



US Army Corps
of Engineers

AD-A226 539

AQUATIC PLANT CONTROL RESEARCH PROGRAM

TECHNICAL REPORT A-90-8

FEASIBILITY OF USING EXPERT SYSTEMS IN AQUATIC PLANT CONTROL

by

L. R. Lawrence

Environmental Laboratory

DEPARTMENT OF THE ARMY

Waterways Experiment Station, Corps of Engineers
3909 Halls Ferry Road, Vicksburg, Mississippi 39180-6199

and

H. Lemmon

US Department of Agriculture
Agriculture Research Service
800 Buchanan Street
Albany, California 94170



**BEST
AVAILABLE COPY**

August 1990
Final Report

**DTIC
ELECTE
SEP 13 1990
E D**

Approved For Public Research; Distribution Unlimited



Prepared for DEPARTMENT OF THE ARMY
US Army Corps of Engineers
Washington, DC 20314-1000

90 09 12 036

Destroy this report when no longer needed. Do not return
it to the originator.

The findings in this report are not to be construed as an official
Department of the Army position unless so designated
by other authorized documents.

The contents of this report are not to be used for
advertising, publication, or promotional purposes.
Citation of trade names does not constitute an
official endorsement or approval of the use of
such commercial products.

**BEST
AVAILABLE COPY**

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			Approved for public release; distribution unlimited.		
4. PERFORMING ORGANIZATION REPORT NUMBER(S) Technical Report A-90-8			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION See reverse.		6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION		
6c. ADDRESS (City, State, and ZIP Code) See reverse.			7b. ADDRESS (City, State, and ZIP Code)		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION US Army Corps of Engineers		8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER		
8c. ADDRESS (City, State, and ZIP Code) Washington, DC 20314-1000			10. SOURCE OF FUNDING NUMBERS		
	PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT ACCESSION NO.	
11. TITLE (Include Security Classification) Feasibility of Using Expert Systems in Aquatic Plant Control					
12. PERSONAL AUTHOR(S) Lawrence, L. R.; Lemmon, H.					
13a. TYPE OF REPORT Final report		13b. TIME COVERED FROM _____ TO _____		14. DATE OF REPORT (Year, Month, Day) August 1990	
15. PAGE COUNT 52					
16. SUPPLEMENTARY NOTATION Available from National Technical Information Service, 5285 Port Royal Road, Springfield, VA 22161					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP			
			Aquatic plants Expert systems		
			Artificial intelligence Management		
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
Research into control of aquatic plants has evolved into different technology areas, including biological, chemical, and mechanical. The knowledge concerning aquatic plants and control strategies is shared by researchers, Corps District and project personnel, and state and local control personnel. As experience with different control methods has increased and the number of personnel involved has also increased, acquiring the most current and relevant information for a specific plant control application has become more difficult. Computer-based expert systems have been used successfully in other technical areas to help manage voluminous information and to help identify solutions to specific problems. Because of those successes and the growing base of knowledge that is not readily available to aquatic plant program managers, this report evaluates the feasibility of using expert systems in the management and control of aquatic plants. (Continued)					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL			22b. TELEPHONE (Include Area Code)		22c. OFFICE SYMBOL

6a. NAME OF PERFORMING ORGANIZATION (Continued).

USAEWES, Environmental Laboratory
and US Department of Agriculture

6c. ADDRESS (City, State, and ZIP Code) (Continued).

3909 Halls Ferry Road
Vicksburg, MS 39180-6199; and
800 Buchanan Street
Albany, GA 94170

19. ABSTRACT (Continued).

An Expert Systems Workshop was held at the US Army Engineer Waterways Experiment Station on 15 February 1989. Workshop participants included field users of aquatic plant control technology, researchers in the discipline, and expert system developers. The workshop discussions identified a number of areas in which an expert system could assist in aquatic plant control. These include: control applications, regulatory considerations, use of new control methods, orientation of new personnel, and dissemination of research findings. It is the consensus of the workshop group that building an expert system for aquatic plant management and control is both desirable and feasible.

A further recommendation by the group is that it would not be wise to attempt to address the overall expert system initially. A more prudent approach would be to build a small prototype expert system that would give users the opportunity to evaluate the capabilities of an expert system and to demonstrate its effectiveness. Participants determined that an excellent prototype would be the knowledge contained in the manual "Aquatic Plant Identification and Herbicide Use Guide, Volume II: Aquatic Plants and Susceptibility to Herbicides" by Westerdahl and Getsinger. Developing the guide into an expert system, even though a prototype, would be immediately useful to the aquatic plant managers and would provide an appropriate technology transfer application in a user-friendly format of state-of-the-art information.

PREFACE

The work reported herein was conducted as part of the Aquatic Plant Control Research Program (APCRP). The APCRP is sponsored by Headquarters, US Army Corps of Engineers (HQUSACE), and is assigned to the US Army Engineer Waterways Experiment Station (WES) under the purview of the Environmental Laboratory (EL). During the period of this study, Mr. J. Lewis Decell, WES, was the APCRP Program Manager, and Mr. E. Carl Brown, HQUSACE, was the APCRP Technical Monitor.

This report evaluates the feasibility of using computer expert systems in implementation and management of aquatic plant control programs by Federal, state, and local agencies. Mr. Larry Lawrence, EL, WES, was Principal Investigator for the evaluation. The evaluation was performed with input by personnel from the US Army Engineer Districts and projects, WES, and developers of expert systems. Dr. Hal Lemmon, US Department of Agriculture (USDA), provided expertise concerning expert systems. Permission to reprint Figure 1 in the text of this report was granted by Addison-Wesley Publishing Co. Material in Appendix A was reprinted from three publications, Science, Computers and Electronics in Agriculture, and AI Applications in Natural Resource Management. The articles were written by US Government employees and permission to reprint was granted by the American Association for the Advancement of Science, Elsevier Science Publishers B. V., and AI Applications in Natural Resource Management, respectively.

The study was supervised at WES by Mr. H. Roger Hamilton, Chief, Resource Analysis Group, EL, and Dr. Conrad J. Kirby, Chief, Environmental Resources Division, EL. Dr. John Harrison was Chief, EL. The report was edited for publication by Ms. Gilda Miller, Information Technology Laboratory, WES.

COL Larry B. Fulton, EN, is Commander and Director of WES.
Dr. Robert W. Whalin is Technical Director.

This report should be cited as follows:

Lawrence, L. R., and Lemmon, H. 1990. "Feasibility of Using Expert Systems in Aquatic Plant Control," Technical Report A-90-8, US Army Engineer Waterways Experiment Station, Vicksburg, MS.

CONTENTS

	<u>Page</u>
PREFACE.....	1
PART I: INTRODUCTION.....	3
PART II: COMPONENTS OF AN EXPERT SYSTEM.....	4
Basic Units of Expert Systems.....	4
Building an Expert System.....	7
Advantages of an Expert System.....	8
Disadvantages of an Expert System.....	8
Current Applications.....	9
PART III: WORKSHOP RESULTS.....	11
Control Applications Technology.....	11
Regulatory Considerations.....	12
New Products.....	12
Orientation of New Personnel.....	13
Research Findings.....	13
Recommendations.....	14
PART IV: SUMMARY.....	15
REFERENCES.....	16
APPENDIX A: INFORMATION PACKAGE.....	A1

FEASIBILITY OF USING EXPERT SYSTEMS
IN AQUATIC PLANT CONTROL

PART I: INTRODUCTION

1. Aquatic plants are important elements in aquatic ecosystems. They can, for example, provide habitat for various kinds of beneficial aquatic organisms. Aquatic plants can also interfere with navigation, recreation, water supply, irrigation, and other uses for water. Consequently, natural resources managers seek the most environmentally acceptable and cost effective means to maintain beneficial levels of aquatic plants.

2. Research for control of aquatic plants is conducted under different technology areas, including biological, chemical, mechanical, and integrated. Field application of the technologies has produced a better understanding of the technologies and management strategies. The knowledge concerning aquatic plants and control strategies is shared by researchers, Corps District and project personnel, and state and local control personnel. As experience with different control methods and the number of personnel involved have increased, acquiring the most current and relevant information for a specific plant control application has become more difficult. Computer-based expert systems have been used successfully in other technical areas to help manage voluminous information and to help identify solutions to specific problems. Because of those successes and the growing base of knowledge that is not readily available to aquatic plant program managers, this report evaluates the feasibility of using expert systems in the management and control of aquatic plants.

PART II: COMPONENTS OF AN EXPERT SYSTEM

3. Expertise consists of knowledge about a particular domain (technical area), understanding of the domain problems, and skill at solving some of these problems. An expert is a person with considerable knowledge of a particular field. That person's knowledge is acquired through formal and informal learning as well as experience (Frenzel 1987). Knowledge in any specialty is usually of two kinds: public and private. Public knowledge includes the published definitions, facts, and theories in textbooks and references. The private knowledge is based largely in heuristics, i.e., a method of education in which the individual relies on personal experiments, observations, practical experience, and rules of thumb to find solutions. Heuristics enable the human expert to make educated guesses when necessary, to recognize promising approaches to problems, and to deal effectively with incomplete data or data with errors (Hayes Roth, Waterman, and Lenat 1983). A relatively new technology for transferring public and private expert knowledge is the computer-based expert system. An expert system is "a computer program which has a wide base of knowledge in a restricted domain, and uses complex inferential reasoning to perform tasks which a human expert could do" (Hart 1986).

Basic Units of Expert Systems

4. The expert system itself consists of three basic units: the knowledge base (facts and rules), the inference engine (control of the use of rules), and the user interface (user/expert system interaction). The knowledge base contains facts and expertise about the domain. The inference engine decides the order of rule execution and makes inferences based on the knowledge base and input from the user. The user interface is the mechanism whereby the user can interact with the expert system (McGraw and Harbison-Briggs 1989).

5. Many of the rules in expert systems are heuristic simplifications that effectively narrow the search for solutions. Many problems of today are not adaptable to the mathematical analysis of algorithmic solutions. An algorithmic method guarantees to produce the correct or optimal solution to a problem, while a heuristic method produces an acceptable solution most of the time.

Knowledge base

6. Knowledge representation is the term used to describe how knowledge is structured in the expert system. The most popular form of knowledge representation is rule-based systems. Rule-based systems make use of IF condition and THEN action statements. An existing IF condition is run through the system to satisfy or match the IF part of a rule. When found, the action specified by the THEN part of the rule is performed and the rule is said to be true or to have fired. When a rule fires, the action inferred is stored in the data base so that it may be used by the inference engine to seek matches in other rules. If it does not fire, it may request additional input from the user. The inference engine searches through the knowledge base until there are no more rules and facts, and then it presents a conclusion. This matching of IF portions of statements can produce what are known as inference chains. An inference chain is the sequence of steps or rule applications used by a rule-based system to reach a conclusion (Waterman 1986).

7. Rules provide a practicable way for describing situations in today's rapidly changing and complex society. In a conventional computer program, the control and use of data are predetermined by the program code. Processing is done in sequential steps and branching (pursuing a new direction for a solution) occurs only at selected points. That type of processing works well for algorithmic solutions and slowly changing data such as solving a set of simultaneous linear equations. Rules, however, work well for data-driven problems with large numbers of branches (different solution directions). Rules enable the program to examine the problem at each step and react appropriately. The rule-based system is capable of explaining what the program did and how the conclusion was reached.

Inference engine

8. The control strategy in the inference engine determines how the rules in the knowledge base will be examined. This is done by a forward-chaining or a backward-chaining sequence. Chaining is the attempt by the inference engine to match facts obtained from the user with IF or THEN statements in the rules. In each case, the inference engine examines each rule in the particular sequence in an attempt to infer new information and thereby identify a solution for the given problem.

9. In many cases, the goal or solution must be assembled or constructed because there may be a large number of possible outcomes. These problems are more suited for forward-chaining. Forward-chaining systems make clear the

distinction between the knowledge base (the information provided by the expert) and the working memory (memory containing facts that emerge as a result of interaction with the user). The premises of the rules in the knowledge base are compared to the contents of working memory, and if they are true, given the information on hand, the conclusions are added to the list of facts and the system examines the rules again. Forward-chaining systems are therefore often referred to as data-driven systems (Harmon and King 1985).

10. Reasoning in a forward-chaining system is a "recognize-act" cycle. First, the rules that can fire, given the contents of the working memory, are recognized and identified sequentially. One rule is selected, and then the action or conclusion is asserted into working memory. The system then proceeds to the next cycle and checks again to determine what rules fire. The identification of aquatic plants can be determined by the following questions:

- a. Is the plant habitat above or below water (to eliminate one of the habitats)?
- b. Does the plant have single or multiple leafs?
- c. Are the leaf veins parallel?

The expert system can include color photographs as examples for the user to view when answering the questions. Even if the user fails to provide part of the description, the expert system can provide a conclusion and state a degree of certainty about that conclusion.

11. Backward chaining is applicable when a problem is caused by a limited number of possible conditions. If the possible outcomes (i.e., the answers or solutions) are known and if they are reasonably small in number, backward chaining is very efficient. Backward-chaining systems are sometimes called goal-directed systems. For example, assume that the user's goal is to determine the best method for controlling an aquatic plant. The system would begin by asking itself: is 2,4-D recommended? It then goes through the rules until it locates one stating that 2,4-D is recommended if spraying is recommended. Another rule states that spraying is recommended if the plant is emergent. The next rule states the plant is emergent if it is waterhyacinth. If the system cannot find a rule stating the plant is waterhyacinth, it simply asks if the plant is waterhyacinth. If the user responds "yes," then the expert system makes the appropriate inferences:

- a. The plant is waterhyacinth, and it is emergent.
- b. The plant is emergent, and spraying is recommended.
- c. Spraying is recommended, spray with 2,4-D.

12. When reasoning with either chaining system, the control strategy guides the solution by determining the order of the rules to be examined and which rule to examine next after a rule has fired.

Building an Expert System

13. The process of building an expert system is often referred to as knowledge engineering. It involves an interaction between the expert-system builder, referred to as the knowledge engineer, and one or more human experts in some problem area. The knowledge engineer extracts from the human experts their procedures, strategies, and rules of thumb for problem solving and builds this knowledge into the expert system as shown in Figure 1 (Waterman 1986).




Figure 1. Knowledge engineering: transfer of knowledge from human expert to a computer program (Waterman 1986, permission to reprint granted by Addison-Wesley Publishing Co.)

14. Researchers at the US Army Engineer Waterways Experiment Station (WES) developed an expert system to assess the feasibility of providing aquatic plant control program managers with recommendations for controlling aquatic plants. The expert system works as follows:

- a. The user is asked at what level he would like to think about his problem: the habitat level, the plant type level, or the species level.
- b. If the habitat level is selected the user is asked to select: above water or below water.
- c. If the plant type is selected the user is asked if the plant type is: emergent, floating, or submerged.
- d. If the species is selected the user is asked if the plant species is: alligatorweed, hydrilla, waterhyacinth, waterlettuce, or watermilfoil.

- e. Finally the user is asked if he wishes to consider that the possible treatments be applied concurrently or sequentially.

15. There are aquatic means of controls that can be applied to the stated problem. Controls are divided into four categories; biological, chemical, mechanical, and physical. Examples of each category are:

Biological: agasicles, n. bruchi, n. eichhornia, sameodes, thrips, white amur.

Chemical: 2,4-D, complexed copper, dichlobenil, diquat, endothall, fluridone, glyphosate, and triclopyr.

Mechanical: dredge and harvester.

Physical: barrier and drawdown.

The expert system eliminates all controls that are not effective for the problem stated. It does this by applying a set of rules formulated by the human experts in the field. The system then lists all possible combinations of controls that can be expected to be effectively applied.

Advantages of an Expert System

16. There are numerous benefits to developing expert systems. They allow you to preserve the expert's valuable knowledge. Should the expert leave, his or her knowledge can be used if it has been acquired and appropriately packaged into an expert system. Expert systems also help to understand how an expert solves a problem or uses knowledge. When creating an expert system, the knowledge engineer determines what knowledge is required and how it is used. Human experts are scarce and, therefore, costly. Expert systems are relatively inexpensive; they are costly to develop (i.e., the knowledge engineering) but relatively inexpensive to operate.

Disadvantages of an Expert System

17. Expert systems in most cases perform well, but there are areas where human expertise is superior to the programmed computer. This is not necessarily a fundamental limitation but more a current state of the art.

Limitations of the expert system are:

- a. Less creative and innovative than human expert.

- b. More programming and additions to the knowledge base are required for new concepts.
- c. Lacks the commonsense knowledge of human experts.

Current Applications

18. Currently, there are approximately 3,000 expert systems operating in the United States. Expert systems have been developed to address many different situations but can be grouped in the following categories and examples:

- a. Interpretation--inferring from input data along with a knowledge base in an attempt to understand the data and provide an explanation.
- b. Predictions--inferring likely consequences of given situations.
- c. Diagnosis--inferring system malfunctions from observations.
- d. Planning--designing actions.
- e. Monitoring--monitor a process and then provide an output control response.
- f. Instructional--evaluates a student's level of knowledge and understanding and can adjust the instructional process to the student's needs.
- g. Control--governing overall system behavior.

19. For example, there are expert systems that diagnose system malfunctions in an automobile electrical system, a high-performance disc drive, and a drill pipe stuck on a drilling rig. A system designed and used to discover a molybdenum deposit will probably exceed \$100,000,000 in value. Others diagnose bacterial infections in hospital patients, configure VAX computer systems (humans tend to forget to order components of the system), control the treatment of postsurgical patients in intensive care units, and monitor instrument readings in a nuclear reactor (looking for indications of an accident).

20. Expert systems have been used for a number of agricultural and natural resource management applications. COMAX is an expert system for cotton crop management used in making decisions about three factors related to cotton management: irrigation schedules, nitrogen requirements, and the crop maturity date (Lemmon 1986). Each day it computes the expected irrigation date, the expected date and amount of fertilization, and the expected date of crop maturity. These are computed daily because, as the predicted weather for each day is replaced by the actual weather for that day, the computed dates are recalculated. Growers believe that the system's ability to pinpoint the

day the crop is mature is the most valuable feature. "Expert Systems for Agriculture" (McKinnon and Lemmon 1985) is a related article reprinted in Appendix A. An expert system for decision support in resource management has been developed for rangeland grasshopper treatment selection (Kemp, Onsager, and Lemmon 1988). The system addresses concerns about more cost-effective treatments, better timing of applications, environmental sensitivity, and a lack of local expertise.

21. The aquatic plant management programs of the Corps, other Federal, state, and local agencies involve many human experts in different technical areas/domains. This results in a large amount of available technology with difficulty in transferring it to potential users. This evaluation was undertaken to determine if it is feasible and desirable to use expert systems for formulating and evaluating solutions to problems in aquatic plant control. At the 23rd annual meeting of the Aquatic Plant Control Research Program it was decided that an expert system workshop would be held. The Expert Systems Workshop was held at WES on 15 February 1989. The attendees were:

Al Cofrancesco, WES
Robert Gunkel, WES
Joyce Johnson, US Army Engineer District, Galveston
Larry Lawrence, WES
Hal Lemmon, US Department of Agriculture
Ron Mediema, Lake Okachoobee, Florida
Robert Rawson, US Army Engineer District, Seattle
Craig Smith, WES

PART III: WORKSHOP RESULTS

22. The objective of the workshop was to determine whether or not expert systems can be used to assimilate the knowledge and reasoning required to control aquatic plants. Workshop participants included field users of aquatic plant control technology, researchers in the discipline, and expert system developers. The evaluation was accomplished by discussion among field personnel explaining aquatic plant control with Dr. Hal Lemmon, developer of several expert systems for the US Department of Agriculture. An information package (Appendix A) was sent to each of the workshop participants for their review prior to the meeting. The package stated the objectives of the meeting, explained what expert systems are and how they function, and provided articles explaining specific applications of the expert systems. These articles were written by US Government employees for publication in Science, Computers and Electronics in Agriculture, and AI Applications in Natural Resource Management. Permission to reprint was granted by the American Association for the Advancement of Science, Elsevier Publishers B.V., and AI Applications in Natural Resource Management, respectively.

23. The workshop discussions identified a number of areas where an expert system could assist in aquatic plant control. These include: control applications, regulatory considerations, use of new control methods, orientation of new personnel, and dissemination of research findings.

Control Applications Technology

24. When a program manager assesses an aquatic plant control problem, the amount of information available on a given control, e.g., chemical herbicide, is overwhelming. There are newsletters, journal articles, knowledgeable personnel, books, and manuals offering relevant information. When a manager requests a literature search, just the number of titles alone can overwhelm him, with no attempt to read or understand the articles. He needs an expeditious answer regarding the application of a control to a specific problem, rather than being inundated with more information than he can possibly consider. Here the expert system truly works, its purpose being to convert information into knowledge that can be used in the field. From user input, the expert system directs the computer to sift through the information

available on a specific herbicide and identify relevant solutions to a particular problem.

Regulatory Considerations

25. The use of virtually all aquatic plant controls are regulated to some extent, and regulations are constantly changing. A control may be legal for application one year, and illegal the following year. Laws governing a treatment often vary in different states. It is difficult for a manager faced with an aquatic plant problem to have a full grasp of the meaning of a regulation and how it applies to his problem in a particular state. Workshop participants determined an expert system could provide the manager with current information regarding the regulation of controls.

26. The expert system could contain knowledge about the regulations appropriate to aquatic plant control. For example, if the use of a certain chemical is illegal in California but legal elsewhere, the expert system determines the state or states involved. If it is California, the chemical would be ruled out as a potential control agent. If the manager notices that a chemical used last year is not recommended by the expert system, the expert system can be queried as to why the chemical was not recommended. The expert system would respond that use of this chemical is no longer allowed in the manager's area. The expert system could also contain considerations of impacts on endangered species. There are sometimes problems with different institutions having different lists. The expert system would search the data base containing all lists and inform the user of proper considerations.

New Products

27. Field personnel pointed out that when a new method is approved for the control of aquatic plants there is usually a deficiency of both knowledge and experience about the product. The manager has innumerable questions about the method's effectiveness and use for the specific waterways for which he is responsible. The information available is often sketchy and too general or, conversely, too voluminous to find the information appropriate to the application. A current, accurate, and well-managed expert system provides the knowledge in a form specifically useful to the aquatic plant manager. The expert system asks the manager about the plants to control and specific details about

the site, such as water temperature, depth of water, or other environmental factors. If the new product is appropriate, the expert system automatically recommends it with instructions for application under local conditions.

Orientation of New Personnel

28. There is a continuous flow of new personnel into aquatic plant management positions as people are transferred, retire, or accept other positions. Considerable time and training are required before a new manager can become knowledgeable in available control methods and their proper use. An expert system allows a new employee to benefit from the accumulated knowledge. That knowledge in the expert system reduces the training requirements for new personnel and allows them to become productive in a shorter time frame.

Research Findings

29. It is extremely difficult for the manager to remain knowledgeable of the latest research developments. Participants pointed out that WES and other organizations conduct research and accumulate experience on methods of control. Technology transfer of the knowledge may be delayed or be ineffective. For example, research may develop treatment schedules and methods superior to the original documented recommendations. A specific treatment method needed may be contained in a technical report or article that addresses a large number of new technologies, thereby escaping the manager's attention. The field application of aquatic plant controls can be thought of as an extension of research, with more experience and knowledge gained by the managers each year. A superior method of using a control may be identified by persons in the field actually doing the work. From experience, a manager may learn that certain techniques that should work, do not work under his particular conditions. However, it is extremely difficult to disseminate this kind of knowledge. An aquatic plant expert system would be under constant development and the latest research and experience incorporated into the system. By accessing the expert system, the manager has the best opportunity for making full use of research results. The workshop participants emphasized that district personnel could not keep the expert system updated. They suggested that a formal procedure be established to evaluate and update information concerning successes and failures during field activities.

Recommendations

30. It is the consensus of the workshop group that building an expert system for aquatic plant management and control is both desirable and feasible, and the overall expert system program should have the following capabilities:

- a. Operate on an IBM PC computer or compatible clone.
- b. Contain current knowledge regarding new products.
- c. Be easy to use, with little training and accessible by the aquatic plant control managers.
- d. Be kept current. This may require a small staff with access to specialists.
- e. Reflect the knowledge gained from experiences of the plant control managers in the field.
- f. Reflect the knowledge obtained from on-going research; not only Corps of Engineers research, but research conducted by other organizations as well.
- g. Reflect the requirements imposed by environmental regulations and be current with new and changing regulations.
- h. Explain the reasoning for the recommendations it makes.
- i. Explain why particular control methods were not recommended.

31. Major expenses for the expert system program for aquatic plant management are the cost of establishing the system and updating the system with new or revised information. This would require a staff with access to specialists. Specialists can be charged with responsibility of updating those portions of the expert system for which they are knowledgeable. As discussed earlier, the operational responsibility for maintaining and distributing a current expert system should be someone other than District personnel.

32. A further recommendation by the workshop group is that it would not be wise to attempt to address the overall expert system initially. A more prudent approach would be to build a small prototype expert system. This would give users the opportunity to evaluate the capabilities of an expert system and to demonstrate its effectiveness. Appropriate applications for the pilot study were considered. Workshop participants determined that an excellent prototype would be the knowledge contained in the manual by Westerdahl and Getsinger (1988). Developing the guide into an expert system, even though a prototype, would be immediately useful to the aquatic plant managers and would provide an appropriate technology transfer application in a user-friendly format of state-of-the-art information.

PART IV: SUMMARY

33. In summary, the workshop participants suggested several areas in the aquatic plant control program that could be improved with the development of expert systems. An expert system for aquatic plant control would provide a single source for knowledge, thus eliminating lengthy searches and conflicting information. An expert system would provide a ready means for field personnel to capture knowledge and pass it forward to other persons in the field. The knowledge should still be verified by the appropriate experts. By placing the new knowledge into an expert system, this knowledge becomes available to other field personnel and increases their effectiveness. Participants determined that an updated expert system would provide the manager with current information regarding the regulation of controls. Aquatic plant control knowledge contained in an expert system would reduce the training requirements for new personnel and allow them to be productive in a shorter time frame. An aquatic plant expert system would be under constant development and the latest research results incorporated into the system as they become available.

REFERENCES

- Frenzel, Louis E., Jr. 1987. Understanding Expert Systems. Howard W. Sams and Company, Indianapolis, IN.
- Harmon, Paul, and King, David. 1985. Expert Systems, Artificial Intelligence in Business. Wiley, New York.
- Hart, Anna. 1986. Knowledge Acquisition for Expert Systems. McGraw-Hill, New York.
- Hayes-Roth, Frederick, Waterman, Donald, and Lenat, Douglas. 1983. Building Expert Systems. Addison-Wesley, Reading, MA.
- Kemp, William P., Onsager, Jerome A., and Lemmon, Hal E. 1988. "Rangeland Grasshopper Treatment Selection: An Expert System for Decision Support in Resource Management," AI Applications in Natural Resource Management, Vol 2, No. 4, pp 1-8.
- Lemmon, Hal. 1986. "Comax: An Expert System for Cotton Crop Management," Science, Vol 233, pp 29-33.
- McGraw, Karen L., and Harbison-Briggs, Karan. 1989. Knowledge Acquisition: Principles and Guidelines. Prentice-Hall, Englewood Cliffs, NJ.
- McKinnon, J. M., and Lemmon, H. E. 1985. "Expert Systems for Agriculture," Computers and Electronics in Agriculture. Elsevier Science Publishers B.V., Amsterdam, The Netherlands.
- Waterman, D. A. 1986. A Guide to Expert Systems. Addison-Wesley, Menlo Park, CA.
- Westerdahl, H. E., and Getsinger, K. D. 1988. "Aquatic Plant Identification and Herbicide Use Guide," Technical Report A-88-9, Vol II, US Army Engineer Waterways Experiment Station, Vicksburg, MS.

APPENDIX A: INFORMATION PACKAGE

Feasibility of Using
Expert Systems
in
Aquatic Plant Management

Feasibility of Using
Expert Systems
in
Aquatic Plant Management Programs

Objectives

- To determine if the knowledge and reasoning required to control aquatic plants can be entered into an expert system that can be used by managers in the field.
- To provide managers, faced with problems in aquatic plant control, with a method for determining the best solution for the problem.

Expert Systems

One method of transferring knowledge held by a specialist or several specialists is by means of an expert system.

Expert systems are special computer software applications that are capable of carrying out reasoning and analysis functions in narrowly defined areas at proficiency levels approaching levels of the human expert.

The study of expert systems is a subfield of the computer science field known as artificial intelligence. Currently there are approximately 3,000 expert systems operating in the United States.

Many expert systems are of the diagnostic type. For example, there are expert systems for diagnosing problems with an automobile electrical system, a high performance disk drive, a diesel locomotive, and a stuck drill pipe on a drilling rig.

An expert system typically performs as follows:

- Asks questions about the problem.
- May instruct the user to perform tests and report the results.
- Diagnoses the problem.
- Recommends an action to solve the problem.

Expert systems are designed by a team consisting of:

- Experts

The person or persons who are experts in the field.

- The knowledge engineer

Expert Systems and Aquatic Plants

The person who can convert the knowledge from an expert into a computer for reasoning and analysis.

- Users

Those supported by the expert system. Users use the expert system and also play an important role in debugging an expert system. They often provide additional knowledge that is added to the expert system. For example, the principal users of the XCON Expert System (to configure VAX computers) are also the experts on how to configure computers.

Rule-Based Expert Systems

There is no limit to the variety of ways that expert systems can be developed. However, in the past few years there has been an acceleration in the popularity of rule-based expert systems.

Rule-based expert systems have many advantages:

- It is easy to think in terms of rules and facts.
- It is easy to enter rules and facts into the computer, thus eliminating time consuming programming.
- It is easy and fast to build a prototype to test the feasibility of using an expert system to solve a problem.
- After gaining experience building a prototype, it is easy and inexpensive to begin again, using another approach.
- After a satisfactory prototype is built, it is easy to modify and extend it to a comprehensive final system.

How a Rule-Based Expert System Works


There are three parts to a rule-based expert system.

- Rules
- Facts
- Inference engine

Graphic Representation of an Expert System

A set of rules and facts are prepared. They contain the knowledge and reasoning required for the expert system to perform.

The expert system requests additional facts from the user.



Graphic Representation of an Expert System

The inference engine applies the rules and the facts and infers (hence the name inference engine) from these a conclusion and recommendations.

Hybrid Expert Systems

A principal disadvantage of a rule-based system is that sometimes the problem or some parts of a problem cannot be expressed in rules. For example, the rate at which aquatic plants grow is expressed better as a mathematical formula or several mathematical formulas depending on temperatures, day length, nutrients, etc. It would be impossible to express this as a set of rules.

In these cases we use hybrid systems. We use rules where they are appropriate, and call in and execute mathematical subroutines when needed. Comax/Gossym, an expert system for the management of cotton, is a hybrid. Comax is the rule-based expert system, and Gossym is a model of the cotton plant. HOPPER, an expert system for control of grasshoppers on rangeland is essentially a rule-based system, but calls upon mathematical programs to compute the rates at which grasshoppers grow and the amount of forage they consume.

Rule-based expert systems also have an educational advantage. It is possible to design the system in such a way that it can explain its recommendations.

For example, if the expert system recommended c as a control for alligatorweed, this recommendation could be questioned and the system

Expert Systems and Aquatic Plants

would explain its reasoning. For instance, alligatorweed is an above-water plant, and the controls that are applicable and effective in a short period of time are a, b, and c, with c being the least expensive.

The system can also explain why a different recommendation was not made.

For example, the user might ask why not use white amur, and the expert system would reply, "White amur will not control plants with growth above water."

Other computer systems can be programmed to explain their results but it is easier with rule-based systems.

Expert System Shells

A wide variety of expert system shells are available. An expert system shell is a system of programs that provide a means for entering rules and facts into the computer plus an inference engine that executes those rules and facts interactively with the user.

Commercial shells range in price from \$100 up to \$60,000. In 1985, the US Department of Agriculture (USDA) purchased two expert system shells (named ART), \$40,000 each, and also purchased \$120,000 Symbolics LISP computers to run them on. Comax, the cotton crop management expert system, was developed in this way.

In 1986, USDA purchased the VP-Expert package for \$100 to run on the PC computer. It was used to develop HOPPER, the grasshopper management program described in one of the attachments.

There is also an excellent shell named CLIPS, developed by NASA, to run on the PC. This shell is patterned after ART and is free to Government agencies.

Workshop on the Feasibility of Using Expert Systems in Aquatic Plant Management

The Aquatic Plant Control Research Program (APCRP) proposes to conduct a 1-day workshop to evaluate the feasibility of using expert systems in aquatic plant management.

Dr. Hal Lemmon from the Agricultural Research Service of the USDA, under contract to the APCRP, will present an overview of expert systems and discuss other systems similar to a possible Aquatic Plant Management expert system.

The participants of the workshop, prior to the workshop, are asked to brief themselves on the concepts of expert systems by reading or browsing the three articles attached to this package.

The participants will be invited to discuss the aquatic plant control problem.

Expert Systems and Aquatic Plants

- We will attempt to identify aspects of these problems as seen from the point of view of the field managers and from the point of view of researchers.
- We will define the scope of the problem, identify specific objectives, and decide criteria for considering the proposed expert system a success.

The point of contact for the workshop is:

Dr. Larry Lawrence
Resource Analysis Group
(601) 634-2778

Articles About Expert Systems

The following articles about expert systems are attached:

- Comax, an Expert System for Cotton Crop Management
- Rangeland Grasshopper Treatment Selection: An Expert System for Decision Support in Resource Management
- Expert Systems for Agriculture

Expert Systems and Aquatic Plants

Dr. Hal Lemmon from the Agricultural Research Service of the USDA, under contract to the APCRP, will present an overview of expert systems and discuss other systems similar to a possible Aquatic Plant Management expert system.

The participants of the workshop, prior to the workshop, are asked to brief themselves on the concepts of expert systems by reading or browsing the three articles attached to this package.

The participants will be invited to discuss the aquatic plant control problem.

- o We will attempt to identify aspects of these problems as seen from the point of view of the field managers and from the point of view of researchers.
- o We will define the scope of the problem, identify specific objectives, and decide criteria for considering the proposed expert system a success.

The point of contact for the workshop is:

Dr. Larry Lawrence
Resource Analysis Group
(601) 634-2778

Articles About Expert Systems

The following articles about expert systems are attached:

- o Comax, an Expert System for Cotton Crop Management
- o Rangeland Grasshopper Treatment Selection: An Expert System for Decision Support in Resource Management
- o Expert Systems for Agriculture

(Reprinted with permission from author and AAAS.)

Reprint Series
4 July 1986, Volume 233, pp. 29-33

SCIENCE

Comax: An Expert System for Cotton Crop Management

HAL LEMMON

Copyright © 1986 by the American Association for the Advancement of Science

Comax: An Expert System for Cotton Crop Management

HAL LEMMON

Expert systems are computer programs that perform at the level of human experts. One expert system, Comax, has been developed that acts as an expert in cotton crop management. The system has a knowledge base consisting of a sophisticated cotton plant simulation computer program, a set of "if-then" rules, and a computer program called an inference engine. Comax determines the best strategy for irrigating, applying fertilizer, and applying defoliant and cotton boll openers. Sensors in the cotton fields automatically report weather conditions to the system, and Comax reevaluates its recommendations daily. Comax was tested on a large farm and demonstrated excellent results in reducing the unit costs of production.

TODAY THREE BALES OF SYNTHETIC FIBERS ARE MILLED FOR every bale of cotton. Further, the synthetic fiber industry has recently adopted a vigorous research program to produce fibers at still lower cost. For cotton to survive, research to lower production costs is imperative (1).

An expert system, Comax (COtton MAnagement eXpert), has been developed that advises cotton growers on crop management at the farm level. The expert system is integrated with a computer model, Gossym (from *Gossypium* and simulation), that simulates the growth of the cotton plant (2). This is the first integration of an expert system with a simulation model for daily use in farm management.

Gossym

Researchers began developing Gossym in 1973. The program was developed over 12 years with contributions from ten scientists at four institutions (3) in two countries. It simulates the growth and development of the entire cotton plant on an organ-by-organ basis: roots, stems, leaves, blooms, squares, and bolls. It also simulates soil processes such as the transfer of water and nutrients through the soil profile. For Gossym to accomplish this, it needs data from mechanical and chemical soil analyses of the farm field to which it is being applied. Such analyses can be performed by state-owned soil test laboratories, the Soil Conservation Service, or commercial laboratories. The specific data required are soil hydrologic properties, soil fertility, soil impedance (resistance to root growth), water release curves, and bulk density.

The model is driven by weather variables. It requires, on a daily basis, such data as the maximum and minimum temperatures, solar radiation, and rainfall. It was developed with SPAR (Soil-Plant-Atmosphere-Research) units, where cotton is grown under highly

controlled conditions and the various rate processes can be determined, but it was extensively tested and validated against field data.

Gossym is capable of running on most computers, including microcomputers. A complete simulation, from emergence to harvest, can be done in 6 to 8 minutes on a VAX 750 computer, in 60 to 90 minutes on a microcomputer (an IBM PC, or equivalent, with a math coprocessor), and in 20 to 30 minutes on an advanced microcomputer (an IBM PC-AT, or equivalent, with a math coprocessor).

The development of microcomputers has expedited the movement of Gossym to the farm to assist in crop management. In 1984 a project to use Gossym on cotton farms was initiated by the USDA Agricultural Research Service in cooperation with the National Cotton Council, and microcomputers were provided for a 6000-acre farm in the Mississippi Delta (4) and a 1000-acre farm in the South Carolina Coastal Plain (5). In 1985 Comax was tested on the 6000-acre farm.

In the research laboratory, a multidisciplinary team of cotton experts provides Gossym with input and interprets its output. Comax was developed to provide the input and to perform the analyses when Gossym is used for practical, on-farm decision making. This is the first attempt I am aware of to integrate an expert system with a simulation model with the objective of optimizing crop production.

Comax

An expert system is a computer system with the capability of performing at the level of human experts in some particular domain. It is possible to build expert systems that perform at remarkable levels (6). While there are several methods for designing expert systems, rule-based systems have emerged as the popular architecture. Deriving their knowledge from relatively easily understood facts and rules, rule-based systems offer surprising power and versatility (7).

Comax is a rule-based expert system that operates Gossym the way a human expert would to determine three factors: irrigation schedules, nitrogen requirements, and the crop maturity date.

As shown in Fig. 1, Comax consists of a knowledge base, an inference engine, Gossym, a weather station, and data (for example, the seeding rate and soil parameters). The knowledge base is a set of rules and facts written in near-English. The inference engine examines the rules and facts to determine what is to be done. It prepares data files accordingly to hypothesize the weather and to hypothesize applications of water and nitrogen. Then it calls Gossym, which reads the data files prepared by the inference engine and simulates

The author is a computer scientist for the USDA Agricultural Research Service, Mississippi State, MS 39762.




Fig. 1. The Comax components. The four components to the right reside in a microcomputer located at the grower's farm.

the growth of the cotton plant under the conditions specified in those files. Results from the simulation (such as the day the simulated crop goes into water stress) are saved as facts in the knowledge base.

The inference engine program and the Gossym program change little if at all. The knowledge base continuously changes as researchers and growers improve management strategies or observe the impact of different strategies.

Software, Hardware, and Data

The software components of Comax are the inference engine and Gossym. The inference engine is written in the LISP computer language, and Gossym is written in FORTRAN. The computer languages were selected on the basis of appropriateness for the task to be performed, LISP being appropriate for an expert system but inappropriate for simulation. The knowledge base, so far, has about 50 rules, the inference engine about 6000 lines of code, and Gossym about 3000 lines of code.

Comax was developed on a Symbolics 3670 computer and is down-loaded, unchanged, to the PC computers where it runs under Common LISP, offered by Gold Hill Computers. Gossym was developed on the VAX 750 computer and is also down-loaded, unchanged, to the PC computers and compiled using the FORTRAN 77 compiler offered by Ryan-McFarland.

The cotton grower who used Comax has a microcomputer (an IBM PC or equivalent) with a math coprocessor and a dot-matrix printer in his office. The system can automatically call the weather station daily by telephone but, if a phone line is not practical, the data may be entered into the computer manually. The microcomputer costs \$4000 to \$7000, depending on the configuration selected. The cost of the weather station is \$4000, which includes solar panels to provide power. Hardware for telephone connection is \$1200.

Comax Rules

Figure 2 shows some of the facts and one of the rules used in Comax. This rule, "find-water-stress-day," is one of the set of rules used to determine the optimum irrigation schedule. The rule is true if every term in the "if" part of the rule matches a term in the facts base. In this case, (run-number ?number) of the rule matches the fact (run-number 1) if ?number is assigned the value 1, and (hypothesized-weather ?weather) matches the fact (hypothesized-weather hot-dry) if ?weather is assigned the value hot-dry. Entries that begin

with a question mark, such as ?number, are treated as variables by the inference engine and are assigned values, as needed, to cause a match.

In the case shown in Fig. 2, the rule is true, and the inference engine will proceed with the actions in the "then" part of the rule. It first prints on the computer screen a message describing the action. Next, it runs the Gossym program using the hot-dry weather scenario. When Gossym is finished, the inference engine examines the results of the run and places new facts into the facts base. One of the new facts will be, for example, (w-stress-day 236), where 236 represents the day of the year the crop went into water stress.

The final action of the inference engine is to assert a new fact, (set-hypothesis-irrigation), into the facts base. The purpose of this new fact is to cause another rule, which is called "set-up-hypothesized-irrigation" and is not shown in the figure, to be true. That rule, a lengthy one, determines the day that irrigation should be applied. Conceptually, it does this by taking the water stress day, subtracting the application time given in the fact (irrigation application-time 4), determining the amount of water to be applied from the fact (irrigation amount 1), and asserting a new fact (hypothesized-irrigation 232 1). However, there are actually other considerations, such as how soon to harvest and how many days since the last irrigation, which this rule also considers.

Comax recomputes the optimum management scenario each day, prints a daily report that recommends crop management procedures and, if it is desired, summarizes the intermediate simulations to explain the basis for the recommendations. Comax can show the results of simulations either by tabular reports or by graphs on the dot-matrix printer.

Operating Comax on the Farm

Comax is designed to run continuously throughout the crop year on a dedicated microcomputer. Each day it computes the expected irrigation date, the expected date and amount of fertilization, and the expected date of crop maturity. These are computed daily because, as the hypothesized weather for each day is replaced by the actual weather for that day, the computed dates change.

Determining irrigation requirements. Comax begins each day by determining the expected irrigation date. It does this by running Gossym with a hypothesized weather scenario, noting the date the crop goes into water stress and subtracting the number of days it takes to apply the irrigation. Some irrigation systems, the center-pivot type, for example, take several days to apply water. Comax uses three different types of hypothesized weather scenarios: (i) normal weather, (ii) hot-dry weather, and (iii) cold-wet weather. The weather scenarios are specific to each farm. Comax first runs Gossym with the hypothesized hot-dry weather scenario. This establishes the earliest date that irrigation would be required. Comax then runs Gossym with the normal weather scenario to determine the most likely date that irrigation will be required. The results are presented in a report printed at the end of the daily Comax operation.

The report states, for example, that today is 1 July and irrigation will be required on 10 July if subsequent weather is hot and dry or on 17 July if subsequent weather is normal. The next day, 2 July, the hypothesized weather for 1 July is replaced with the actual weather for 1 July, and the irrigation requirement is redetermined. If 1 July was a cold and wet day, the new report may state that irrigation is required on 12 July if subsequent weather is hot and dry (instead of 10 July as reported the day before) or on 19 July if the subsequent weather is normal (instead of 17 July). Conversely, if 1 July is actually a hot and dry day, the irrigation date for hot-dry weather will still be 10 July, but the irrigation date for the normal weather

er hypothesis will be earlier, perhaps 15 July instead of 17 July.

Determining nitrogen requirements. With cotton, it is important not to overfertilize, not only because of the obvious economic waste but also because overfertilization can cause the plant to be in an undesirable state at time of harvest. To determine the nitrogen requirements, Comax first ensures that there is no water stress by calculating an additional series of irrigation dates. After each calculation Comax determines the day the simulated crop went into water stress and, on the basis of the assumption that the grower would irrigate to relieve that stress, it hypothesizes a date and amount of irrigation. It then runs Gossym again to determine the next date that the crop will be in water stress. This process is repeated until the end of the season is reached, and the result is an hypothesized irrigation schedule that should prevent the crop from ever being in water stress. This schedule is only for use in determining nitrogen requirements and is never followed. The actual irrigation schedule to be followed is determined as described in the previous section.

Comax is now ready to determine the minimum amount of nitrogen that can be safely applied. It does so by making a series of Gossym runs with the cold-wet weather scenario, to simulate the minimum plant growth and thus to estimate the minimum nitrogen requirement. Comax again makes a series of these Gossym runs and, after each run, the day the crop went into nitrogen stress is noted. Comax then enters into the calculation a predetermined amount of nitrogen, and runs Gossym again. If nitrogen stress occurs again, the amount of nitrogen hypothesized is increased. When too much nitrogen is applied, there will be an undesirable effect: after the bolls are mature, the plant will begin to grow vigorously. If such undesirable growth (shown in Fig. 3, row 4, third graph) occurs, Comax reduces the amount of nitrogen. This process is repeated until Comax has determined the amount of nitrogen just sufficient to relieve nitrogen stress. This value is printed in the Comax daily report and represents the minimum amount of nitrogen the grower should apply.

The process is repeated with the normal weather scenario. This tells the grower the most probable nitrogen requirement. Finally, the process is repeated a third time with the hot-dry weather scenario, and the result tells the grower the maximum nitrogen requirement. From these three figures and from his own assessment of the weather the grower decides the amount of nitrogen to apply.

The grower's safest strategy is to assume the cold-wet weather scenario will hold and apply the minimum amount of nitrogen. If the weather turns out to be better than this, the grower can apply additional amounts of nitrogen later in the season. The penalty for underestimating the nitrogen requirement is only the cost of applying the additional nitrogen. The penalty for overestimating the nitrogen requirement is the cost of the excess nitrogen plus, at harvest, the loss from its undesirable effects, which can be substantial.

There is an additional risk that nitrogen applied too early in the season can be lost because of leaching. Such a loss varies with soil conditions, rainfall, and irrigation. Gossym is capable of identifying the amount of nitrogen lost in this way.

Farms that do not have irrigation systems are handled in a different, simpler manner. Farms with trickle irrigation require a different set of rules, a problem which will be addressed this year.

Determining harvest date. Comax also informs the grower when the cotton is mature so he can apply defoliant and boll openers. This is particularly important in such locations as the Mississippi Delta, where early rains can physically damage the cotton, induce boll rot, and make the ground so muddy that the mechanical cotton pickers cannot operate. Near the end of each season the grower must decide either to wait until it is certain the cotton has reached its

FACTS

```
(run-number 1)
(hypothesized-weather hot-dry)
(irrigation amount 1)
(irrigation application-time 4)
...
```

RULE find-water-stress-day

```
IF
  (run-number ?number)
  (hypothesized-weather ?weather)
THEN
  (printout "Finding water stress day") .
  (run-gossym ?number ?weather)
  (assert (set-hypothesized-irrigation))
```

Fig. 2. Four of the facts and one of the rules used in Comax. The rules are discussed in the text.

maximum yield or to proceed with the harvest before the rains begin. With Comax, the farmer knows weeks in advance when his crop will mature. This can only be an approximation because of uncertainty in the weather; but as each day passes, the hypothesized weather is replaced by the actual weather, and the projected maturity date becomes more reliable.

Comax in operation. An example of the operation of Comax as it selects nitrogen and irrigation schedules is shown in Fig. 3. The graphs in each row are the results of a Gossym simulation run by Comax. In the first graph of each row, the circles represent nitrogen applications. The first three applications are actual, but the fourth application (on the first graph of rows 3, 4, and 5) is hypothesized by Comax. On this farm the grower has applied 55, 60, and 30 pounds of nitrogen per acre at the time of planting and at 33 and 63 days after the plants emerged, respectively. The line shows the nitrogen stress, computed as the ratio of the nitrogen used to the nitrogen needed by the plant for full growth of all organs. In the second graph of each row, the jagged line represents a measure of water stress in the plant, and the vertical bars indicate the amount of water applied or that is expected to be applied by either rain or irrigation. The third graph of each row shows the height of the plant, the number of squares (unpollinated flower buds), and the number of bolls. The number of squares increases with time and then decreases as some squares are shed (because of stress) and others turn to bolls. The fourth graph of each row shows the development of the predicted yield. The final yield, in bales per acre, is printed above the curve.

The first row of graphs were produced by Comax just after the third application of nitrogen. The second row of graphs is the last of a series of Gossym runs in which Comax has directed its attention to the water stress problem and hypothesized a heavier irrigation schedule with no additional nitrogen. The second graph of this row shows that increased irrigation resulted in reduced water stress and in intensified nitrogen stress. With increased water, the simulated plant has the capacity for increased growth, and therefore it needs even more nitrogen. Even though irrigation is increased, there is no increased yield.

In the third row, Comax has hypothesized an application of 30 pounds of nitrogen per acre. The nitrogen stress is reduced, and the yield is increased.

In the fourth row, Comax has hypothesized an additional 60 pounds of nitrogen per acre. The nitrogen stress is eliminated, and the yield has increased correspondingly. However, the third graph

of this row shows that, after the bolls have all matured, the cotton plant has had a spurt of new growth and that it has started adding new squares that will never mature. At the point where the yield levels off, the crop should be harvested since no more cotton would be expected and delay would increase the risk of harvest losses due to inclement weather. To harvest cotton with modern equipment, it is necessary to apply a defoliant; however, this model plant would be so robust that the defoliant would not be as effective as it should be. The rules of Comax will cause this hypothesis to be rejected.

In the last row, Comax has selected 40 pounds of nitrogen per acre in conjunction with the indicated irrigation applications. This provides the maximum yield subject to the constraint of no secondary growth.

Constraints, such as irrigation capacity and the time required to irrigate, are provided for in the knowledge base. For example, on a

field with pivot irrigation a typical constraint may be that 1 inch of water can be applied in 4 days. Constraints are considered on a farm-by-farm basis; as a consequence, the knowledge base varies somewhat from farm to farm.

Results from a Pilot Test

Comax was tested on the Mitchener farm (4) so that we could acquire experience in its practical operation under realistic conditions (8). In mid-July 1985 Comax predicted the need for nitrogen at the rate of 50 pounds per acre, as shown in the last row of Fig. 3. As a result, the grower, who had not planned to apply any additional nitrogen, applied 20 pounds per acre throughout the farm except on a 6-acre test plot where no nitrogen was applied on alternate eight-




Fig. 3. Graphs produced by Comax from the results of Gossym simulations, showing the process whereby Comax reduces the water stress and then the nitrogen stress, as described in the text.

row strips. Comax predicted an additional 200 pounds of cotton lint on the cotton treated with nitrogen, with no delay in the date of maturity. At the end of the season, the test plots were picked, some by hand and some by machine. Although cotton is no longer picked by hand for commercial purposes, some rows of the test plot were so picked to obtain a precise figure to compare with the yield predicted by Comax. The hand-picked rows showed a net increase of 180 pounds per acre of cotton, and the machine-picked rows a net increase of 115 pounds per acre. The additional cotton (machine-picked) had an economic value of about \$71 per acre, the cost of the nitrogen was \$4 per acre, and the cost of application was \$5 per acre. Allowing for the cost of processing the additional cotton, there was a net gain of over \$60 per acre on this 6000-acre farm.

The grower believes, however, that it is the system's ability to pinpoint the day the crop is mature that is its most valuable feature. In the previous year (1984), the system predicted a maturity date of 1 September for the crop. Instead, the grower elected to use the widely accepted rule that a crop is not mature until 60% of the bolls are open and delayed harvesting until 21 September. Rain began on 6 October, and it was not possible to complete the harvest until November, which resulted in a loss of both yield and quality. The grower now believes that the maturity date of 1 September was correct and that, if the harvest had begun on that date, cotton production would have increased by approximately 4.3 million pounds and the quality would have been improved by an amount worth an additional \$0.11 per pound.

Future Outlook

During the coming crop year (1986), testing and development of Comax is continuing with 15 growers in five states and with a total cultivation of over 50,000 acres of cotton.

In the United States, there are 10 to 12 million acres (varying from year to year) of cotton on 30,000 farms. Approximately 1300 farms (4%) are of 1000 acres or more and account for 33% of the cotton, whereas 4000 farms are of 500 acres or more and account for 58% of the production (9). The former are obvious candidates for Comax; the latter are probable candidates.

REFERENCES AND NOTES

1. R. J. Kohel and C. F. Lewis, Eds., *Cotton* (American Society of Agronomy, Madison, WI, 1984), pp. xi-xii.
2. D. N. Baker, J. R. Lambert, J. M. McKinion, S.C. *Agric. Exp. Sta. Tech. Bull.* 1089 (December 1983).
3. The scientists and their institutions are: D. N. Baker, J. M. McKinion, R. E. Fye, USDA Agricultural Research Service, Mississippi State, MS; F. D. Whisler, J. A. Landivar, D. R. Reddy, S. G. Kharche, A. Ben-Porath, Mississippi State University, Mississippi State, MS; J. R. Lambert, Clemson University, Clemson, SC; and A. Marani, Hebrew University of Jerusalem, Rehovot, Israel.
4. Mitchener farm, Sumner, MS.
5. McCoy farm, Sumner, SC.
6. R. O. Duda and E. H. Shortliffe, *Science* 220, 261 (1983).
7. R. Davis, *ibid.* 231, 957 (1986).
8. F. Mitchener, *Cotton Grower* 2, 42 (1986).
9. A. Jordan, National Cotton Council, Memphis, TN, personal communication.
10. I thank D. N. Baker, A. R. Grable, J. M. McKinion, and G. S. Hasegawa for their critical reading of the manuscript and for their many helpful suggestions.

Abstract. *Of the insect pests on rangeland in the western United States, grasshoppers are by far the most serious. Appropriate treatment selection for rangeland grasshoppers is a complex problem. Little information is available to land managers who wish to evaluate whether or not to spray, or when to control with what, based on site-specific conditions. In addition, concerns about more cost-effective selections, better timing of applications, uncertainty of action thresholds, environmental sensitivity, and a lack of local expertise all indicated the need to develop a rangeland grasshopper treatment selection support tool for ranchers and land managers. A prototype rule-based expert system was developed that requires minimal computer capabilities and yet is powerful enough to allow consideration of a wide array of environmental and economic scenarios. The expert system was developed to utilize site-specific input easily obtained by potential users. Output from the expert system provides users with appropriate treatment selections and benefit/cost ratios. This system was made available to extension agents, land managers, and ranchers early in the fall of 1988.*

Rangeland Grasshopper Treatment Selection: An Expert System for Decision Support in Resource Management¹

William P. Kemp,² Jerome A. Onsager,² and Hal E. Lemmon³

Management intensity of rangelands in the western United States has significantly increased in recent years. There is also growing awareness among ranchers and land managers concerned with rangeland pests that control activities should be selected and scheduled to maximize efficacy and/or save forage with minimal environmental disturbance. Of the potential insect pests, grasshoppers are by far the most serious that land managers and ranchers face on rangeland. Several recent studies estimate forage losses to exceed 20 percent of that available annually (Hewitt 1977; Hewitt and Onsager 1982, 1983). Though the magnitude of the grasshopper problem was recognized over 100 years ago, little information is available to land managers who wish to evaluate whether or not to spray, as well as when to control with what, based on site-specific conditions. This lack of information is due at least in part to the complexity of the problem. There are

over 25 species of common grasshoppers, each with its own biology, feeding on the wide array of forage plants in western rangelands.

The Problem

Even though few tools exist to help land managers make grasshopper management decisions, millions of acres of rangeland are sprayed annually with chemicals. At present, the USDA Animal and Plant Health Inspection Service-Plant Protection and Quarantine (APHIS-PPQ) is charged with the control, not man-

¹ Paper presented at the Third Workshop on AI and Related Topics, USDA Forest Service, Northern Training Center, Missoula, Montana, 20-21 April 1988.

² U.S. Department of Agriculture, Agricultural Research Service, Rangeland Insect Laboratory, Bozeman, Montana 59717-001.

³ U. S. Department of Agriculture, Agricultural Research Service, 800 Buchanan, Albany, California 94710.

agement (control is a subset of management activities), of rangeland grasshoppers on federal lands, and has been operating in this role since the late 1940s. APHIS-PPQ also participates in grasshopper control on state and private lands subject to minimum block-size constraints and cost sharing. According to the APHIS-PPQ Final Environmental Impact Statement (APHIS 1987), eight grasshoppers per square yard (9.6 per square meter) is an administrative action threshold used throughout rangelands of the western United States. Given the complexity of the grasshopper problem and the differences in ecotypes and economics, it is unlikely that any one grasshopper density is an appropriate trigger for all control activities. Nevertheless, eight grasshoppers per square yard has been used by land managers and ranchers as both an economic injury level and an economic threshold. The economic injury level (EIL) is the lowest density of grasshoppers that will cause economic damage (Pedigo et al. 1986, Stern et al. 1959). The economic threshold (ET) is the time (expressed as pest density) when economic damage will probably occur in the future if no control is imposed, and pest numbers are merely an index of that time (Pedigo et al. 1986). Use of a single-density measure as an EIL and ET no doubt has resulted in unneeded (non-economically justifiable) control activities in some areas and years, as well as failure to recognize other areas and years where control activities would have resulted in economic gain.

The weakness of eight grasshoppers per square yard as an EIL or ET is further illustrated when one understands the source of this figure. APHIS (1987) states that eight grasshoppers per square yard is the density above which grasshoppers compete with cattle for forage. Parker (1939) is cited as the source of this information. However, when describing population trends of grasshoppers from year to year, Parker stated:

The following year may again result in twice the population of the previous year and there would be only eight per square yard, which is enough to cause slight injury to crops but is not enough to cause much comment.

Parker (1939) also stated that densities of 24 to 32 or more grasshoppers per square yard are high enough to cause severe damage to crops and should be considered outbreak densities. It is alarming that Parker's original statement of an ET for crops has been misinterpreted over the years to be a generally accepted EIL for rangeland grasshoppers.

Given that a rancher or land manager has determined that some sort of control is needed for rangeland grasshoppers, he/she is still faced with two additional problems. First, he/she is faced with selecting a treatment from an array of chemicals or biological insecticides applied as sprays or on carriers such as bran bait. Second, he/she must decide, based on the appropriate biotic variables, the correct timing of the control measure of choice. The complexity of the problem is increased by the fact that the treatment selection is influenced in part by aspects of timing.

For more than 50 years, the Rangeland Insect Laboratory (RIL) in Bozeman, Montana, has been involved in the management of rangeland grasshoppers and has published numerous articles on their biology, ecology, and management. For example, recent work by Hewitt and Onsager (1982) resulted in the development of a method to estimate potential forage consumption by grasshoppers within a given year based on initial densities. This work was conducted for three years at one site in Montana. Hewitt and Onsager (1983) and Onsager (1986, 1987a, 1987b) provide new perspectives on timing of control activities to maximize efficacy. Perhaps the only existing work on estimating EILs for rangeland grasshoppers is Onsager (1984), which provides methods

useful for determining when to use malathion or carbaryl for grasshopper control. Torell et al. (1987) developed a spreadsheet program for assessing the economics of rangeland grasshopper control programs. However, neither Onsager (1984) nor Torell et al. (1987) consider directly the ability of a specific site to produce forage or how that forage production capability influences the determination of an EIL for rangeland grasshoppers. There also does not presently exist any single source (guide, computer program, etc.) that a rancher or land manager can use for rangeland grasshopper management.

The major objective of the work described here was to develop a simple and easily applied computer decision support tool that land managers could use to develop site-specific control decisions for rangeland grasshoppers. A secondary objective of this effort was to determine the applicability of expert system technology for this resource-based decision support problem. Our intention was to incorporate the influence of land capability, insect densities, expected uses of forage, and weather to illustrate the importance of using currently available and easily obtainable information to improve the rangeland grasshopper treatment selection process in the western United States.

Methods

System Design Considerations

The decision support tool described herein was developed with several *a priori* constraints. First, required site-specific inputs must be easily obtainable. Second, the resulting system should be small enough to run quickly on IBM-PC or compatible machines. Third, the model should be flexible enough to cover a wide range of scenarios that can be encountered by the user. Finally, the system should be a significant improvement over existing methods.

These constraints were developed from meetings with agricultural extension agents, as well as APHIS-PPQ, Forest Service (USFS), and Bureau of Land Management (BLM) personnel.

Approaching the problem from a simulation modeling perspective, it was obvious that the first and second constraints above would be extremely limiting. Also, if a simulation model were developed from those few locations where detailed ecological studies had been conducted, the resulting system would likely have limited value beyond those sites. There also exists a number of significant data gaps that would prevent the development, at this time, of a widely applicable decision support system based solely on simulation models. Therefore, we considered the possibility of using expert system technology.

In assessing the appropriateness of expert system technology for this problem, we considered criteria similar to those suggested by Stock (1987). First, expertise on rangeland grasshopper treatment selection is scarce. However, of five to seven recognized experts nationwide, three currently reside at the RIL. The three experts at the RIL have a total of more than 75 years of experience with rangeland grasshoppers. Further, all three of these experts use heuristics gained from their long experience when called upon to make a recommendation about treatment selection, and have worked cooperatively for a long enough period so as to be in general agreement about the appropriate treatment selections. There also exists an adequate number of test cases and potential users to validate individual components as well as the entire system.

The Way Decisions are Made

Our human experts used similar methods to solve a specific grasshopper treatment selection problem. In general, the expert, when confronted with a potential problem,

would first consider all possible treatments as potentially applicable. Next, the expert would ask a series of specific questions, the answers to which would successively exclude different treatment options (i.e., environmentally sensitive areas, predominant grasshopper type, current local weather conditions). Other questions posed by the expert were aimed at determining whether it was too early or too late to treat with certain options (i.e., development stages of grasshoppers present, percent of grasshoppers in the adult stage). This method of considering all options until sufficient information invalidated all but the best possible subset is called *contra-indication* and is commonly used in expert systems where diagnoses are made (Lemmon, unpublished).

To develop an expert system within the *a priori* constraints above, we used VP-Expert.⁴ VP-Expert is a rule-based expert system shell that has a number of powerful development features that permit rapid prototyping and debugging (Latham 1988). System requirements for VP-Expert are minimal (greater than 256K RAM, one DS-DD diskette drive, and DOS Ver. 2.0 or later) and the inexpensive and unlimited annual runtime dispensing license made this shell desirable from a distribution standpoint.

System Overview

We separated the problem into two parts. First, the expert system determines all possible treatments considered acceptable according to the rules established by our human experts. Selections are based on scientific and technical reasons. Then the system determines, for each acceptable treatment, the cost of the treatment and the value of the benefit, and ranks them according to their benefit/cost ratio. This approach allows us to consider a treatment that is not applicable today, but that might be applicable at a later date. For example, it may be prefer-

able to apply treatment B next week, instead of applying treatment A now, even though additional forage losses will result prior to treatment B.

A series of meetings with the lead expert (J. A. Onsager) established the general pattern in which decisions were made. An initial prototype was developed and presented to him for criticism. This process was repeated four times during construction, each time improving accuracy and flexibility. Other experts (G. B. Hewitt and J. E. Henry) were questioned on specific aspects related to their expertise as the prototype was developed.

Information and heuristics used by the experts were captured in the form of if/then propositions (about 100 at present) within the VP-Expert context (Fig. 1). VP-Expert uses backward chaining as its problem-solving method, but we were able to forward chain using the FIND statement and subsequently used both problem-solving methods in the final system. As noted previously, the method of contra-indication was used throughout in the initial selection of possible treatments. That is, instead of using rules to determine what treatment could be used, the rules were written to determine what treatments could not be used. This is the same approach that the human expert used to determine the appropriate treatments. The user is queried (Fig. 2) for information on:

- the state; for example Montana.
- existence of environmentally sensitive areas; for example, lakes and streams.
- current weather conditions; for example, rainy.
- the predominant grasshopper and development stage; for example, spur-throat and fourth instar.

A treatment could be contra-indicated for the present but considered at a later date; for example, in the

⁴ Paperback Software International, 2830 Ninth Street, Berkeley, California 94710.

A) RULE WEATHER 1

IF weather = cool_wet

THEN contra_ind_now = malathion
 contra_ind_now = acephate
 contra_ind_now = carbaryl_bait
 contra_ind_now = *Nosema*_bait

BECAUSE "Some control options work better than others under different weather conditions. For example, malathion and acephate efficacy is significantly reduced under moist conditions. Baits, when wet, become unattractive to grasshoppers.";

B) RULE START 2

IF sens_env = yes

THEN contra_ind_now = malathion
 contra_ind_fut = malathion
 contra_ind_now = acephate
 contra_ind_fut = acephate
 contra_ind_now = carbaryl_spray
 contra_ind_fut = carbaryl_spray
 contra_ind_now = carbaryl_bait
 contra_ind_fut = carbaryl_bait

ELSE start = yes

BECAUSE "If the area in question is environmentally sensitive, applications of chemical sprays or chemical treated baits cannot be considered now or in the future.";

Figure 1 Form of if/then propositions used to capture system knowledge and examples of method of contra-indication within the if/then context.




Figure 2. Simplified conceptual flow of information through an expert system for rangeland grasshopper treatment selection showing two levels of user input.

case of wet weather (Fig. 1A). A treatment could also be contra-indicated for the future if environmental constraints were too limiting (Fig 1B). The users' answers to these queries are combined with the knowledge captured in the rules to arrive at a possible solution set of treatments (Table 1).

Minimal questioning was employed throughout to prevent a number of potential user-interface problems (Schmoldt 1987). However, if a user does not know the answer to a particularly important question, a second level of reasoning is pursued to obtain an answer at a reduced level of certainty. At present, certainty factors are not available to the user, though we are exploring appropriate ways to use them in future versions of this system.

For the computation of benefit-cost ratios, the user is asked for estimates of grasshopper densities, the value of an Animal Unit Month (AUM) of forage equivalents, and is asked to modify default application costs if they differ from what he/she knows about the particular situation. This design was employed so that a user could explore the relative differences in actual treatment costs. For example, a rancher may not find it economical to treat with malathion if he/she must be responsible for the entire cost of application, but this may not be true if he/she participates in a state or federal cost-sharing program that reduces the rancher's per-acre obligations.

At present, the benefit-cost ratio computations are simplistic; they only consider expected forage replacement costs and application costs. Data from Onsager (1984) were used in part to develop a function that computes expected forage destroyed as a function of grasshopper density. This value is then converted to AUM equivalents and a dollar value is computed for the expected loss based on user input.

At any point in the session, a user may query the system on WHY a

Table 1. Possible selected treatments.

Treatment	Application Date
malathion spray	today
malathion spray	one week from today
carbaryl spray	today
carbaryl spray	one week from today
carbaryl bait	today
carbaryl bait	one week from today
acephate spray	today
acephate spray	one week from today
Nosema bait	today
Nosema bait	one week from today
do nothing	

Table 2. Summary of input data for a simplified example run.

Variable	Value
state	Montana
sensitive environment	no
current weather conditions	hot-dry
predominant development stage of grasshoppers	adult
predominant type of grasshopper	spur-throat
period of grasshopper hatch	two weeks
proportion of grasshoppers in adult stage	<75%
objective of treatments	save forage

Table 3. Benefit-cost ratios computed from standard runs (Table 2) over a range of possible densities and AUM values (cost of malathion application \$2.00/acre).

	Density per square yard			
	8	15	20	25
AUM Value				
\$ 5.00	.28	.48	.61	.71
10.00	.56	.97	1.21	1.43
15.00	.85	1.45	1.82	2.14

particular question is being asked, and the system will respond with answers specific to that question. The user may also select WHAT/IF scenarios to rerun a session (without exiting) and evaluate the results of changes in specific inputs. Upon completion of a session, the user may request a hard copy summary of results (Fig. 3).

Results and Discussion

Verification and validation of components of the system were conducted after the initial prototype was developed. A number of modifications were made as a result of meetings with appropriate domain experts (J. A. Onsager, grasshopper population dynamics, chemical control; G. B. Hewitt, grasshopper ecology; J. E. Henry, biological control of grasshoppers). Once alterations were made and final computational testing of the prototype was conducted, the system was deemed suitable for on-site testing. System development time, from initial meetings to final prototype, consisted of about nine months. This system was sent out for field testing at about 10 test locations during the fall of 1988. Input that we obtain from users (from a questionnaire sent with the system) will contribute to system improvement.

Action Threshold

Some potential impacts of the system can be illustrated with a very simple example. We will consider the problem, stated earlier, of the arbitrary and generally accepted action threshold for rangeland grasshoppers. Table 2 contains simplified input (selected or direct) for a particular scenario. Using this reduced set of input values, we examined the range of benefit-cost ratios of only one of the possible selected treatments (malathion) that resulted from variable inputs for AUM value and grasshopper density (Table 3). Note here that the selected treatment also

HOPPER Summary of Input Data

Variable	Value
state	Montana
sensitive environment	no
current weather conditions	hot-dry
development stage	adult
species	spur-throat
hatch	under two weeks
objective	save forage
proportion population adult	under 75%
density, grasshoppers per square yard	15/sq yd
animal unit month equivalent	800 lbs.
animal unit month value	\$15.00
malathion cost	\$2.00/acre
acephate cost	\$2.00/acre
carbaryl spray cost	\$3.00/acre
carbaryl bait cost	\$2.25/acre
Nosema bait cost	\$4.50/acre

Explanation of values:

- a. Values shown were entered by the user, or were inferred by HOPPER from values entered by the user.
- b. Values left blank were not needed in the consultation.

HOPPER Recommendations

Applying malathion now has a benefit/cost ratio of 1.45
 Applying malathion later (a week from now) has an approximate benefit/cost ratio of 1.22

Applying acephate now has a benefit/cost ratio of 1.45

Applying carbaryl spray now has a benefit/cost ratio of 0.97.

Figure 3. Example of output from an expert system for rangeland grasshopper treatment selection that may be requested by the user.

influences benefit-cost ratios through heuristics related to treatment-specific efficacy rates. The interaction of density and AUM value on the resulting benefit/cost ratios is striking. The AUM values selected for consideration represent a reasonable array of expected replacement costs over the land types considered in this system. However, for rangeland in the simplified case presented here, treatment of eight adult grasshoppers per square yard with AUM value of \$5.00 is not economical under any of the

densities considered. Generally speaking, however, as the value of an AUM (the replacement costs of forage lost to grasshoppers) decreases, greater densities are required in order to justify the costs of control (assuming fixed application costs). Even in this simple example, the inappropriateness of a single EIL is apparent, especially since it is common for forage replacement costs to vary monthly during some parts of the year. It also suggests that if the rancher's actual treatment applica-

tion costs were reduced, through, for example, participation in a government-sponsored program, the benefit-cost figures could change to the point where it could be profitable to treat, even though densities and AUM values were low (Table 3).

In this example, our results are similar to those of a more detailed economic model developed by Torelli et al. (1987). However, the Torelli et al. model is designed more for research than management. In addition to requiring a large amount of input from the user, it allows only comparisons of benefit-cost ratios for treatments selected by the user *a priori* and does not consider the environmental conditions that our system does. It is very important to base grasshopper treatment selection on current and site-specific environmental factors, as well as benefit-cost ratios (Onsager 1987b, Torelli et al. 1987). The system that we have developed offers land managers and ranchers the opportunity to consider a wide range of site-specific environmental factors, as well as associated costs of all control treatments (both biological and chemical) currently registered for grasshoppers on rangeland.

We found the use of contra-indication very helpful in developing this system. We expect that this method will find additional use in other pest management systems (i.e., forests, crops) where treatment selection is the end goal. Also, the use of VP-Expert will no doubt be more common in the future. The rapid prototyping capabilities as well as other features (i.e., excellent editor, trace features) make this a surprisingly complete development tool for about \$100. The CHAIN feature will allow very large systems to be developed, if necessary, without the need for a great deal of RAM (say, 640K). This package continues to improve concurrent with user demands.

Work is continuing on the development of subsystems, such as detailed phenology (Kemp 1987a, Kemp and


Onsager 1986) and probability of outbreak (Kemp 1987b) models that can easily be linked to the present system architecture. Future plans also include linking geographic information system attributes for the purpose of expanding forecasting capabilities.

Acknowledgments

We thank G. B. Hewitt and J. E. Henry at the USDA/ARS Rangeland Insect Laboratory for their advice during system development and manuscript reviews.

References

- APHIS. 1987. Rangeland grasshopper cooperative management program final environmental impact statement. U.S. Department of Agriculture, Animal and Plant Health Inspection Service-Plant Protection and Quarantine, 87-1. Washington, D.C. 20250.
- Hewitt, G. B. 1977. Review of forage losses caused by rangeland grasshoppers. U.S. Department of Agriculture, Agriculture Research Service Miscellaneous Publication number 1348.
- Hewitt, G. B., and J. A. Onsager. 1982. A method for forecasting potential losses from grasshopper feeding on northern mixed prairie forages. *Journal of Range Management* 35: 53-57.
- Hewitt, G. B., and J. A. Onsager. 1983. Control of grasshoppers on rangeland in the United States - a perspective. *Journal of Range Management* 36: 202-207.
- Kemp, W. P. 1987a. Predictive phenology modeling in rangeland pest management. Pages 351-368 in: *Integrated Pest Management on Rangeland: A Shortgrass Prairie Perspective*, J. L. Capinera, editor. Westview Press, Boulder, Colorado.
- Kemp, W. P. 1987b. Probability of outbreak for rangeland grasshoppers (Orthoptera: Acrididae) in Montana: application of Markovian principles. *Journal of Economic Entomology* 80: 1100-1105.
- Kemp, W. P., and J. A. Onsager. 1986. Rangeland grasshoppers (Orthoptera: Acrididae): modeling phenology of natural populations of six species. *Environmental Entomology* 15: 924-930.
- Latham, D. 1988. VP-Expert: a review. *AI Applications in Natural Resource Management* 2(1): 60-61.
- Onsager, J. A. 1984. A method for estimating economic injury levels for control of rangeland grasshoppers with malathion and carbaryl. *Journal of Range Management* 37: 20(3)-203.
- Onsager, J. A. 1986. Stability and diversity of grasshopper species in a grassland community due to temporal heterogeneity. Pages 101-109 in: *Proceedings, 4th Triennial Meeting, Pan American Acridological Society*, 28 July-2 August 1985, Saskatoon, Saskatchewan, Canada.
- Onsager, J. A. 1987a. Current tactics for suppression of grasshoppers on range. Pages 60-66 in: *Integrated Pest Management on Rangeland*, J. A. Onsager, editor. U.S. Department of Agriculture, Agriculture Research Service, ARS-50.
- Onsager, J. A. 1987b. Integrated management of rangeland grasshoppers. Pages 196-204 in: *Integrated Pest Management on Rangeland: A Shortgrass Prairie Perspective*, J. L. Capinera, editor. Westview Press, Boulder, Colorado.
- Parker, J. R. 1939. *Grasshoppers and Their Control*. U.S. Department of Agriculture, Farmers Bulletin no. 1828.
- Pedigo, L. P., S. H. Hutchins, and L. G. Higley. 1986. Economic injury levels in theory and practice. *Annual Review of Entomology* 31: 341-368.
- Schmoldt, D. L. 1987. Implementing design considerations in a forest pest diagnosis expert system. In *Proceedings, First Annual Conference on INSIGHT/LEVEL5 Expert Systems Applications*. Melbourne, Florida, 9-10 November 1987. USERS.PRL, Ramsey, New Jersey.
- Stern, V. M., R. F. Smith, R. van den Bosch, and K. S. Hagen. 1959. The integrated control concept. *Hilgardia* 29: 81-101.
- Stock, M. 1987. AI and expert systems: an overview. *AI Applications in Natural Resource Management* 1(1): 9-18.
- Torell, L. A., K. M. Garrett, J. H. Davis, and E. W. Huddleston. 1987. A spreadsheet program for assessing the economics of rangeland grasshopper control programs. *Western Computer Consortium Journal of Computer Applications* 2: 49-57.



William P. Kemp is a research entomologist at the USDA/ARS, Rangeland Insect Laboratory in Bozeman, Montana. He also has adjunct appointments in the Department of Biology and in the Department of Entomology at Montana State University. He received a B.S. in wildlife management and forestry from the University of Maine in 1976, an M.S. in entomology from Michigan State University in 1978, and a Ph.D. in forestry, wildlife, and range sciences from the University of Idaho in 1984. Since then, he has worked on grasshopper ecology with emphasis on developing manager-oriented decision support tools.



Jerome A. Onsager is Research Leader at the USDA/ARS Rangeland Insect Laboratory in Bozeman, Montana. He received a B.S. (1958), an M.S. (1960), and a Ph.D. (1963) in agricultural entomology from North Dakota State University. He joined ARS in 1963, studied biology and control of wireworms at Yakima, Washington (1963-1969), and led vegetable insect investigations at Charleston, South Carolina (1969-1974). Since transferring to his present position in 1975, he has published over 50 technical papers on biology, population dynamics, sampling, and integration of chemical, biological, and cultural tactics for control of rangeland grasshoppers. He is also Research Manager for a federal nine-

agency cooperative Grasshopper IPM Experimental Demonstration, with responsibilities for accelerated research on parasites, predators, pathogens, population dynamics, and an expert system.



Hal E. Lemmon is a computer scientist with the USDA/ARS, Agriculture Research Service. He received his B.S. and M.S. in mathematics, and Ph.D. in chemical engineering from the University of Utah. He is the author of COMAX, an expert system for cotton management that has made a significant impact on production aspects in that industry. He is currently working on a textbook for LISP programming.

SUBSCRIPTION INFORMATION

AI Applications in Natural Resource Management is published quarterly (January, April, July, and October) for developers and users of artificial intelligence technology in natural resource fields. Its purpose is to provide timely and informative articles that represent the state of the art of AI in natural resource management and to keep readers posted on AI-related matters.

Prepaid individual subscriptions \$25; billed individual subscriptions \$30; government agencies, libraries, private companies, and foreign subscriptions (including Canada) \$45 US. Send subscriptions to: AI Applications, P.O. Box 3066, Moscow, Idaho 83843 USA. (A limited number of charter subscriptions is still available. If you'd like your subscription to start with volume 1, number 1, Summer 1987, please ask.)



Lorraine Ashland

Computers and Electronics in Agriculture, 1 (1985) 31-40
Elsevier Science Publishers B.V., Amsterdam - Printed in The Netherlands

(Reprinted with permission from the authors and Elsevier Science
Publishers B.V., Physical Sciences and Engineering Division.)

EXPERT SYSTEMS FOR AGRICULTURE*

J.M. McKINION and H.E. LEMMON

USDA-ARS Crop Simulation Research Unit, P.O. Box 5367, Mississippi State, MS 39762
(U.S.A.)

(Accepted 6 May 1985)

ABSTRACT

McKinion, J.M. and Lemmon, H.E., 1985. Expert systems for agriculture. *Comput. Electron. Agric.*, 1: 31-40.

Recent advances in computer technology have been made possible the development of Expert Systems. Expert Systems are special computer software applications that are capable of carrying out reasoning and analysis functions in narrowly defined subject areas at proficiency levels approaching that of a human expert. The prime targets for the development of expert systems applications in agriculture are the narrowly defined subject areas which have experts available for solving problems. All commercial crop production systems in existence today are potential candidates for Expert Systems. These Expert Systems would take the form of integrated crop management decision aids which would encompass irrigation, nutritional problems and fertilization, weed control-cultivation and herbicide application, and insect control and insecticide and/or nematicide application. Additional subject areas of potential are plant pathology, salinity management, crop breeding, animal pathology, and animal herd management. The advantage of Expert Systems is that once developed they can raise the performance of the average worker to the level of an expert.

INTRODUCTION

Artificial Intelligence (AI) as a field of research consists of four main subtopics: robotics, natural-language interpretation, computer vision, and expert systems (ES) (Rich, 1983; Hayes-Roth et al., 1983. Until recently, AI research required the use of dedicated million-dollar computers. Now the advent of supermicrocomputer LISP machines, which cost much less than mainframes and in some cases have a much larger memory address space than the old mainframes, have made AI cost-effective for a variety of

*A contribution of the USDA-ARS Crop Simulation Research Unit in cooperation with the Agronomy Department, Mississippi State University and the Mississippi Agricultural and Forestry Experiment Station.

The use of company names and brand names is for information only and does not represent an endorsement or warranty by the USDA.

applications. LISP machines provide high-quality graphics interfaces which improve the people/machine communications and make it far easier to develop large, complex AI applications. AI experts have also changed their point of view and narrowed their focus to a smaller, more practical domain of problems. As these factors have evolved, AI applications have begun to move from the laboratory to the commercial domain (Anonymous, 1983).

EXPERT SYSTEMS

Of the four principle areas of research in AI, Expert Systems currently offer the most promise for immediate applications solving computer programs that achieve a high level of performance in some specialized problem domain considered to be difficult and requiring specialized knowledge and skill. They have the following characteristics:

- (1) heuristic — they employ judgemental as well as formal reasoning in solving problems;
- (2) transparent — they have the ability to explain and justify their line of reasoning;
- (3) flexible — domain-specific knowledge is generally separate from domain-independent inference procedures, thus knowledge updating is made considerably easier than in conventional programming.

The emphasis in Expert Systems (ES) is on symbolic representation and inference rather than the numerical approach of traditional programming languages. ES contain two components. One of these is called the knowledge base. The knowledge base contains in some symbolic manner the knowledge of facts, judgements, rules, intuition, and experience about a particular problem area. The other component is called an inference mechanism. It can interpret the knowledge in the knowledge base. It can also perform logical deduction and knowledge base manipulations. The objective of an Expert System is to raise the performance of the average worker to the expert level (Santarelli, 1984).

KNOWLEDGE BASE

The inference mechanism is essentially static. However, the knowledge base grows and expands as the expert behind it adds more knowledge to it. The knowledge base, like a database, stores information. The comparison of the commonality of a database and a knowledge base end here. The combination of symbolic representation of knowledge within the knowledge base, various kinds of knowledge-base structures, and relationships between the structures, make it possible to represent common sense information.

Some of the ways used to represent knowledge in a knowledge base are scripts (used mostly in natural language systems), logic, processes, rules, frames, and semantic nets. In general, any knowledge that can be represented by one method can be represented by the others. The choice

of method depends on how the knowledge engineer chooses to think about the knowledge and which representation lends itself most efficiently to retrieval and deduction of facts.

A semantic net uses both predicates and attributes to represent objects and to show relationships between the objects. A typical representation might be ROBIN IS-A BIRD, SPARROW IS-A BIRD, BIRD IS-A ANIMAL. In this case ROBIN, SPARROW, BIRD, and ANIMAL are nodes in a network and the links in the network represent the relation IS-A. The network as a whole forms a taxonomy (Kinnucan, 1984).

Another symbolic knowledge representation structure found in knowledge bases is called frames. Instead of memory areas called fields, which a database uses to hold information about its data, frames have variable-sized memory areas called slots. The slots may contain standard attributes, like databases do and they may also contain hypotheses that relate to the expert program's function, rules about program situations and actions to take, subprograms, and pointers to other frames. This slot-to-frame transition creates a hierarchy not found in databases (Ham, 1984).

The most common form of knowledge base representation is rule-based. A rule is a conditional statement that specifies an action that is supposed to take place under a certain set of conditions. Rules in an AI program can be somewhat similar to if-then statements in conventional programming languages. However, most conventional programs contain only a relatively small number of possible paths at each step that calls for branching. In contrast, the conditionality embedded in AI problems is so great that the number of paths that can be exploited explodes combinatorially. In conventional programming, the rules are imbedded directly into the program and consequently require considerable effort to develop, debug, and maintain. In a rule-based system the rules are entered into the knowledge base without programming. The programmer does not have to worry about where the rule fits in the structure, system developers can add, modify, and delete rules with ease. Since system developers do not have to be concerned with proper sequencing and consistency, they can explore and rapidly prototype complex, ill-specified, ill-understood, changeable requirements — a characteristic of AI systems. The point of all this is that while a problem which is amenable to an expert system must be narrowly defined in order to be tractable, it does not necessarily have to be well understood. Using AI, crude prototyping can be rapidly developed which can hopefully be iterated on until a viable system emerges.

INFERENCE ENGINE

The inference engine solves a problem by interpreting the domain knowledge contained in the knowledge base. An inference engine is essentially a computer programmed to process symbols that represent objects. The computer reasons by processing these symbols. The most important symbol-

processing operations are matching two character strings, joining or separating two strings, and substituting one string for another. As conceptually simple as this is, such operations allow for automatic reasoning. Two reasoning mechanisms are commonly used in rule-based inference engines, either alone or in combination. In forward (data-driven) inferencing (also called forward chaining), the system attempts to reason forward from the given facts to a solution. In backward (goal-driven) inferencing (also called backward chaining), the system works backward from a hypothetical solution (the goal) to find evidence supporting the solution. Often this requires formulation and testing of intermediate hypotheses (subgoals).

LISP MACHINE

AI development generally requires large quantities of computer resources. Research in AI in the past was performed on large mainframe computers that were dedicated to this effort. With the advent of the supermicrocomputer, the lowering of computer hardware costs, and the steady increase in computer performance, LISP machines have now become widely available at minicomputer prices. However, there are still distinct differences between the LISP computers and conventional computers.

LISP computers are usually single-user machines. They are significantly different in computer architecture. LISP computers typically use a tagged architecture. A tag is placed in front of the computer word to designate the data-type. Special hardware allows data-type checking to be carried out at run-time, not just at compile time. This is important in a dynamic LISP environment to be compatible with the flexibility of the LISP language (Winston and Horn, 1981) and the generic nature of most functions to operate on many different data types. Run-time data-type checking ensures that the data types match the instruction. In this way, erroneous operations such as 'add this number to this character string' is avoided. A characteristic of a LISP environment is that when objects existing in that environment are terminated (made to be inaccessible) they do not automatically disappear and free up the memory space used. The computer must collect these unused objects and recover the memory. This process is called garbage collection. Some LISP machines have hardware assisted garbage collection. These machines also have the primitive instructions for the LISP language implemented directly into the hardware. All of these factors add up to a very significant advantage for the LISP machines. They are able to run LISP programs 5 to 10 times faster than conventional computers which have the same number of instructions per second rating. The LISP machines have multitasking capability along with high resolution graphics. To further enhance the very highly interactive programming environment, these machines have integrated a hand-held pointing device (mouse) to invoke software functions. The display is usually a bit-mapped screen that can be updated from 1 to 10 MHz data rate (Rich, 1983). A list of LISP machines

is given in Table 1. It is the authors' opinion that the price of the different computing systems is a very close approximation of the system's computational power. The quality of the software and the support is another matter.

TABLE 1

List of Artificial Intelligence computers

Manufacturer	Model	Approximate cost (US\$)
Symbolics	3670	100 000
Xerox	1132	130 000
LISP Machine	Lambda	72 000
Texas Instruments	Explorer	60 000
Tektronics	4404	15 000

LISP AND SOFTWARE DEVELOPMENT TOOLS

The language of AI is LISP (an acronym for LISt Processor). LISP has been in use since 1958 making it one of the oldest high-level languages in existence. LISP is a symbolic manipulation language and can handle predicate calculus logic. LISP programs are collections of independent procedures called functions. However, the developers of an Expert System may not need to program at all. There are many software development tools becoming available which give the AI programmer great power in developing software (Verity, 1984). Rapid and efficient development of Expert Systems is enhanced if a powerful development system is available which meets the needs of the developer. A list of these development systems is shown in Table 2.

TABLE 2

List of knowledge engineering tools

Vendor	Name	Price (US\$)
Intellicorp	KEE	60 000
Inference Corp.	ART	60 000
Carnegie Group	SRL+	70 000

LISP ON NON-LISP MACHINES

It is only in recent years that LISP machines have been available. Prior to this, AI systems were developed on mainframe computers or on large

minicomputers. Many large LISP systems have been successfully developed on non-LISP machines such as the IBM 370 or the DEC VAX. This includes expert systems for computer configuration and Expert Systems for equipment failure diagnosis. There is no question that AI systems can be successfully developed on non-LISP computers. With the explosive growth of the microcomputer market, it is no surprise that LISP languages are beginning to appear in the market place. The question arises as to whether viable AI systems can be developed on microcomputers.

It seems to be the consensus of veteran AI developers that microcomputers will make an important contribution in two areas, first in training, and second in downloading and executing of AI systems that were developed on larger computers. The microcomputer LISP programs are excellent for learning LISP and the concepts of AI programming. There are excellent LISP tutorial packages offered which compliment the training process. So LISP on microcomputers has the potential of playing a major role in AI training. Downloading of an expert system developed on a LISP computer onto a microcomputer is the approach that is being taken in the development of COMAX, which will be discussed later. The prototyping is on the Symbolics 3670 and the final system will be downloaded onto a microcomputer which accomodates a subset of COMMON LISP.

Developing large AI systems on a microcomputer is another question however. Most AI developers feel that the microcomputer does not yet have sufficient power to support the development of serious AI systems. Undoubtedly with time this will also change.

One additional question is in regard to the Expert Systems Tools which are beginning to appear for microcomputers. Here again the consensus is that these tools could provide excellent learning experiences, and that limited expert systems can be developed on them, but that they are not yet suitable for the development of a serious expert systems. As certainly as the same technology that has presently given us LISP machines for development of AI systems, that same technology will in time give us super-microcomputers that can cope with the demands of a serious AI environment (McKinion, 1980).

ARTIFICIAL INTELLIGENCE AND AGRICULTURAL RESEARCH

Let us now address the question of AI and its usefulness, if any, to agricultural research and researchers. There are several components to this question. First is the use of AI computers in research. The powerful LISP computers now on the market are single-workstation computers. They can only be used by one person at a time. These computers were designed to maximize the performance of the individual developing a system, and they typically have capabilities not available in traditional computing. In developing computer programs, for example, the work is done in an editor window. While working on the program it is only necessary to make a single

keystroke to compile the program. The nature of the architecture and the speed of the computer are such that the program is compiled almost before you can take your finger off the compile key. You can bring up the test window, again with a single keystroke, and make a test run. To go back to the editor takes another keystroke, and the cursor is pointing exactly to the point where you left off. Notice that the compilation was made without closing out the editor. The result of all of this speed and power is that software development is enhanced enormously. It is the experience of the authors that the development time is enhanced by a factor of 10. To place this in perspective consider that this means that a system which normally would take a year to develop using traditional methods can be developed in about a month. Systems which would require 3 year to develop can be developed in 3 months. It is also worth mentioning that this capability is not so much because of AI, but because of the power of the computers and of the power of development software used to support AI.

The high speed, large memories (the Symbolics 3670 we are using has 6.5 million bytes of RAM and almost 500 million bytes of disk) and the flexibility of these computers also support research users with large databases to analyze. Data can be moved in and out of the computer with ease. Plots can be made on the high resolution screen giving the researcher unprecedented access to his data. The impact of these machines on research can be awesome.

There are currently 57 experimental and commercial applications of Expert Systems as listed by Feigenbaum and McCorduck (1983) in their book, *The Fifth Generation*. Expert Systems have been applied already in a diverse number of disciplines: chemistry, medicine, genetic engineering, mineral exploration and others. Table 3 gives a generic classification of the Expert Systems that have been developed.

Regarding the use of expert systems and research, there seems to be several opportunities. For example, there is now a commercially available and highly successful expert systems to advise on experiment planning for determining DNA sequences. Other expert systems have been developed for experiment planning but have had limited success.

One of the first expert systems, DENDRAL, grew from a research need.

TABLE 3

Current knowledge-engineering applications

Medical diagnosis and prescription	Signal interpretation
Equipment failure diagnosis	Mineral exploration
Computer configuration	Military threat assessment and targeting
Chemical data interpretation and structure elucidation	Crisis management
Experiment planning	Advising about computer system use
Speech and image understanding	Very large scale integrated circuit design

The problem was to determine the three-dimensional structure of organic molecules. DENDRAL takes spectrographic data from nuclear magnetic resonators and mass spectrographs, coupled with empirical formulae and basic chemical knowledge and infers with phenomenal success the molecular structure. This expert systems took 18 years to develop and is highly successful.

Given the areas in which Expert Systems have already been developed, where in agriculture are the opportunities for application of this new technology? The answer to this question should come in part from the definition of Expert Systems: Computer software applications that are capable of carrying out reasoning and analysis functions in narrowly defined subject areas at proficiency levels approaching that of a human expert. The two key words in this definition are *narrowly defined*. The problem to be dealt with must also have at least one expert on the subject who has solved the problem. That is to say Expert Systems deal with applications and not with research. Expert Systems typically evolve most successfully where research talent in artificial intelligence is combined with subject expertise required to build a knowledge base for a specific application (Battelle Today, 1984).

The first and foremost opportunity for using Expert Systems technology in agriculture is with integrated crop management. The operations costing the most and having the greatest potential effects on crop yield should be addressed in these systems. Farmers, farm managers, extension specialists, county agents, Soil Conservation Service agents and others have to make high-risk decisions concerning management of their crops on irrigation, tillage, fertilization, pesticide applications and herbicide applications. Not only are the timings of these events important, but also the quantity or type are important. The USDA Agricultural Research Service, Crop Simulation Research Unit at Mississippi State, MS is currently developing a CrOp Management EXpert (COMAX) advisory system based on the dynamic cotton crop simulation model GOSSYM (Baker et al., 1983). This Expert System will incorporate the knowledge of developers of the cotton model which predicts crop growth and yield in response to external weather variables, soil physical parameters, soil fertility, and pest damage and the practical knowledge of the extension specialists. A production rule system (also known as IF-THEN rules) is being constructed. The COMAX, which is written in LISP, calls the FORTRAN model GOSSYM to acquire information to be put in the COMAX knowledge base. COMAX exercises the cotton model to find the optimum recommendation for management decisions on a daily basis to maximize cotton yields while minimizing user input to the crop system. Risk analysis will also be considered because some farmers can afford higher risk for possible higher payoffs than other farmers who cannot afford any risk. While the COMAX system is being developed on a LISP computer, the system will be downloaded onto a microcomputer for use. COMAX will use the full resources of the microcomputer and run

steadily 24 h a day. The COMAX system will use weather data acquired from an automatic weather station to bring the model of the crop up to current status. COMAX will then generate weather scenarios which will be fed to the model GOSSYM, which then predicts the growth and development of the crop. COMAX, by using the weather scenarios, management decisions, cultural practice information, and other non-automated information (insect scouting reports), then determines an optimum recommendation for today.

The list of Expert Systems for crop management is only limited by the number of agronomic crops. The Crop Simulation Research Unit is also developing simulation models of soybean and wheat crops, and these crops are also candidates to follow the COMAX system. There are many other areas in agriculture ripe for Expert Systems technology. Plant pathology, weed control, pest management, irrigation management, salinity management, crop breeding, and many other fields, some as stand-alone Expert Systems and others as adjunct advisory systems to crop management Expert Systems. Expert Systems are suitable for any task which requires judgement and manipulation of facts (Santarelli, 1984).

TECHNOLOGY TRANSFER

Perhaps one of the greatest problems today is that of transferring new technology from the laboratories of research to practical application. Expert System technology is the ideal conduit of new knowledge from the agricultural scientists' laboratory to usage at the farm level, the ultimate consumer of agricultural research. Expert Systems will not be static devices; they will be under continual development and improvement. As new knowledge is discovered, this information will need to be incorporated into the knowledge base, calling for a continuing commitment of Expert Systems developers. Expert Systems derive their power from knowledge rather than from a single powerful technique. "In the knowledge is the power" is the key concept of Expert Systems developers.

REFERENCES

- Anonymous, 1983. Artificial intelligence: a solution whose time has come. *Syst. Software*, 2(12): 150-177.
- Baker, D.N., Lambert, J.R. and McKinion, J.M., 1983. GOSSYM: a simulator of cotton crop growth and yield. *S.C. Agric. Exp. Stn. Tech. Bull.* 1089, 134 pp.
- Battelle Today, 1984. Number 40, November 1984, p. 5.
- Feigenbaum, E.A. and McCorduck, P., 1983. *The Fifth Generation*. Addison-Wesley, Reading, MA, 275 pp.
- Ham, M., 1984. Playing by the rules. *PC World*, 2(1): 34-39.
- Hayes-Roth, F., Waterman, D.A. and Lenat, D.B., 1983. *Building Expert Systems*. Addison-Wesley, London, 444 pp.
- Kinnucan, P., 1984. Computers that think like experts. *High Technol.*, 4(1): 30-42.
- McKinion, J.M., 1980. Dynamic simulation: a positive feedback mechanism for experimental research in the biological sciences. *Agric. Syst.*, 5: 239-250.
- Rich, E., 1983. *Artificial Intelligence*. McGraw-Hill, New York, NY, 436 pp.
- Santarelli, M.B., 1984. Artificial intelligence: impact on business. *Software News*, 4(1): 30, 31, 36.
- Verity, J., 1984. AI tools arrive in force. *Datamation*, 30(15): 44-53.
- Winston, P.H. and Horn, B.K.P., 1981. *LISP*. Addison-Wesley, Reading, MA, 430 pp.