The Potential Application of Expert Systems in Dairy Extension Education[1]

TERRY R. SMITH
Department of Animal Science
Cornell University
Ithaca, NY 14853

ABSTRACT

The application of artificial intelligence technologies encompasses a number of fields with opportunities for exploiting these tools to solve real world problems that traditional programming environments do not offer. Expert systems offer a structured approach to knowledge representation and use techniques to manipulate data in ways that may generate inferences not explicitly programmed. When encoded in an expert system, rules help guide the user through masses of data and the expert's reasoning strategies and rules captured in the system. An expert system approach to problem solving provides a flexible yet structured approach to many problems that extension specialists ("experts in a field") now solve relatively routinely. The data-intensive nature of dairy herd management analysis offers numerous opportunities to apply expert system concept to the monitoring and controlling of herd performance. The ability to query the expert system rule structure during a consultation provides the user the opportunity to view the flow of the rules used during the session, thereby increasing the user's expertise and providing an instructional experience. The expert system development process is necessarily iterative and therefore demands a highly flexible programming environment. An overview of factors to consider when evaluating the potential for using an expert system for a particular application and factors to consider when selecting an expert systems programming environment will be discussed.

INTRODUCTION

The objectives of this paper are: 1) to describe what an expert system (ES) is and how it is used to capture and ultimately deliver a base of knowledge to a target audience, 2) to describe what should be considered when evaluating the potential application of ES technology in the field, and 3) to discuss some of the greatest opportunity areas for applying ES in dairy extension education.

ARTIFICIAL INTELLIGENCE

The field of artificial intelligence (AI) has often been broadly defined as the study of how to make computers solve problems traditionally thought to require human intelligence. Commercially available products that have benefited from the application of AI to real world problems and situations have evolved from early computer science laboratories that focused on AI. Artificial intelligence has found applications in robotics, computer vision, speech recognition and voice synthesis, computer-assisted instruction, computer-aided software engineering, and ES applications, to name a few. The potential role of AI in agriculture was examined in a recent study supported by USDA and carried out at Purdue University (12). Whereas the field of AI is not clearly defined in agriculture, forms of AI that offer potential for application in agriculture are expert systems, natural language processing, and robotics. The Experiment Station Committee on Organization and Policy established a broad set of new research initiatives including computer technology for agricultural management and robotics and sen-

Received December 7, 1988.
Accepted April 21, 1989.
[1]Support provided by NC-119 Regional Research Project and the Agricultural Experiment Station, Cornell University.
[2]Invited paper.
[3]Current address: Department of Dairy Science, University of Wisconsin, Madison 53706.

sors in agriculture. The specific objectives include the development of expert systems for management decision making.

The basis for the application of AI is symbolic processing, which is the representation and manipulation of knowledge and information encoded as symbols. The symbols are used to represent objects, concepts, properties, and relationships in much the way that humans visualize the world and reason. Symbolic processing is in contrast to numerical processing in which data and information are reduced to numbers and formulas.

EXPERT SYSTEMS

Expert systems are computer programs containing a wide base of knowledge (knowledge base) in a restricted subject area or domain. They use complex inferential reasoning algorithms (comprising the inference engine) to perform tasks done by human experts to solve real world problems that traditional programming environments do not readily offer. Expert systems contain both declarative and procedural knowledge. Declarative knowledge is that typically found contained in databases and managed using database management software. Procedural knowledge is the course of action or rules. Expert systems offer a structured approach to knowledge representation and use techniques to manipulate knowledge in ways that may generate inferences not explicitly programmed (3). Rules help guide the user through the knowledge base during a consultation with the ES. An ES approach to problem solving provides a flexible, although structured, approach to many problems that extension educators now solve routinely. Responding to individual requests from decision makers or other educators often consumes a significant proportion of a specialist's time. Spahr et al. (17) described the potential for integrating data collected automatically from specialized animal sensors with ES technology to improve the decision-making ability of the dairy manager. Several authors have described how ES technology has been applied or could potentially be applied to a variety of agricultural problems (2, 5, 8, 10, 11, 16).

An important attribute of most ES applications is the ability provided to users to query the ES rule structure during a consultation session, thereby increasing the users' expertise and providing them with a better sense of the ES problem-solving strategies. The nature and functionality of ES development environments (often referred to as ES shells) are such that the end user has the ability to query the system any time during a session. This ability allows the user to "ask why" the system is asking a particular question. The query feature also allows the user to "ask how" the system has come to a particular conclusion or recommendation by requesting mapping of the rule-firing sequence and display (often graphically) for the user's review. It may also be possible for the user to modify a previous response and continue the session from the point at which it was halted. Some ES have been developed primarily for training purposes and to provide the user with the ability to view the expert's reasoning process. Instructional benefits obtained may outweigh the ES value as a decision aid since the user can change responses to ES queries and view the changes and the resulting conclusions or recommendations (3, 13).

Another feature of most ES applications is a natural language interface. A natural language interface provides the user (as well as the developer) of an ES with the ability to interact with the system in a relatively free-form question and response format using English sentences or phrases and provides users with little or no computer experience more direct and easy access to the ES application. Jones (6) used AI techniques to develop a natural language interface to a dairy herd management database application.

APPROACHING A POTENTIAL EXPERT SYSTEMS PROJECT

Expert systems applications provide the unique opportunity to capture the knowledge, expertise, and heuristics ("rules of thumb") of an expert before that person retires or moves to a different position. The ES also have the potential to capture, organize, and deliver expertise to the "field" without requiring the presence of the subject matter specialist. The potential benefits of developing ES applications for use in instructional or problem-solving situations must be weighed against the real cost associated with developing and maintaining the computer software.

The conventional software development process typically involves the end-user at the beginning of the project, but in the case of the ES project, experimentation with a module or a prototype can begin almost immediately after a discussion of the requirements with the expert(s) and the user(s). A small prototype can be developed to establish that the approach and the problem can be handled using an ES. Depending upon the size and scope of an application, the ES development process may require large amounts of the domain expert's time to acquire and represent adequately the person's knowledge and approach to the specific problem. Experts place the development time of a moderate to large size ES development project (200 to 1000 rules), requiring only limited computer skills to use it, at 5 to 10 person-yr (12). However, many newer commercial ES shells provide for rapid prototyping, particularly for smaller, very focused applications. Although persons in the academic world have typically not approached projects from a "return on investment" standpoint, the real costs and potential returns associated with any project must be recognized. Estimating the return on investment may be an important means for justifying a project. Return on investment is equal to the annualized cash benefits divided by the total investment or budget and is often used as a measure of the value of a project. Estimating the potential benefits and converting them to dollars will be a challenge in most situations. Recognizing the economic realities of software development efforts and the expected returns from delivering an expert's knowledge base to the field in a readily accessible and easy to use format will be important to long-term success of such efforts.

Computer applications have had and will continue to have a dramatic impact on the way the academic and research communities develop and deliver new knowledge. Recognizing the need to develop standards for publications with specific reference to computer software, the ADSA Journal Management Committee appointed a committee to develop such standards. The standards (15) are designed to serve as guidelines for authors of scientific publications that focus on computer software applications submitted to the *Journal of Dairy Science*. Individuals and project teams developing ES with application to dairy science should be encouraged to publish their findings in the *Journal of Dairy Science*. Emphasis should be placed on problem-solving procedure and performance of the ES in the field or in an instructional setting depending upon the purpose for which the ES was developed. Most ES programming environments provide the developer with the opportunity to capture consultation sessions to a structured data file for subsequent evaluation of the performance of the system. This could potentially provide the developer with quantitative information relative to the use being made by end-users of the system to solve particular problems and a characterization of their use of the system.

DESIGNING AN EXPERT SYSTEMS APPLICATION

The design phase of any project is critically important to the eventual success of the effort. The design and development of an ES application are not much different from traditional software development strategies. The initial stage requires the development team to identify the characteristics of the problem clearly. The major objectives of this first step are to ensure that the characteristics of the problem are well understood and a domain expert or domain experts are available and committed to the project. Another important component of the design phase is determining whether the project is technically feasible and whether an ES approach is most suitable to the problem. The next step involves conceptualizing the application and beginning to characterize the underlying knowledge requirements. This process helps to identify potential knowledge representation strategies. Having identified the basic requirements and conceptual approach to be taken, a structure should be developed to organize the knowledge. The tools that allow developers to represent the valuable knowledge and experience of domain experts are known as ES shells. The ES shells provide the developer with a programming environment and programming tools necessary to represent effectively the expert's knowledge and reasoning strategies in a computer system. The ES shell houses the reasoning strategies and methods used to represent the way an expert reasons through a problem and makes a conclusion or decision. One of the most intuitive ways to represent knowledge, the "if. . . .then" rule construct, is also probably the most commonly used approach in ES shells. The reasoning methods in the ES shell allow the end-user to navigate through the rules stored in the knowledge base. The inference engine of the ES shell uses the knowledge in the selected representation scheme to solve problems. A flexible inference engine provides a robust programming environment.

The two most common inferencing techniques are forward and backward chaining. Forward chaining is the simplest inferencing strategy and starts with an initial state and progresses through the rule base toward a goal state. This approach infers solutions from assertions existing in the knowledge base. This inference engine tries to infer everything possible from a given set of assertions. Forward chaining is initiated when the antecedent or "IF" portion of a forward rule matches a set of assertions in the knowledge base. When the rule is matched and ready to fire, the inference engine adds the results to the knowledge base. Forward chaining lends itself to applications with many possible goals (e.g., planning and design).

Backward chaining is used to generate assertions that satisfy a specific goal. Backward chaining narrows the search procedure to a specific goal, which allows for better utilization of resources in the system. Backward chaining searches backward from the goal to the conditions making the goal true. If the information is not known, the inference engine would search to find a rule that matches this goal and would find a rule that ends with the "then. . . .goal", and back track given the IF portion until the last assertion matches the goal. Backward chaining lends itself well to diagnostic applications. A flexible inference engine should be able to deal with a given knowledge representation scheme in more than one way. Because forward chaining creates new assertions or facts with the addition of a goal, backward chaining can operate on the same set of rules and facts. The integration of forward and backward chaining is often referred to as goal-oriented forward chaining. This approach allows the user to fire forward rules in an attempt to satisfy backward chaining goals. This flexibility is typically available in the higher end (more expensive) shells.

Another important feature of an ES shell is the ability for the developer to organize the knowledge base and to control the logical sequencing of a consultation session. A frames provision within a shell organizes a knowledge base into a treelike structure of subproblems. This provides the developer additional ways to represent knowledge to the system by integrating declarative and procedural representations, which increases the likelihood of suiting a developer's way of thinking about a problem. Each frame carries characteristics about itself, parameters, and rules involved in the subproblem solution. Parameters and rules also carry their characteristics, thus continuing the object-oriented structure of the system. Frames provide a familiar way to organize and structure information similar to that of a relational database. Instead of entering facts into the system in an unstructured manner, it is useful to be able to structure them. For example, the developer may want to group all the information about reproductive performance into one segment of the system's knowledge base. Organizing this related information into a frame provides this structure. The use of frames to organize a knowledge base into its component subproblems offers a more modular structure to the development process and therefore lends the development process to multisite development opportunities.

Meta-rules are another feature of some ES shells used to control rule order or to produce a more efficient and coherent consultation. This "meta-level" knowledge allows the system to make deductions about itself, to determine the most efficient strategy of operation, and to reorder rules considered, using the most relevant rules first. Citrenbaum et al. (1) discuss the features of many of the more popular commercially available expert system shells.

Because both the data and the expertise captured in an ES are often uncertain, assertions made without complete certainty can be assigned a certainty or confidence factor. The domain expert will often be challenged to obtain consistent estimates of subjective probabilities, as it may be difficult to separate probabilities from significance to the particular problem or situation. However, the ability to incorporate uncertainty into a knowledge base, as well as provide the user with the opportunity to apply certainty factors to their responses to queries from the system, makes for a robust decision support environment.

It is usually important to the development process to develop an initial prototype of the ES application. Depending upon the complexity of the application, the initial prototype may be developed by the domain experts themselves. A goal of ES research is to develop system building tools, which can allow the expert in a particular discipline or domain to build an ES without the need to consult with a "knowledge engineer" or builder of ES (12). With the price of ES shells available for the personal computer environment ranging from $99 to more than $10,000, it is often advantageous to select a quick-implementation shell (usually less cost) to develop the initial prototype. Developing an initial prototype has the advantages of causing the project team to begin to deal with the complexity of a particular problem and at the same time may serve to convince administrators or funding agencies of the potential of such an application without expending significant resources early in the development process. The next step involves expanding the initial prototype to include the full set of rules and corresponding structure including "user-friendly interfaces" to be contained in the final prototype. Depending upon the shell selected for the initial prototype, this may be the point at which a more flexible shell would have to be used. The next step in the ES development process is that of validating the system. Validation strategies depend to a great extent upon the ultimate end-user of the system but are obviously critical to the successful implementation of the ES application. The domain expert or experts must be satisfied with the performance of the system. This includes evaluating the consistency, completeness, and functionality of the system and the ease with which additions and modifications can be made to the system to improve performance.

EXPERT SYSTEMS AND CONVENTIONAL PROGRAMMING ENVIRONMENTS

Expert systems can begin to tackle problems requiring judgmental decisions that have not been easily incorporated into decision support systems using more traditional high level programming languages (e.g., FORTRAN, COBOL, BASIC, PASCAL). Electronic spreadsheet (e.g., Lotus 1-2-3™, Multiplan™, Excel™) have provided decision makers with tools that can be readily adapted to a wide range of situations and problems. The user interface of most spreadsheets is a familiar tabular format for entering data and formulae and therefore enhances the interaction between the software and the user. Although decision aides based upon electronic spreadsheets have been used effectively in dairy extension education, the potential for applying ES and ES interfaced to spreadsheet or databases will be significant as applications are developed and tested. Levins and Varner (9) used a rule-based approach in a compiled BASIC application to evaluate the reproductive management program of a dairy herd and recognized the potential for using an ES approach to enhance the application.

Many of the ES shells on the market for use in the PC environment include interfaces to some of the popular electronic spreadsheet and database management systems. This feature offers the opportunity to use the computational and database structure of an existing application in conjunction with the rule base of an ES, which may reduce the cost associated with moving the spreadsheet or database portion of the application into the ES application. Another important consideration may be the ability to integrate a computerized decision aid or a simulation model programmed in a high level language with an ES. Most shells provide the ability to pass information between the ES environment and other programs (13). However, several language interfaces provide access only to whatever arguments are explicitly passed to them through subroutine or function calls, thereby limiting interfacing to short procedures and data files. In some applications, the inferencing feature of an ES may play a minor role in the overall application, in which case the external program needs to have the ability to drive the inference engine much like a subroutine. There are also integrated systems on the market that include database, graphics, telecommunications, word processing, and spreadsheet capabilities fully integrated within the ES environment (e.g., "GURU™", mdbs, Inc., Lafayette, IN). When selecting an ES shell to be integrated with existing applications one must be sure to evaluate the flexibility for the current project as well as likely future needs (13).

APPLYING EXPERT SYSTEMS TO DAIRY MANAGEMENT PROBLEMS

The DHI program continues to be an integrally important component of most dairy extension education programs. As the management reports become increasingly sophisticated and the databases continue to expand in scope, the need to develop improved strategies for delivering this increasingly complex information to the field represents a formidable challenge.

Databases are often described from three viewpoints: physical, logical, or conceptual. The physical view relates to how data are stored within the computer system; this is the view that systems and application programmers have of the data within a database. Database management system software provides the logical view of the data to the application programmer and oftentimes to the end user. The third view of the data contained in a database is conceptual, the way the end user visualizes data and data relationships. Databases can become more accessible to the end user by including conceptual data definitions that allow the logical view of the data to approximate the conceptual view of the data more closely. An ES may offer the opportunity to make a database appear more accessible to the end user, particularly through its use of a natural language interface with the ES environment.

All nine Dairy Records Processing Centers (DRPC) in the US have addressed the need to provide users with direct access to their herd management database and related databases by developing programs for accessing and capturing data and reports from the farm (14). The DRPC are providing remote access to the databases using a variety of approaches (14). Coupled with the processing capabilities of microcomputers, microcomputer-based on-farm data collection applications have been developed for monitoring and controlling many aspects of the dairy production system (17). Decision support applications accommodating some on-farm processing of data downloaded from the DRPC mainframe is a real need and has already been given significant attention by DRPC, university faculty, and agribusiness professionals (2, 7, 10). The DHI databases are very dynamic and therefore require the dairy manager to monitor closely the individual cow and herd information in order to make sound management decisions. Expert systems have the potential to provide the end user with an expert's view and evaluation of a herd's DHI database. The USDA-ES National Core Parameters Committee described 10 major dairy management areas and charged each of 10 subcommittees to identify core parameters that could be used to monitor and evaluate management performance. This effort provides a framework for standardizing the computation and reporting of herd management performance measures across DRPCs (4). This will enhance the opportunity for multistate and national program efforts, thereby increasing the expertise and spreading the real cost of development and benefits to be accrued across a broader base. Knowledge acquisition and ES development costs can be amortized over the broad base of DHI participants. The remote access capabilities to these large databases will provide the opportunity to move data directly into the ES application and therefore reduce the chances that the end user will lose the motivation to employ the system by having to enter large amounts of data at each session. Due to the dynamic nature of the DHI database, routine execution of the ES application will be necessary and could even be automated to monitor changes in the data. Significant opportunities exist to apply an ES approach to the evaluation of dairy herd and farm management data using the knowledge and skill available from many dairy extension specialists in the US. Dairy farmers often suggest that they do not use their DHI records due to lack of time and lack of understanding as to the interpretation of the management reports. An ES approach for both education and problem-solving applications should be an important emphasis for those providing educational support for the DHI program.

In the future, the potential to link ES with mass media technologies such as interactive video disk, CD-ROM, and others that provide context-sensitive access to large repositories of full text and graphics databases presents some tremendous opportunities to extend an extensive knowledge base to a broad audience.

CONCLUSIONS

Expert systems provide the extension education system with some important opportunities as well as challenges. Production agriculture will continue to face challenges that will require producers to become better managers. Systems that provide the decision maker with ready access to the latest technical information and that dynamically monitor and evaluate the performance of all aspects of the management system will be essential to the future success of most dairy farm businesses. The application of computer hardware and software that can facilitate the manager's role should be expected to result in a more profitable business.

REFERENCES

1 Citrenbaum, R., J. R. Geissman, and R. Schultz. 1987. Selecting a shell. AI Expert 2:30.

2 Conlin, B. J., G. R. Steuernagel, J. K. Reneau, and R. R. Peters. 1988. Designing dairy management decision support systems using current management information principles and technology. J. Dairy Sci. 71(Suppl. 1): 200. (Abstr.)

3 Gilkes, A., and A. Nekhom. 1986. Applying expert systems to on-line fault isolation. TI Eng. J. 3:19.

4 Heald, C. W. 1988. Identification of Dairy Herd Improvement Core Parameters. J. Dairy Sci. 71(Suppl. 1): 145. (Abstr.)

5 Jones, J. W. 1985. Using expert systems in agricultural models. Agric. Eng. 66:21.

6 Jones, L. R. 1988. A natural language interface for retrieval of information from a dairy database. Ph.D. Thesis, Univ. of Illinois, Urbana.

7 Lazarus, W. F., W. A. Knoblauch, and T. R. Smith. 1986. Managing information on a dairy farm. New York State Food Life Sci. Q. 16:12.

8 Lemmon, H. 1986. COMAX: An expert system for cotton crop management. Science 233:29.

9 Levins, R. A., and M. A. Varner. 1987. An expert diagnostic aid for reproductive problems in dairy cattle. Comput. Electron. Agric. 2:47.

10 McMackin, L. 1987. Growing with expert systems. Penn State Agric., Penn. State Univ. 87-1031:24.

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

12 Morrison, J. B. 1985. The forthcoming role of artificial intelligence in the agricultural science and education system (executive summary). Purdue Univ., West Lafayette, IN.

13 Pederson, K. 1988. Connecting expert systems and conventional environments. AI Expert 3:26.

14 Smith, T. R. 1987. Remote Access to Dairy Herd Improvement Information in the United States. Proc. Natl. Invit. Workshop on Microcomputer Use in Dairy Manag. Dallas, TX.

15 Smith, T. R., D. Braund, R. Lamb, M. McGilliard, and R. Brown. 1987. Report of the American Dairy Science Association Subcommittee on standards for publishing with reference to computer software. J. Dairy Sci. 70: 209.

16 Smith, T. R., and N. R. Scott. 1986. The electronic revolution on the dairy farm. New York State Food Life Sci. Q. 16:9.

17 Spahr, S. L., L. R. Jones, and D. E. Dill. 1988. Expert systems–their use in dairy herd management. J. Dairy Sci. 71:879.