos are present. This is expressed using the logical sum ("OR", disjunction) of two complexes, each describing one of the alternatives:

[Leaf Spot Size $\geq$ 1/16"][Leaf Spot=With Yellow Halos]
v
[Leaf Spot Size < 1/16"][Leaf Spot=Without Yellow Halos]

The symbol V denotes logical 'OR'. It is assumed that the conjunction of selectors is always evaluated first, before the disjunction.

A decision rule whose condition part is a disjunction of complexes is said to be in disjunctive normal form (DNF). This form allows one to express any possible logical relationship (although not always in a concise way). Due to this generality and also to the simplicity of rule interpretation by humans as well as by computers, the DNF is an attractive format for representing the condition part of rules. One of the two sets of diagnostic rules used in PLANT/ds is in DNF.

The set of DNF rules was created by a general purpose inductive learning program, called AQ11 (Michalski & Larson, 1975). Such a program takes in as data examples of decisions (in this case it used 340 diagnoses made by plant pathologists) and from them creates decision rules for each class (in this case, a disease). Here is an example of a decision rule derived by this program:

[Plant Stand=Normal][Precipitation>Normal][Seed=Abnormal]
[Severeity=Minor][Plant Height=Normal][Seed Size=Normal]&
[Leaf Spots=No Yellow Halo][Seed Discoloration=Present]&
v
[Condition of Leaves=Normal][Seed=Abnormal]&
[Seed Size=Normal]
::> [Soybean Disease = Purple Seed Stain]

The second set of rules used in PLANT/ds was created by formally representing plant pathologists knowledge about relationships between symptoms and diseases. These expert-derived rules were more complex than the inductively-derived DNF rules, and therefore required more advanced formalism for their representation. The formalism is described in the following subsections. The complete set of expert-derived and inductively derived diagnostic rules is given in (Michalski & Chilausky, 1980).

As in the case of conjunction, there is also more than one way to interprete disjunction. Two methods were tested:

- MAX: The evidence degree of a disjunction of complexes is the maximum of the evidence degrees of the complexes,
- PSUM: (Probabilistic sum) The evidence degree of the disjunction of two complexes is $a + b - ab$, where a and b are evidence degrees of the two complexes. The evidence degree of the disjunction of more than two complexes is computed by the repeated application of the above rule.

2.5 Implicative Statements

From the formal viewpoint, any logical condition can be expressed in the DNF form. When one wants to express in this form, however, the diagnostic processes of plant pathologists, the DNF rules may be very clumsy, and may have no direct relationship to human descriptions. An important additional construct which facilitates expressing expert's descriptions is the implicative statement. An implicative statement is used when one wants to state that if some condition is present, then some other condition must be present. If the first condition is not present, then the other condition is irrelevant. There are many instances of implicative statements in the soybean diagnostic rules, for example, the rule for Downy Mildew contains the implicative statement:

[Time of Occurrence = September v October]=>
[Seed Mold Growth = Present]

This condition states that if the disease is occurring in September or October, then the seeds should appear moldy, otherwise seed mold growth is irrelevant. Implicative statements are evaluated by evaluating the logically equivalent disjunctive statement:

$$\sim r_1 \vee P_1$$

where $\sim P_1$ (negation of P_1) is evaluated as $1-d(P_1)$, and $d(P_1)$ is the evidence degree of P_1. The above means that if the condition P_1 is true ($d(P_1)=1$), then the evidence degree $d(P_1 \Rightarrow P_2) = d(P_2)$, which agrees well with our intuition. However, if P_1 is not true ($d(P_1)=0$) then the evidence degree $d(P_1 \Rightarrow P_2) = 1$. In this case we observe a difference between the formal interpretation and an intuitive one, in which we would simply ignore the whole implicative statement. A complex is considered to be a special case of an implicative statement in which the condition before the "=>" is always 'true', and therefore can be omitted together with the implication sign.

2.6 Linear Modules

In describing a disease, it is sometimes important to express the idea that certain groups of symptoms are more important for a diagnosis than other groups. For example, when diagnosing "Downy Mildew" the important symptoms are abnormal leaves, leaf spots without yellow halos, mildew growth on the lower leaf surface, and abnormal seeds with mildew growth, and time of occurrence in September or October. If any of these conditions are not present, then most probably Downy Mildew is not the problem. On the other hand, the conditions such as premature defoliation and presence of leaf malformation are confirmatory, but not crucial for the diagnosis. To expresss such relations, a construct called the linear module is used. A linear module has the form

$$q_1 \cdot C_1 + q_2 \cdot C_2 + q_3 \cdot C_3 + \ldots$$

where C_1, C_2, C_3, ... are conditions of the form considered so far (e.g. a complex, implicative statement or a product of implicative statements) and q_1, q_2, q_3 ... are coefficients indicating the relative significance of the conditions. The operator "$\cdot$" is interpreted as the arithmetic product of a coefficient q_i and the evidence degree of a C_i. A linear module is a very flexible form for expressing descriptions of diseases or decision processes in general.

2.7 Decision Rules

A general format of decision rules is:

$$\text{LINEAR MODULE} ::> \text{DECISION} : \alpha$$

where α is the degree of certainty that DECISION is correct if the conditions specified by the LINEAR MODULE are completely satisfied. The overall evidence degree for the DECISION is computed as the product of α by the evidence degree provided by the LINEAR MODULE.

In our study we used two-part linear modules:

$$q_s \cdot C_1 + q_c \cdot C_2$$

where coefficients q_s and q_c represented the relative importance of the significant (C_1) and confirmatory (C_2) part, respectively ($q_s + q_c = 1$). In experiments described here we assumed that $q_s = 0.8$ and $q_c = 0.2$. Thus, 80% of the evidence for a disease was assumed to come from conditions in C_1 and 20% from conditions in C_2. So, if any condition in the significant part is not satisfied, then the evidence degree of the whole linear module is greatly reduced. But if the same condition is in the confirmatory part, and is not satisfied, then the evidence degree of the whole linear module would be reduced only slightly. The values 0.8 and 0.2 were selected experimentally.

3. EVALUATING DECISION RULES

To diagnose a diseased plant, the expert-derived rules or the inductively-derived rules or both are evaluated using values of descriptors obtained by questioning the system's user. It is also possible to use a combination of two groups of rules. In the latter case, both groups of rules are used in series. First, the inductively-derived rules are applied to eliminate all but a few (e.g., five) most probable diseases, and then the expert-derived rules are used to determine the final diagnosis.

To simplify the explanation of how the system conducts rule evaluation let us first assume that the values of all descriptors for a diseased plant are already known. A diagnosis is determined by finding the rule whose condition part best matches the characteristics of the diseased plant, i.e., the rule with the highest evidence degree for the given values of the descriptors.

In general, there may be more than one rule with the maximum degree of evidence, or there may be rules whose evidence degrees differ only slightly. The system resolves this problem by giving alternative diagnoses when the evidence degree varies from the maximum value within an experimentally determined range δ (we used $\delta = 0.25$). In addition, PLANT/ds suggests advice only if the evidence degree of the diagnosis is above a certain threshold of acceptance. This threshold was determined experimentally to be 0.65 for expert-derived rules, and 0.8 for inductively-derived rules. Also, the best evaluation scheme for the expert-derived rules was found to be:

AVE: the average function for conjunction and
MAX: the maximum function for disjunction,

and for inductively-derived rules:

AVE: the average function for conjunction, and
PSUM: the probabilistic sum for disjunction.

The evidence degree computed for rules should not be taken as the statistical probability (frequency) of the correctness of the diagnosis. The evidence degree is simply an indicator of the degree to which a description of the diseased plant matches a decision rule using a given evaluation scheme.

To illustrate the rule evaluation process, the rules for Downy Mildew and Powdery Mildew will be evaluated for a particular diseased plant. The values of the descriptors for this plant are:

Time of Occurrence	- August
Precipitation	- Normal
Temperature	- Normal
Damaged Area	- Whole Fields
Condition of Leaves	- Abnormal
Leaf Spots	- Without Yellow Halos
Leaf Mildew Growth	- On Upper Leaf Surface
Premature Defoliation	- Present
Seed Mold Growth	- Absent
Leaf Malformation	- Absent
Condition of Seed	- Normal
Condition of Stem	- Normal

(the other variables are not relevant for these diseases and therefore are not listed). First we will evaluate the diagnostic rule for Powdery Mildew. The rule is:

$q_s \cdot$[Time of Occurrence = August..September]&
[Condition of Leaves = Abnormal]&
[Leaf Mildew Growth = On Upper Leaf Surface]
+
$q_c \cdot$[Precipitation < Normal][Temperature => Normal]

::> [Soybean Disease = Powdery Mildew]

The significant factors for Powdery Mildew are occurrence late in the season, in September or August, the leaves are abnormal and there is mildew on the top side of the leaves. The confirmatory evidence is less than normal preciptation and with temperatures that are at least normal if not above normal.

Since the time is August, the first selector, [Time of Occurrence = August..September], is satisfied, and so has the evidence degree 1. The other two selectors in the significant part (with coefficient q_s) are also satisfied, and so also have the evidence degree 1. The evidence degree of the whole significant part is the average value of these evidence degrees, i.e., 1. The first selector in the confirmatory part, [Precipitation < Normal], is not satisfied, because precipitation was normal. So it has the degree of evidence 0. The last selector, [Temperature $\geq$ Normal], is satisfied; so it has degree of evidence 1. The evidence degree for the confirmatory part is the average of 1 and 0, i.e., 0.5. The degree of evidence of the whole rule is (0.8 • 1.0) + (0.2 • 0.5) = 0.9. The same calculation is done for Downy Mildew. The rule is:

q_s•[Time of Occurrence: @ T7][Precipitation>Normal]&
[Damaged Area=Whole Fields]&
[Condition of Leaves=Abnormal]&
[Leaf Spots=Without Yellow Halos]&
[Condition of Stem=Normal]&
[Leaf Mildew Growth=On Lower Leaf Surface]&
([Time of Occurrence=September∨October] =>
[Seed Condition=Abnormal][Seed Mold Growth=Present])
+
q_c•[Premature Defoliation=Present]&
[Leaf Malformation=Present]

::> [Soybean Disease = Downy Mildew]

$$\text{where} \quad T_7 = \begin{cases} 1.0 - \text{July or August} \\ 0.8 - \text{June or September} \\ 0.7 - \text{October} \\ 0.1 - \text{Other} \end{cases}$$

The evidence degree of the first selector is 1 as determined by the weight function T7. The evidence degrees for the next five selectors are 1, 1, 1, 1 and 0 respectively. The next three selectors are a part of an implicative statement. All three selectors in it are false, but the whole statement has a degree of evidence of 1. This can be understood if the statement is translated to English. It says "if the disease occurred in either September or October, then there should be abnormalities in the seed and there should be mold growing on the seeds". Since it is neither September nor October, the statement evaluates to the evidence degree 1.

Returning to the evaluation of the rule, the significant part has the evidence degree: (1.0 + 1.0 + 1.0 + 1.0 + 1.0 + 0.0 + 1.0)/7 = 0.86 (the implicative statement is treated as one selector). Thus, the Downy Mildew rule is satisfied with the evidence degree: (0.8 • 0.86) + (0.2 • 0.5) = 0.79.

Now we can compare the two diagnostic evaluations. Powdery Mildew has degree of evidence 0.9, while Downy Mildew 0.79. Both values are above the acceptance threshold of 0.65, so neither is eliminated. Powdery Midew has a higher degree of evidence, so it is the first choice diagnosis (assuming that no other rule evaluates with a higher degree of evidence). Since Downy Mildew's evidence degree differs less than 0.25 from Powdery Mildew's, the latter is an alternative diagnosis.

Suppose now that leaf mildew growth was on the lower instead of the upper leaf surface. This variable is found in the significant part of both rules, so a change in the diagnosis can be expected. The degree of evidence becomes 0.64 for Powdery Mildew and 0.9 for Downy Mildew. The degree 0.64 is below the acceptance threshold, so Powdery Mildew is eliminated. This leaves Downy Mildew as the only diagnosis (unless there were some other rules with a higher degree of evidence).

Instead of using the averaging function for evaluating a complex, the minimum function could have been used. In this case, the degree of evidence for Powdery Mildew would be:

$$0.8 \cdot \text{MIN}\{1,1,1\} + 0.2 \cdot \text{MIN}\{0,1\} = 0.8,$$

and for Downy Mildew:

$$0.8 \cdot \text{MIN}\{1,1,1,1,1,1,0,1\} + 0.2 \cdot \text{MIN}\{1,0\} = 0.$$

Even if leaf malformation was present, the evidence degree for Downy Mildew would only be 0.2. Thus, using the minimum function leads to a much stronger separation between the evidence degrees for the two diseases (perhaps even too strong). The overall performance of expert-desired rules, however, determined on all known cases of the considered diseases, turned out to be better using the averaging function than the minimum function.

4. ITERATIVE EVALUATION OF DIAGNOSTIC RULES

In the interactive mode, PLANT/ds derives a diagnostic advice through asking a user a series of questions, determined on the basis of the rules in the knowledge base. As we mentioned earlier, the system contains two sets of rules: expert-derived and inductively-derived. To compute a diagnosis, only one set of rules is used, or a combination of both.

Let us assume at the beginning that only a single set is used. At the first step, several standard questions are asked, the answers to which enable the system to eliminate a number of other potential questions, and reduce the scope of rules under consideration (such rules are called the candidate rules). An example of a standard question is "Condition of Leaves." If the answer is "Normal," all questions concerning leaf abnormalities become irrelevant. Through a procedure, called approximate evaluation, some diagnostic rules can be eliminated without evaluating all the descriptors in them. The procedure computes the upper bound on the evidence degree of each partially evaluated rule by assuming that all the unevaluated selectors have the evidence degree 1. If so computed evidence degree of a rule falls below the threshold, the rule is eliminated.

Each subsequent step of the system generates a question to the user, and upon receiving an answer applies the procedure of approximate evaluation. Questions are determined by finding the most frequently occurring descriptor in the candidate rules. Obtaining a value of such a descriptor will tend to eliminate the largest number of rules. This descriptor is determined by a counter measuring the number of occurrences of each descriptor in the candidate rules. Whenever a rule is eliminated, the appropriate counters are updated. When the counter associated with a given descriptor drops to zero, the descriptor is eliminated from further consideration. When several descriptors occur with comparable frequency, the one is selected that is "semantically closest" to the descriptor asked about in the previous question(s). This is done by arranging descriptors into a hierarchy, and computing the "semantic distance" as the path length between the corresponding nodes. This facility makes the system act more like a human expert, who tends to ask a series of related questions, rather than jump back and forth from questions of one category to another.

The process ends when a single diagnosis, or a few alternative ones are determined. When the system uses both, inductively-derived and expert-derived rules, it proceeds in two stages. The first stage uses only inductively-derived rules, until the set of candidate rules reduces to a specified size (e.g., 5 rules). Then the system switches to the expert-derived rules. The justification for this is that the inductively-derived rules are simpler to evaluate, while expert-derived rules contain more details about the diseases.

5. AN EXPERIMENT

The rule evaluation process described in the previous two sections was implemented in two computer programs. One program (the interactive version) was designed to be used by a grower or anyone else with an unknown case of soybean disease. It has several facilities to help a user to understand how the program operates, and how an advice is being computed. The second program (the batch version) is for experimental purposes only, to evaluate the program performance on many cases of disease at once. The collected statistics about rule performance are used to locate the rules that need further refinement and to determine the best evaluation scheme.

Table 1 gives a summary of the results from using the batch version using 340 cases of soybean disease and applying separately the inductively-derived and expert-derived rules. The label "% 1st Choice Correct" tallys how often the diagnosis with the highest evidence degree was the correct one. The label "% Correct" tallys how often the correct diagnosis was among generated alternatives. Label "% Not diagnosed" indicates how often the system could not identify a case of disease. The "Indecision Ratio" measures the precision of diagnostic advices. It is defined as the average number of alternative diagnoses generated per case of disease. A low indecision ratio is desirable, but it does not imply correctness. "Threshold" is the minimum degree of evidence that a rule must achieve in order that a diagnostic advice is made (the threshold of acceptance).

Type of Rules	% 1st Choice Correct	% Correct	% Not Diagnosed	% Ratio	% Threshold
Expert-Derived	71.8	96.9	2.1	2.90	0.65
Inductively-Derived	97.6	100.0	-	2.64	0.80

Table 1. Summary of results from PLANT/ds

6. CONCLUSION

The paper presented an experimental expert system PLANT/ds for advising farmers and other users on the diagnosis of soybean diseases common in the state of Illinois. The design of the system was oriented toward a novice user, not trained in the computer technology. In one of the modes of operation, the user can interact with the system just by touching appropriate places on the screen of the terminal. The knowledge base of the system contains two types of diagnostic rules: expert-derived, which were obtained by formally representing the diagnostic knowledge of a plant pathologist, and inductively-derived, which were obtained by applying a general inductive learning system (AQ11) to several hundred cases of the diseases. The Illinois Cooperative Extension Service is planning to use the system in selected offices in the state.

ACKNOWLEDGMENTS

The initial version of PLANT/ds was developed by R. L. Chilausky in collaboration with R. S. Michalski. The authors thank Paul Baim, Albert Boulanger, and Carl Uhrik for comments and criticism. The authors gratefully acknowledge a partial support of this work by a grant from the United States Department of Agriculture, No. 3215112344 and in part by a grant from the National Science Foundation, No. 79-06614.

REFERENCES

[1] B. G. Buchanan and E. A. Feigenbaum, "Dendral and Meta-Dendral, their applications dimension," Artificial Intelligence J., 11:5-24 (1978).

[2] R. Davis, "Applications of Meta Level Knowledge to the Construction, Maintenance and Use of Large Knowledge Bases," STAN-CS-76-552, Department of Computer Science, Stanford University, Stanford, California (July 1976).

[3] R. O. Duda et. al., "A Computer-Based Consultant for Mineral Exploration," Final Report, SRI International, Menlo Park, California (September 1979).

[4] Gladwin, C. H., "A theory of real-life choice: applications to agricultural decisions," in: P. Barlett, ed., Agricultural Decision Making, Academic Press, Inc., (New York 1980).

[5] R. M. Michalski, "Variable-valued Logic: System VL_1," Proceedings of the Fourth International Symposium on Multiple Valued Logic, Morgantown, West Virginia (May 1974).

[6] R. S. Michalski and R. L. Chilausky, "Learning by Being Told and Learning from Examples: An Experimental Comparison of the Two Methods of Knowledge Acquisition in the Context of Developing an Expert System for Soybean Disease Diagnosis," International Journal of Policy Analysis and Information Systems, Vol. 4, No. 2, (1980).

[7] Michalski, R. S., Davis, J. H., Bisht, V. S., and J. B. Sinclair, "PLANT/ds An experimental computer consulting system for the diagnosis of soybean diseases. (Abstract) Phytopathology 71: (1981).

[8] D. Michie, "Knowledge-Based Systems," Report No. 1001, Department of Computer Science, University of Illinois, Urbana, Illinois (January 1980).

[9] D. Michie (editor), Expert Systems in the Micro Electronic Age, Edinburgh University Press (1979).

[10] J. D. Myers and H. E. Pople, "INTERNIST: A consultative diagnostic program in internal medicine," Proceedings of the 1st Annual Symposium on Computer Applications in Medical Care, IEEE, New York (1977).

[11] E. H. Shortliffe, Computer Based Medical Consultation. MYCIN, Elsevier North-Holland, New York, 1976.

[12] S. M. Weiss, C. A. Kulikowski, S. Amarel and A. Safir, "A model-based method for computer-aided medical decision-making," Artificial Intelligence. 11(12):145-172 (1978).

[13] Zadeh, L. A. "A Fuzzy-Algorithmic Approach to the Definition of Complex or Imprecise Concepts." Memorandum no. ERL-M474, Electronics Research Laboratory, College of Engineering, University of California, Berkeley, Oct. 74.