Review of the Utilization of Intelligent Systems for Plant Protection Project

Richard E. Plant
Departments of Biological and Agricultural Engineering and Plant Sciences
University of California
Davis, CA 95616 USA
replant@ucdavis.edu

June 16, 2006

Introduction

The Utilization of Expert Systems in Plant Protection (UISPP) project is an interagency effort to develop and implement expert systems for use in crop pest management in barley, wheat, and chickpea. The project's primary partners are ICARDA and CLAES, with secondary participation of IRRI and ICRISAT and tertiary participation of other agencies. The goal of the project is to address the need of NARS Extension services for up-to-date information on crop protection. I will assume that a reader of this report has also read the Project Proposal and other supporting documentation, and therefore I will not provide a detailed description of the project. Nevertheless, it is useful to summarize some its most important aspects, both for completeness and to provide a context for the conclusions and recommendations of the report.

The primary goal of the project is to develop and implement software for knowledge transfer on the subject of pest management of barley, wheat, and chickpea. A secondary goal, or perhaps more appropriately termed, an explicitly recognized side effect, is the development of the human and technical resources to efficiently create knowledge transfer software on other relevant subjects. Software is to be created in two forms, expert systems (that is, consultation programs that respond to direct, specific inquiries) and e-learning systems (that is, software that proactively transfers general knowledge about a particular subject). Both of these systems would be implemented using the classical expert system paradigm in which the domain independent knowledge and data management are separate from the domain specific knowledge. All knowledge management software for a particular domain would access the same knowledge base for that domain. The primary target users are Extension workers. Secondary target users are NARS scientists and possibly some farmers.

This report is based on observations made during a visit to ICARDA headquarters in Aleppo on May 22-23, 2006 and to CLAES headquarters on May 24-25, 2006, as well as review of supporting documents provided by the project leaders. The report is organized as follows. I will first provide a brief list of the terminology as I use it. Following that, in the interest of avoiding any suspense, I will immediately "get to the point" and give my conclusions and recommendations regarding the overall project. I then list some observations that I noted while on my visit, in the hope that this view by an impartial

outside observer will be useful to the project participants. Following that I provide some constructive criticisms of what seemed to me to be the weak points of the system. Last, I give some more detailed recommendations on the way I feel the second phase of the project would best be carried out.

Terminology

In order to ensure that there are no misunderstandings resulting from technical terms, whose usage is often inconsistent, the terms used in this report are defined here.

Data, information, and knowledge. As I use these terms in this report, these terms have the following interpretation. Data are recorded observations pertaining to a process or phenomenon. Information is data that have been placed into a context. Knowledge is a generalized interpretation of the behavior of the process or phenomenon obtained by integrating multiple instances of information. It should be noted that in a rule based expert system the knowledge base is generally encoded as the propositions and conclusions of rules, possibly including certainty of the conclusion, and that this encoding may be stored in a database.

Domain. The subject of the expert system. In this case the domain is pest management of barley, wheat, or chickpeas.

Inference engine. The software responsible for acting on the knowledge base and developing the conclusions to be displayed to the user.

Recommendation

I strongly support continuation of this project into Phase 2. I have over the years developed a heavy skepticism about large scale expert system projects such as this one. Only a very small fraction of these projects even make it to the implementation stage, and of those that are implemented the vast majority (including my own CALEX/Cotton system) do not last more than a few years. There are a number of causes for this failure, and indeed an analysis of these causes was the subject of a special issue of the journal *Agricultural Systems* a few years ago. Some important reasons for failure are the following:

- *Failure of the supporting agency to commit long term funding.* In many cases the development of the expert system is supported by research funds. The maintenance of an existing system, however, is not research but extension, and in many cases there has been no provision for continuing extension support. This was, for example, the fate that befell the CALEX/Cotton system.
- *Lack of commitment on the part of the developers in the ultimate success of the system.* To some extent this was also a problem with CALEX/Cotton. Often the main goal of the primary developers is recognition of the research component of the development effort through the publication of journal articles. The developers may be unwilling or unable to make the sacrifice to their professional careers necessary to shepherd the system through the inevitable problems associated with maintenance of software (operating system updates, customer support, etc.).

- *Failure to transition from research to Extension personnel.* Even if support funds have been provided, there may be no qualified extension staff to take over maintenance of the system once the researchers have left the project. This was the fate that befell the Gossym-COMAX project, for example.
- *Failure to properly assess the needs of the intended clientele.* The developers of the expert system may have inadvertently created a system that solves problems not currently of sufficient interest to the intended clientele to warrant the effort to acquire and use the system.
- *Poor user interface*. The system may be too complex, clunky, or uninformative. This is particularly characteristic of systems whose user interface is designed by computer scientists rather than professional interface designers.
- *Failure to commit sufficient resources to marketing the system*. Even if the product is a good one it may be defeated by a lack of interest or ability on the part of the developers or the sponsor to market the system to the intended clientele.

There are very few projects that do not suffer, in most cases fatally, from at least one of these flaws. I came into the UISPP project predisposed to encounter one or more of them in sufficient measure to doom the system. I was presently surprised to observe that this seems not to be the case. Although the project is not perfect, and there are some worrisome aspects, I nevertheless feel that the positive features of the project sufficiently outweigh the negative features that there is a very good chance of short term success and, if long term support is committed, a reasonable chance of long term success. For this reason, if funds are available then I support their being expended to support Phase 2 of the project.

The Phase 2 proposal should include a plan and realistic estimated budget for long term implementation after the initial development process is complete. If the project is funded for Phase 2, the funding agency should determine whether there exists the capacity to support maintenance of the system in the long term. I believe that benefit can be derived from the project even if it cannot be supported in the long term (or if it does not survive for one or another reason such as those given above), and therefore I recommend support for Phase 2 even if these funds cannot be guaranteed. Whether or not long term funds are available will, however affect the research and development priorities as discussed below.

General Observations

I was very much impressed with the enthusiastic support displayed at all levels by the domain participants in the project. The scientists with whom I interacted, both the ICARDA scientists in Aleppo and the NARS and university scientists in Cairo, were uniformly positive in their discussion of the system. Even more impressively, the Extension workers with whom I met in Cairo were uniformly positive in their feelings toward the system. I cannot emphasize enough what a positive sign this is. At one, obvious level, an expert system will not be successful unless the domain users buy into it. There are, however other, more subtle considerations that may be even more significant. Novel software such as an expert system can be perceived as intimidating or even threatening by non-specialists, and they can respond by sending subtle negative or skeptical signals about it. I did not detect any such signals, indeed, the signals, both overt

and subtle, were nothing but positive. This indicates to me that the individuals responsible for developing the system have done an excellent job of engaging the domain scientists and potential users, and that the potential for successful adoption of the system is relatively high.

The primary intended use of the UISPP expert systems is to provide a means of dynamically transmitting knowledge to Extension workers. This is clearly stated in the proposal, and was repeatedly expressed in meetings with the project leaders and participants. There are, however, a number of secondary uses that can add value to the project at very little added cost. The project leaders are obviously aware of these, and indeed mentioned them on numerous occasions. I did not detect, however, any formal provisions in the project devoted explicitly to exploiting these secondary uses. If these provisions are not explicitly made early in the project there is a risk that the opportunity to exploit these secondary uses to the fullest will be lost. The following is a list of some of the secondary uses I have in mind, and how provision might be made to take advantage of them.

- The knowledge base provides an archive of the knowledge of many scientists that will be available after these scientists have retired or left their current employment. Since the primary knowledge base developed at ICARDA is maintained as a Microsoft Access file, this archive can be accessed by setting up a Microsoft Access front end query system.
- The development of the knowledge base provides a means of integrating knowledge from disparate disciplines. The project has initiated a procedure of regular interdisciplinary meetings of scientists in which the contents of the knowledge base are discussed and modified. This procedure is excellent and should be maintained, with support provided for participants' travel. For this reason the geographic extent of the implemented expert system should be consistent with the funds available to support these meetings.
- The formal encoding of knowledge provides a means of identifying knowledge gaps. Many of these gaps will become evident in the process of knowledge acquisition. However, one of the tasks of the CLAES verification group should be to develop algorithms that identify logical gaps in a general knowledge base.
- The knowledge base provides a means of organizing knowledge. A part of the development of the ontology in the second phase should be a hierarchical structure for organizing the knowledge base. The developer of the ontology should interact with participating scientists to develop an understanding of how these scientists mentally organize their own hierarchy of thought in conceptualizing their understanding of the biology of the system.

There have over the last 40 years been developed a large number of crop simulation models. Most of these models are developed for general circulation with the idea that individuals outside the development effort will use the models for some particular purpose. Nevertheless, the general feeling about crop simulation models that the primary beneficiaries of the development of a crop simulation model are the developers themselves. The process of model development forces the developers to think critically about crop growth and development and to try to fit together all of the processes

involved. The developers in this manner identify knowledge gaps and develop a better organizational structure for the knowledge that they do have. In a similar way I suspect that thus far for this project the primary beneficiaries of the development of the knowledge base have been the scientists who collaborated to develop it. This is not a bad thing, and indeed the enforced discipline of a group of individuals meeting to link together bits of knowledge into a coherent whole can be a very rewarding experience. The domain scientists with whom I met in Aleppo and Cairo were very enthusiastic about the project and I suspect that this was in part because they recognized the benefit of assembling the knowledge base.

Another point worth noting is that the developers have been very wise in building flexibility into their database. Their use of the Concept, Property, Value Indicator model for knowledge representation seems to work very effectively for this use. The ability to incorporate new information readily into the knowledge base is a very important one, and I would expect that the system is well positioned to do this.

Problems

As I mentioned in the Introduction, I did detect some features of the project that seem to me to warrant some constructive criticism or suggestions.

1. There appears to be at this time a partial disconnect between the domain scientists and expert system developers. This is to some extent inevitable due to the geographic separation between the principal domain scientists, located in Aleppo, and the expert system developers, located in Cairo. The effects of this separation manifest themselves in the potential to create expert system software that, while optimal from the point of view of artificial intelligence theory, are less effective in their application to the specific domain. The very rapid and effective development of the knowledge base through the action of seconding a CLAES scientist to ICARDA during the knowledge acquisition phase suggests a way to ameliorate the effects of the geographic separation of the two systems. This is to place a young pest management scientist, such as a postdoc, in CLAES during the period of development of the expert system. This person should be assigned to the primary task of guiding the development of the interrogation and explanation systems, but also including any communication between the user and the inference engine, and in the development of the ontology. The secondary task would be to serve as a readily available domain consultant for the developers. In addition, or as an alternative, I suggest retaining a person with a strong computer science background and placing this person at the ICARDA headquarters in Aleppo. The primary responsibility of this person would be to serve as a conduit of information between the computer scientists at CLAES and the domain scientists at ICARDA.

UISPP Team Response:
This is true in traditional development of ES but we claim that we have developed a generic tool for agricultural tasks and what is needed is to build the knowledge base of this tool in a well-defined format. We came up with this tool after so many years of interaction with domain experts. The solution we follow is untraditional and what is needed now is to evaluate the developed ES with domain experts to get their reactions.

Modifications will be implemented in the tool such that new experts systems can be developed better, faster and with minimum interaction between the developer and the domain experts.

2. There appears to be a wasteful redundancy of effort. There is an existing web-based knowledge acquisition tool developed by CLAES staff seconded to ICARDA during the first phase of the project, and yet there are still CLAES staff committed to developing duplicative knowledge acquisition tools. It is not clear to me why this duplicative effort is perceived to be necessary, but it should be discontinued unless there is very strong justification.

UISPP Team Response:
The idea was to let two groups develop the same tool in order to have different views and ideas for developing the KA tool and then to choose the better one or to take the best idea or component from each one and have a better one than any of them. This practice is done when it is required to create a new piece of software whose specifications are not well defined before hand. In effect, the CLAES tool includes some functionality that does not exist in ICARDA tool and visa versa.

3. One of the most common tendencies among software developers is to react defensively when their product is criticized. I suffer from this tendency myself. It manifests itself in a desire, when the system is criticized, to defend the system and debate the critics, or to try to explain why the system works as it does, instead of accepting the criticism and trying to modify the system to address the perceived problems. This is a universal problem and the CLAES personnel are no exception, but in interacting with domain scientists and users they need to be aware of this tendency so that they can resist it.

Specific Suggestions

1. The project leaders verbally expressed the intent that this system have worldwide implementation, and the Phase 1 prototype includes a query to the user requesting the user's country and providing a selection list that includes dozens of countries from Argentina to Zimbabwe (more or less). While this is a laudable ambition, it is misplaced this early in the project and threatens to dilute the effort and thereby jeopardize success. It is far better to do a very good job in a couple of countries initially than a mediocre job in many countries. I recommended that the initial implementation be in only two or at most three countries. One of these should be Egypt that the other should be a country where there is either sufficient ICARDA presence or a sufficiently capable and committed NARS to provide good support for the initial implementation. By the way, a good example of annoying redundancy is asking the user to specify his or her country every time the system is used.

UISPP Team Response:
There is a misunderstanding here. It is not intended to develop an expert system, e.g. for Barley worldwide. We included data for some countries as our scientists added varieties of barley for these countries. What we want ultimately to avail on the Web, is a Generic Agriculture ES shell that can be used easily by scientists in any country to develop their

ES's for Plant Protection if they want. The Barley ES developed was a demonstration of this approach. This will be an international public good.

2. The project should consider using Extension workers directly in the expert system development. As I have already mentioned, I was especially impressed by the commitment to the system of the Extension workers I met with, and their positive attitude toward the system. The project should exploit this by including these individuals directly in the development effort. They are particularly useful in the development of the user interface. A meeting could be arranged in which interface developers and Extensionists interact to work on this interface so that it best suits the Extensionists needs and wants. This will provide the double benefit of making use of the Extensionists' knowledge, and of giving these individuals a sense of ownership of the system and commitment to its success.

UISPP Team Response:
Accepted.

3. There is a critical need for version control and verification of the knowledge base. The Knowledge Acquisition Tool developed last year in ICARDA makes a very good first step in this direction by providing an automated means of notifying the primary domain expert whenever someone edits the knowledge base. Moreover, the knowledge base as it currently exists includes the source of the knowledge for each record. This provides a good foundation for maintaining traceability in the knowledge base. The existing knowledge base developed in 2005-06 should be considered as the base, and the versioning system should be such that an audit trail can be generated for any rules such that any modification made to the rule can be traced back to this original version. This audit trail should include the date and the author of each change. Moreover, when a change in the knowledge base is proposed there must be an automatic notification of an authorized domain expert so that this expert can review the proposed change. Finally, it would be useful to include in the data file that represents the rules a field in which a citation of a source document may be given.

UISPP Team Response:
Accepted, and is a good idea.

4. The developers have made a good start at a system that avoids the problem of "brittleness" (that is, the tendency of expert systems to be unable to find any solution in the absence of complete knowledge). They have done this by subdividing conclusions into "suspected" and "confirmed." The way in which the system interacts with the user to arrive at these conclusions, and the way in which the conclusions are displayed, is, however, very confusing. The developers should work with domain experts to determine the manner in which these individuals approach a diagnostic problem, and should develop a user interface that is as consistent with this way of thinking as possible.

UISPP Team Response:
Accepted and this is what we intend to do.

5. I don't recall if this was discussed in my visit, but I would suggest that the developers include an electronic messaging system within the web interface when the system is implemented. One of the most important by-products of a project such as this is increased communication among all participants, and the application software should do everything possible to encourage and facilitate this communication.

UISPP Team Response:
Good idea and accepted.

6. The project has a number of second phase objectives, and the principal investigators need to set priorities for these. My own recommendation is the following prioritization.

First priority. Finalize the inference engine for the expert system. The creation of the pest management knowledge bases is a significant achievement, and these knowledge bases are a significant resource. The inference engine will permit this knowledge base to be accessed and thereby give it value. This task should be essentially completed already.

Second priority. Develop a good user interface for the expert system. The knowledge base will be accessible once the inference engine is completed, but only to those individuals willing to make the commitment to deal with a relatively clunky user interface. To make the system more widely accessible, and to involve more people in the project, domain scientists and Extensionists should be given a direct role in the development of the user interface, including the explanation system. This should be done in the context of meetings similar to the meetings that occurred in the first year when the knowledge base was developed.

Third priority. The third priority is to finish the development of the online knowledge management system. This should provide the capability to track and store edits of the knowledge base and to verify the knowledge base when needed. This will be necessary when NARS scientists begin to participate in the development of local versions of the knowledge base.

These first three priorities are essential to the successful completion of the project. The remaining priorities are less so.

Fourth priority. The development of an ontology for the knowledge base will be useful. It will, if done correctly, provide a means of organizing the knowledge on a hierarchical manner somewhat akin to an object – oriented database.

Fifth priority. The e-learning module is independent of the expert system and would indeed be a good second use of the knowledge base. However, expending CLAES resources on its development would leave them unavailable for use on the expert system project. If IRRI can be persuaded to develop an e-learning package using the knowledge base then this would be appropriate, but if they are not willing to do so then perhaps this module should be set aside for now.

Sixth priority. The development of an intelligent search system is not really necessary for the project and can easily be abandoned if necessary.

UISPP Team Response:
The overall conclusion and recommendations of the report are generally accepted and will be taken care of in the second phase of the project. The team notes particularly that the report has highlighted the knowledge management aspect of the project and the important achievements made, which we think are closely aligned with the overall goals of the ICT-KM program, which had the vision to support the first phase of this project

Acknowledgements

I want to express my thanks to Zaid Abdul-Hadi, Ahmed Rafea, and their co-workers in Aleppo and Cairo for their hospitality and patience with my innumerable questions.