ISTITUTO AGRARIO
DI S. MICHELE A/A
1992
GER
1
BIBLIOTECA

POMI: An Expert System for Integrated Pest Management of Apple Orchards

Alfonso Gerevini, Anna Perini,
and Francesco Ricci
Istituto per la Ricerca Scientifica e Tecnologica
38050 Povo, Italy

Diego Forti, Claudio Ioriatti, Luisa Mattedi,
and Ambrogio Monetti
Istituto Agrario S. Michele
38010 S.Michele all'Adige, Italy

ABSTRACT. This paper describes an expert system for apple orchard integrated control. The system consists of two parts. The first provides help to the agricultural producer in classifying producer observations. In the second part, an interactive abductive module suggests explanations for those observations. Classification has been implemented with the aid of discrimination net techniques. The abductive module exploits results from the parsimonious covering theory of Peng and Reggia. The system has been implemented in KEE with extensive use of graphics. In future work, a module responsible for treatment suggestion will extend the system.

The traditional model of agricultural development, based on massive introduction of innovative technologies, has created great progress, but has also often led to a disruption of the agriculture-environment equilibrium. The general reevaluation of ecological-environmental problems has increased interest in forms of agriculture characterized by reduced environmental impact. Within this framework, particular emphasis has been given to integrated pest management and, more generally, to integrated protection. One objective is to maintain the population of phytophagous pests and damage at tolerable levels. The introduction of the concept of tolerance threshold in agronomic practice clearly demonstrates that a major reduction in phytosanitary treatments is possible. Moreover, chemical products are selected from those authorized by the governing legislation, as provided by the International Organization for Biological Control (IOBC). They are limited to the indispensable minimum, favoring the self-defense of plants and operating with respect for beneficial organisms. The impact of that methodology is based not just on producers' acceptance of these technical views, but on their direct involvement and their ability to put into effect specific recommendations in various operating sectors. Our project is aimed at favoring the full application of these integrated protection criteria by providing a knowledge-based system for the use of technicians.

Several research projects aimed at analogous goals are on the way. AI techniques are more and more often used to cope with the complexity of environmental and agricultural problems (Plant et al. 1991). CALEX/COTTON (Goodell et al.1990) is a cotton management software package developed by the University of California Integrated Pest Management group. CALEX integrates AI techniques with conventional methods. rbWHIMS (Olson et al.1990) is the result of another research project aimed at the efficient management of cotton arthropod pests. It comprises an object-oriented, rule-based system that incorporates the heuristic expertise of extension and research entomologists. Henrion and Cooley deal with uncertainty in a decision analysis approach to the diagnosis and treatment of root disorders in apple trees (Henrion 1989, Henrion and Cooley 1987), and POMME (Roach et al.1985), a rule-based system written in PROLOG, addresses apple orchard management.

THE TRADITIONAL MODEL OF AGRICULTURAL DEVELOPMENT, BASED ON MASSIVE INTRODUCTION OF INNOVATIVE TECHNOLOGIES, HAS CREATED GREAT PROGRESS, BUT HAS ALSO OFTEN LED TO A DISRUPTION OF THE AGRICULTURE-ENVIRONMENT EQUILIBRIUM.

The system we present here addresses the first phase of the complex process of apple orchard integrated control. In that phase, one has to detect all the insect populations in the field and approximate the populations' dimension. In general, whereas pests are not directly observable, one can easily detect the damages produced (Alford 1984). The problem is first to identify as precisely as possible all the types of damage, then to identify the pests that can account for this damage. This problem is a typical example of the diagnostic reasoning process.

The system we have developed consists of two parts. The first provides help to the agricultural producer in classifying producer observations. In the second part, an interactive abductive module suggests explanations for those observations. The next section (The Apple Pests Domain) illustrates some key problems that arise in describing producer observations. The subsequent two sections describe the classification task and the abductive module. We have exploited the Peng and Reggia (1990) scheme of abductive reasoning that is based on the parsimonious covering theory. The last section in this paper is concerned with implementation, which was performed with KEE (Knowledge Engineering Environment) (Fikes and Kehler 1985).

This is part of a working project whose ultimate goal is to provide a comprehensive support system for apple orchard management. In the future, the system will be extended to include estimates of the dimension and the stage of evolution of the colonies detected, including insects that do not cause damage, but that may be useful antagonists to the pest species. Weather conditions will also be incorporated, as well as some information about the location and the history of the orchard under consideration. The goal is to give precise advice about available treatments compatible with the principles of integrated protection.

The Apple Pests Domain

One of the typical tasks in apple orchard management is the control of plant health through regular inspections at the orchard, particularly during critical vegetative phases such as plant sprouting and bloom. The leaves, the blossoms, or the fruit are carefully examined for possible anomalies. Phytophagous activity, as well as stressful atmospheric conditions, are among the principal causes of apple orchard damage. Here we concentrate on the problem of diagnosing apple damage produced by pests.

Usually the pests causing damage are not directly observable or recognizable, but some deformation or abnormal appearance of some plant parts can be noticed and reported to an entomologist for diagnosis and determination of the most appropriate treatment. A crucial step in this process is deriving, from the farmer, the most specific damage description from observations collected at the orchard. This process is illustrated by the following example of a farmer-expert interaction.

During the month of July, the farmer noticed some anomalous leaf and fruit characteristics: two or more leaves webbed together, forming a shelter between them, as well as small cavities just below the skin of the young apples. Based on these observations, the expert can guess the presence of leaf rollers causing the webbed foliage; the fruit damage indicates that larval burrowing has occurred. Because at least three different pest species could be active on fruits in July (oriental fruit moth, codling moth, European corn borer), the expert needs more specific damage feature observations. He/she cuts some damaged apples; a deep gallery reaching the fruit seeds is visible. The expert also analyzes the most recently damaged fruits in order to characterize the entry hole. Red-ringed entry holes are noticed, leading to a diagnosis of codling moth as the only fruit damage agent.

As this example shows, obtaining a satisfactory diagnosis requires the effort to point out typical damages due to a specific pest or at least to identify damage features that allow exclusion of some pests from a set of possible agents. For example, the presence of deep galleries reaching the fruit seeds leads to the exclusion of a diagnosis of oriental fruit moth larvae, which typically feed in pulp-producing galleries just below the fruit skin. The ringed entry holes, a typical feature of damage produced by codling moth larvae, allowed the distinction between European corn borer and codling moth to be made.

The process of making the orchard observations description more and more specific in order to minimize the set of identifiable pests is what we call "classification." A few problems related to it are:

- The *terminological redundancy* in making a qualitative description of damage. For example, "tissue damage," "deep or superficial galleries," and "holes" are terms that could be used in the example described above. While the slight differences in the meanings of these terms might not seem relevant to a non-expert, they can become important discrimination elements for the entomologist who takes into account pests' specific behavior and seasonal population development.
- The *partial description* problem. The observations collected from the producer are often too generic and can lead to an uncertain diagnosis. Further observations at the site are necessary in order to refine it. In the example, the observation of a burrow into the fruit tissue during the month of July suggested three different pests as possible damage agents.
- The *time period* during which the observations are performed. Because there is a great deal of knowledge available about pest development with respect to plant seasonal phenology, timing is one of the most useful discriminating elements. Particular care has to be taken in identifying the distinct time periods in the framework of the task of diagnosis; for example, the phenological stages related to plant sprouting and bloom can take place during different months depending on the geographic location of the orchard.

Classifying Observations

This section briefly describes the knowledge acquisition phase, the main features of the knowledge base, and the classification procedure that has been implemented. The first phase of the project has been characterized by strong interaction with the domain's experts in order to clarify the description of the problem and to identify the basic concepts and relations. The knowledge acquisition phase included both interviews with the experts and consultation of practical booklets (Baggiolini et al. 1983, Mattedi et al. 1989). This literature describes the seasonal observations the farmer should make in the orchard and the damages produced by pests. This information was collected in a database that was revised after each interaction with the experts.

We focused our attention on a restricted set of pests and their associated damages, and we selected a list of significant and directly observable features characterizing the damage description. These features were then ranked according to their importance, that is, their "discrimination capacity." Within this context, a feature is more discriminating than another if the related damage's description denotes a smaller set of possible causes (pests). Among the most relevant features are:

- The time period in which the observations have been gathered. Different pests can cause the same damage, but in different periods of time, so knowing the time period during which the observations have been taken obviously restricts the set of possible causes. Usually damage observations are performed over 17 different time periods corresponding to 10 phenological phases (from the "closed buds" stage to the "fruitlet" stage) and over seven months (from June to December).
- The kind of damage, for example, anomalous coloring or cavities under the fruit skin.
- The damage location, for example, fruit, buds, shoots, flower stems, leaf stalks.

To cope with the problem of terminological redundancy, the system forces the user to encode a natural language description into its corresponding formal specification. The system is quite rigid in the sense that it constrains the user to follow a rigid sequence of discrimination steps, but, in this way, it also avoids problems of possible ambiguity in the classification of the damage.

The damage descriptions have been represented with the aid of a *discrimination tree* (Charniak et al. 1987) in which the most generic description is placed at the root node and the most specific descriptions at the leaf nodes. Discrimination trees are basic tools in diagnostic problem solving and have been widely used both in AI systems and in classical statistically based applications (Quinlan 1988). A discrimination tree consists of a tree with a question (test) at each branching point that determines which branch is to be followed. For example, at a node a question could be: "What is the location of the observed damage?" For each possible answer (such as leaves, fruits, etc.), there is a child node that organizes all the objects (damage descriptions) corresponding to the answer.

THE ULTIMATE GOAL OF THE PROJECT IS TO PROVIDE A COMPREHENSIVE SUPPORT SYSTEM FOR APPLE ORCHARD MANAGEMENT.

In our discrimination tree, the leaves correspond to specific damage descriptions and identify a particular set of pests that, without any further data, can all be responsible for the observed damage. For example, the leaf node corresponding to the damage "presence of pest frass on the external surface of the fruit" is caused by the clerk larva during the "fruitset" and "fruitlet" stages, and by the codling moth during July and September.

Figure 1 shows the path in the discrimination tree along which the user reaches the description of the damage "red-ringed entry holes." The user can consult the discrimination tree starting from the root, answering the first question and then proceeding to the other questions. He/she goes deeper into the tree until a node has been reached that is either a leaf node or is associated with a description that cannot be further refined (partial description). At each step, the user can also backtrack to the previous node, select that partial description, or choose another branch in the discrimination tree.

In order to handle particular cases in which a complete damage description cannot be provided by the user, and to take into account the "partial description" problem described in the previous paragraph, the system also permits the creation of an instance of "partial description" (not leaf nodes). In fact, each non-leaf node of the discrimination tree represents a partial or incomplete description of a specific damage. The system associates with that node the set of all pests that can possibly cause that damage. That set is computed as the union of the set of pests associated with the leaves of the subtree rooted at that node. The observations provided by the user, as previously described, are classified in specific or partially described damages and, together with the associated pests, become the input data of the diagnostic phase.

The Pest Diagnosis

The result of the classification module consists of a set of observations classified by the system. The classification phase associates with every observation a set of pests that can cause that observed damage or symptom. The abductive module, described in this section, takes these observations incrementally, and, using results from the parsimonious covering theory (Peng and Reggia 1990), builds lists of

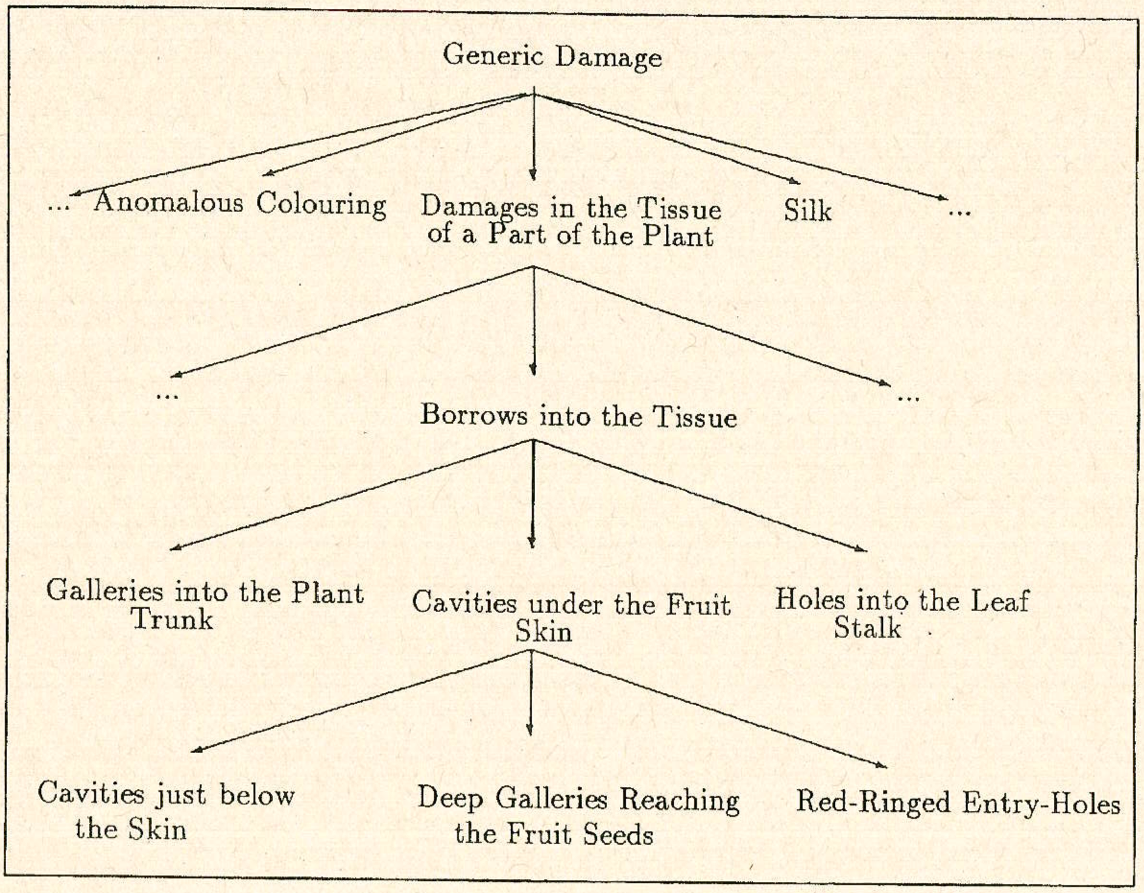

Figure 1.
Part of the discrimination tree.

insects that are associated with the observations.

Our approach to the diagnostic problem is an implementation of the well-known hypothesize-and-test process. This procedure consists of three steps:

- based on some initial observations, the program forms a first set of explanations (abduction);
- these explanations suggest tests and further observations; and
- the first set of explanations is revised to account for tests and new observations.

In our application, an explanation consists of a set of pests which can causally account for the observed damages and can be present in the orchard. The process of generating a plausible explanation for a given set of observations or facts is called *abduction*. Abduction reverses the order of the classical deductive reasoning process. Whereas deduction moves from premises to conclusions, abduction looks for causes that can explain some findings. In general, the abduction process is not simple, because an observation can have many different causes; this is the case in our application too. Damage, such as "small cavities just below the skin of young apples," can be produced by a number of different insects. A first hypothesis including every pest that can account for that damage has to be refined by consideration of some other observation.

Moreover, listing all the insects that can produce a given set of observations is not an adequate procedure; this naive approach would yield an extremely large set of insects. We should require stronger evidence to include an insect into an explanation. The two basic principles we have followed come from the parsimonious covering theory (Peng and Reggia 1990). They are:

- an insect is included in an explanation if it is a possible cause of some observation (covering);

and

- an insect is excluded from some explanation if there is some other insect that might cause the same damage and that can also take into account an additional type of damage (parsimonious).

To proceed we need some formal definitions (Peng and Reggia 1990). Let P and D be the set of all pests and all damages, respectively, that are known by the system. The knowledge base enables one to define the following functions:

- *effects(p,t)* is the subset of D containing all the damages produced by pest p at time t;
- *causes(d,t)* is the subset of P containing all the pests that can produce damage d at time t.

Here we refer to time t as one of the 10 phenomenological stages or one of the months from June to December. The functions *effects* and *causes* may be extended to sets in the following way:

$$effects(A,t) = \bigcup_{p \in A} effects(p,t), \text{for all } A \subseteq P,$$

and

$$causes(B,t) = \bigcup_{d \in B} causes(d,t), \text{for all } B \subseteq D.$$

We shall say that a set of pests $A \subseteq P$ is a *cover* of a set of damages $B \subseteq D$ if $B \subseteq$ *effects(A,t)* (at time t). In other words, a set of pests covers a set of damages if it can causally account for those damages. For example, the set whose elements are leaf rollers, oriental fruit moth, and codling moth is a cover of the set of damages "two or more leaves webbed together forming a shelter between them" and "small cavities just below the skin of young apples" during the month of July.

As we have previously observed, covering cannot be the only criterion defining an explanation. We should require the satisfaction of some additional constraint. In our application we have imposed the irredundant one. A cover $C \subseteq P$ is an *explanation* of a set of damages $B \subseteq D$ if it is *irredundant*, that is, if none of its proper subsets is also a cover of B. In the example previously mentioned, the set leaf rollers, oriental fruit moth, and codling moth is not an irredundant cover of the two damages "two or more leaves webbed together forming a shelter between them" and "small cavities just below the skin of young apples." In fact, either the oriental fruit moth or the codling moth may produce the small cavities.

Finding all the explanations of a given set of damages is not a simple goal. The parsimonious covering theory provides a tool for finding all the irredundant covers of a given set of damages. It is beyond the scope of this paper to present a full description of that theory; we refer the interested reader to the original book by Peng and Reggia (1990). In our application, the system provides the user with two functions:

- *solve* $(A) \subseteq \wp(P)$ is the set of all the explanations of a given set of damages $A \subseteq D$;
- *revise* $(E, causes(d)) \subseteq \wp(P)$ is the set of all explanations that explains the same set of damages explained by explanations in $E \subseteq \wp(P)$ plus $d \in D$.

$\wp(X)$ denotes the set of all subsets of X. The function *solve* consists of a recursive call to the function *revise* on a set of damages. Every abduction/hypothesis phase produces a call to the *revise* function on the additional damage detected. In many cases, a single call to *solve* on the set of detected damage is enough; that is, only one irredundant explanation is found. Whenever more than one explanation is found, the user, with the help of the system, has to check in turn whether the other damages predicted by the explanations can be observed. If this is the case (i.e., an additional damage is observable), the user may call *revise* on the previous hypothesis (set of explanations) and on the current additional damage. This process can be iterated many times until the set of explanations is refined to a single element, that represents the correct explanation. This process is illustrated in the example described above in the section on the apple pests domain. When the function *solve* is called on the damages "two or more leaves webbed together forming a shelter between them" and "small cavities just below the skin of young apples," three explanations are produced: 1) leaf rollers and oriental fruit moth; 2) leaf rollers and codling moth; and 3) leaf rollers and European corn borer. A next call with the two additional damages "deep gallery reaching the fruit seeds" and "red ringed entry holes" refine the set of explanations to a single element, namely leaf rollers and codling moth.

We have tested the knowledge base and the abduction procedure against many real-world cases, with the aid of an agricultural technician. We conclude that it is quite easy for a technician to observe a sufficient number of damages so that the abduction algorithm can produce a small number of explanations that can be easily refined to a single explanation. Moreover, the irredundant criterion is not too narrow, as it does not exclude any reasonable candidate explanation.

System Implementation

The POMI system has been implemented in KEE 3.1 (Knowledge Engineering Environment, IntelliCorp, Inc.) on a SUN4 Sparcstation. The KEE system (Fikes and Kehler 1985) belongs to the family of software tools for building knowledge-based systems which provide different knowledge representation and reasoning schema, such as frame-based representation and production rules, together with a powerful development interface based on menu and browsing facilities. The object-oriented graphics toolkit ("KEEpictures") allows end-user interfaces to be built employing standard pictures, bitmaps, and animation techniques.

The POMI knowledge base contains two object hierarchies, one corresponding to the damage discrimination tree and the other to the pest descriptions and the system interface.

- *Pests as objects.* Each pest is represented as an object characterized by its Latin name and the damages caused, i.e., a set of pointers to the objects belonging to the damage discrimination tree. At present, 35 pests have been described.
- *The damage discrimination tree.* This consists of a tree having a generic damage description at the root and the most specific damage description at the leaves. It has been implemented as a damage object hierarchy (Fig. 2). Each damage is an object defined by an extended description of the damage itself and for each time period by the pointers to the pests which can cause the damage in that period. The damage hierarchy consists of 200 objects.
- *The POMI interface.* In the current phase of the project, the user of the POMI system is the agricultural technician who can be asked by the farmer to make a diagnosis explaining a set of observed damages. The system interface design rests on the decision to make both the knowledge-base information and the information resulting from the interaction accessible at each point during the user-system interaction. To get the information, the user clicks the mouse on the appropriate buttons of the interface (Fig. 3). Navigation along the damage discrimination tree is achieved by making the right choice from a menu that may be activated at each node of the tree.

Finally, the diagnostic module has been implemented as a collection of LISP functions.

Conclusions and Future Work

This paper describes POMI, a pest management expert system for apple orchards. The system focuses on two main problems, classification of damage observations and their diagnosis. The main problems faced in classification are related to the identification of a precise language terminology in the description of damages. The same damage can be defined using different descriptions and terminology, and this language is far from being easily formalizable. This problem has been the main obstacle in an early implementation of the system, written in LOOM (MacGregor 1988), a hybrid knowledge representation language with an automatic classifier based on the concept of "term subsumption." The current implementation of the classification module is based on a discrimination tree. In future work, different descriptions of the same damage manifestation will be available, allowing the user to navigate along multiple paths connecting the root node with the same leaf node.

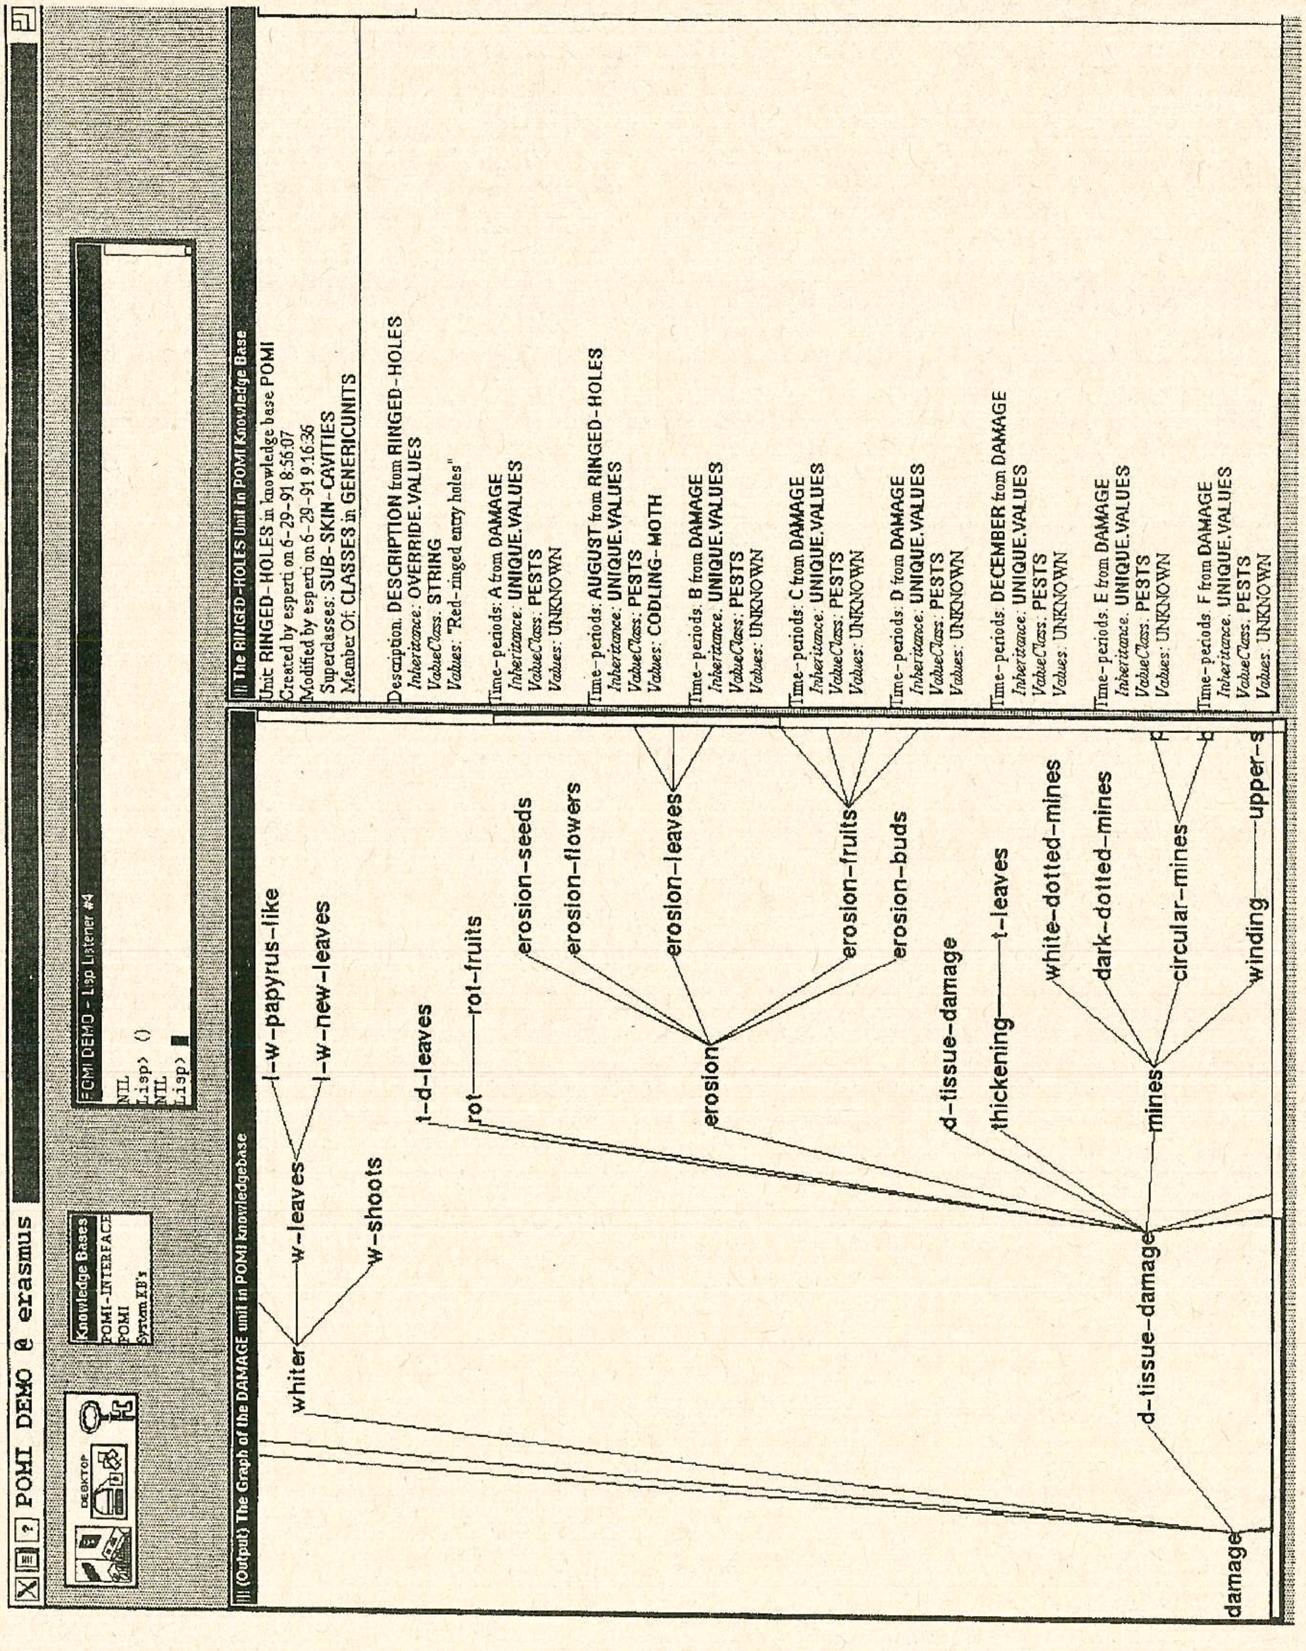

Figure 2. *Partial view of the damage hierarchy and display of a damage object.*

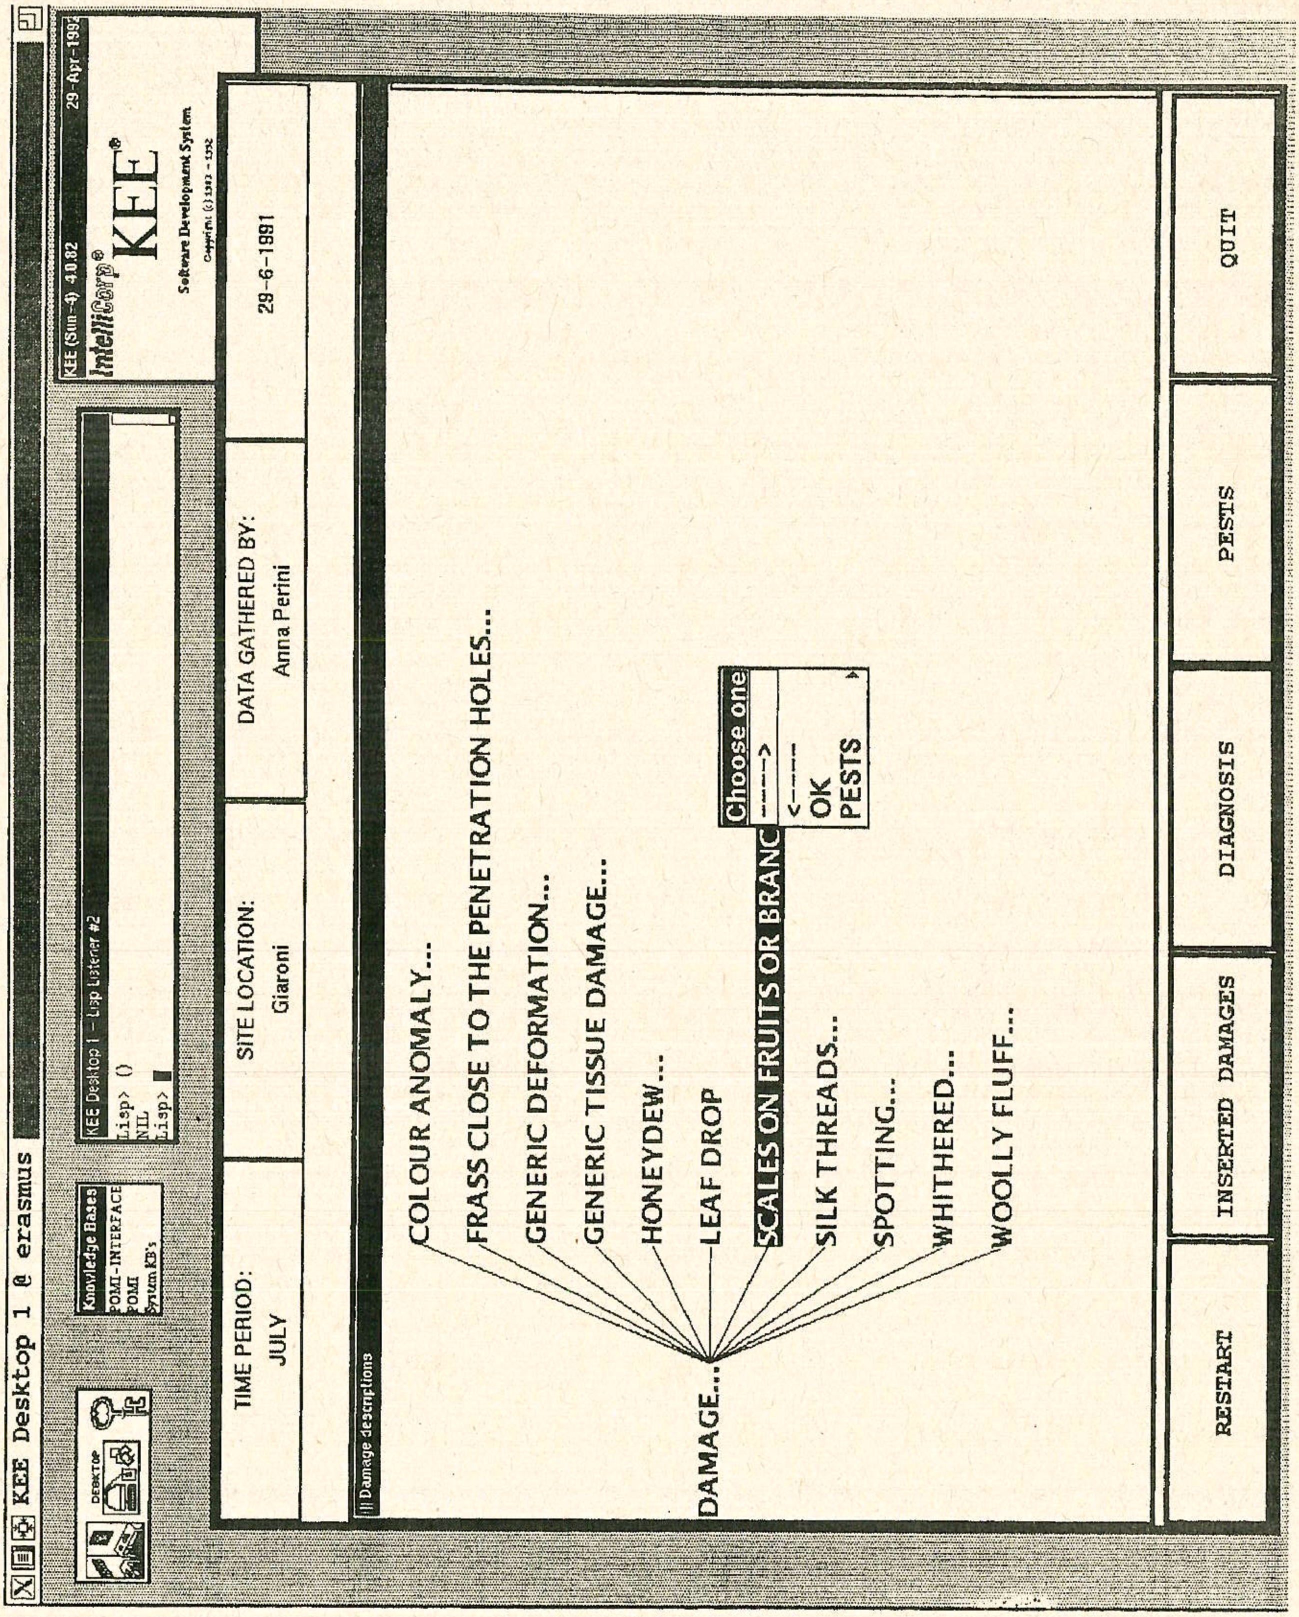

Figure 3. *The interface main panel.*

The kind of classification provided by the system is "semi-automatic" in the sense that it constrains the user to follow a rigid sequence of discrimination steps and does not classify automatically the user's damage description. Case-based reasoning (Slade 1991) is a methodology that we are investigating to study the potential for an automatic classifier sensitive to the problems of the terminology mentioned above and able to interpret the user's descriptions on the basis of a memory containing a set of previous cases.

Future work on the diagnostic step will concentrate on an additional testing phase, with examples from real cases provided by the experts. This work will help us examine the correctness of the results produced by the system and the adequacy of the techniques used in the implementation of the diagnostic module. Finally, a further development of the project will comprise the design and implementation of a module for helping the user select a treatment compatible with the basic tenets of integrated control.

References

Alford, D. V. 1984. Color atlas of fruit pests. Wolfe Publishing, Glasgow, UK.

Baggiolini, M., E. Keller, H. G. Milaire, and H. Steiner. 1983. Guida per i rilievi periodici dei nemici del melo. Organizzazione Internazionale di Lotta Biologica contro i nemici delle piante, Verona, Italy.

Charniak, E., C. K. Riesbeck, D. V. McDermott, and J. R. Meehan. 1987. Artificial intelligence programming. Lawrence Erlbaum Associates Publishers, Hillsdale, New Jersey.

Fikes, R. E., and T. P. Kehler. 1985. The role of frame-based representation in reasoning. Communications of the ACM 28(9): 904-920.

Goodell, P. B., R. E. Plant, T. A. Kirby, J. F. Strand, L. T. Wilson, L. Zelinski, J. A. Young, A. Corbett, R. D. Horrocks, and R. Vargas. 1990. CALEX COTTON: an integrated expert system for crop production and management in California cotton. California Agriculture 44(5): 18-20.

MacGregor, R. M. 1988. A deductive pattern matcher. Pages 403-408 in: Proceedings, Seventh National Conference on Artificial Intelligence, Saint Paul, Minnesota, August 1988. Morgan Kaufmann, San Mateo, California.

Henrion, M., and D. R. Cooley. 1987. An experimental comparison of knowledge engineering for expert systems and for decision analysis. Pages 471-476 in: Proceedings, Sixth National Conference on Artificial Intelligence, Seattle, Washington, July 1987, Morgan Kaufmann, San Mateo, California.

Henrion, M. 1989. Some practical issues in constructing belief networks. Pages 161-173 in: Uncertainty in Artificial Intelligence 3, L.N. Kanal, T.S. Levitt, and J.F. Lemmer, editors. North-Holland, Amsterdam, Netherlands.

Mattedi, L., M. Varner, and D. Forti. 1989. Natura e Agricoltura. Centro studi rotaliani Mezzocorona, Mezzocorona, Italy.

Olson, R. L., J. L. Willers, and T. L. Wagner. 1990. A framework for modeling uncertain reasoning in ecosystem management. II. Bayesian belief networks. AI Applications in Natural Resource Management 4(4): 11-24.

Peng, Y., and J. A. Reggia. 1990. Abductive inference models for diagnostic problem solving. Springer-Verlag, New York.

Plant, R. E., and N. D. Stone. 1991. Knowledge-based systems in agriculture. McGraw-Hill, New York.

Quinlan, J. R. 1988. Decision trees and multi-valued attributes. Pages 305-318 in: Machine Intelligence II, J. E. Hayes and D. Michie, editors. Oxford University Press, Oxford, UK.

Roach, J. W., R.S. Virkar, M. J. Weaver, and C. R. Drake. 1985. POMME: a computer-based consultation system for apple orchard management using prolog. Expert Systems 2(2): 56-68.

Slade, S. 1991. Case based reasoning: A research paradigm. AI Magazine 12(1): 42-55.

Alfonso Gerevini has been a research scientist with the Artificial Intelligence Department at the Istituto per la Ricerca Scientifica e Tecnologica, in Trento (Italy) since 1988. He received his doctoral degree with honors in Computer Science from the University of Milano in 1987. Since 1988 he has been engaged in the development of knowledge-based systems in the fields of office automation and decision support in medicine and agricultural management. His research interests include knowledge representation, uncertainty management, temporal rea-

soning and constraint satisfaction problems. He is a member of AI*IA (Associazione Italiana per l'Intelligenza Artificiale).

Anna Perini received her doctoral degree in physics from the University of Trento in 1981. She worked in the area of statistical mechanics and computer simulation applied to solid state physics problems, with a fellowship from the Istituto per la Ricerca Scientifica e Tecnologica, Trento, Italy. Since 1985 she has been a Research Scientist in the Artificial Intelligence Department at IRST, working in AI applications in the domain of office automation and of agricultural management. Her research interests include constraint satisfaction problems and temporal reasoning. She is a member of AAAI and AI*IA.

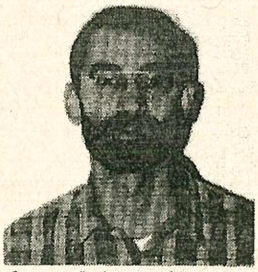

Francesco Ricci received his doctoral degree in mathematics from University of Padova, Italy, in 1983. He later attended Scuola Normale Superiore, Pisa, Italy, and from 1986 to 1987, he was with Enichem, S.p.A., in Milan, where he developed expert systems and office automation applications. Since 1988 he has been Research Scientist in the Artificial Intelligence Department at Istituto per la Ricerca Scientifica e Tecnologica, Trento. His current research interests include constraint satisfaction problems, expert systems, object-oriented programming, game theory, and automated learning. Ricci is a member of the American Association for Artificial Intelligence, AI*IA (Italian Association for Artificial Intelligence), and received the best degree award from the University of Padova in 1984.

Diego Forti graduated in agricultural science from the University of Piacenza in 1969. He has worked as an agronomist and taught plant pathology at the Agricultural Institute of San Michele a/A, Trento. Since 1986 he has been the manager of the Plant Protection Department at the Agricultural Institute of San Michele a/A.

Claudio Iorlatti was born in Trento, Italy, in 1958. He obtained his doctoral degree in agricultural science from the University of Padova in 1983. From 1984 to 1985 he worked at the ISF (Istituto Sperimentale di Frutticoltura) on genetic breeding of cherry. From 1985 to 1986, with a study grant of CNR (Centro Nazionale delle Ricerche), he was at RAC de Changins (CH), working on a program of research for developing new strategies in IPM in fruit and vine-growing. Since 1986 he has been a researcher at the Istituto Agrario S.Michele, Trento, where he works in the Department of Plant Protection. His current research interests include biological and biotechnical means of pest control, assessment of side effects of pesticides, and development of a complete package for integrated fruit production.

Luisa Mattedi graduated from the Istituto Agrario S. Michele, Trento, in 1976. She then took some specialization courses in technique and principles of integrated plant protection at the Station Federale de Recherches Agronomiques de Changins in Switzerland and at the Landesanstalt fuer Planzenschutz in Stuttgart, Germany. From 1979 to 1990 she worked as agricultural expert in plant protection (vineyard and apple trees) for the extension service in Trentino. Since 1990 she has been working at the laboratory of plant protection of the Istituto Agrario S. Michele, Trento.

Ambrogio Monetti graduated in agricultural sciences from the Catholic University of Milan in 1986. He is a member of the Biometric Society and works as a biometrician at the

Istituto Agrario S. Michele, Trento. Since 1990 he has had primary responsibility for the computing system of the Institute. In 1990 and 1991 he was a visiting scientist in the Statistical Department of Rothamsted Experimental Station (UK). He has interests in chemometrics and multivariate statistical analysis, mainly classification and identification, and expert systems in agriculture.