21

Towards an archaeological methodology for expert systems

Katharine G. Baker
Research Centre for Computer Archaeology, North Staffordshire Polytechnic

21.1 Introduction

Expert systems are a borrowed technique. Archaeology has adapted and adopted techniques from many disciplines; in this case the borrowing is from computing although the broad area of artificial intelligence also has strong links with areas such as cognitive psychology. With any wholesale adoption one problem is that, as archaeologists, we may lack a detailed understanding of the borrowed technique and so we may attempt to exploit it without really appreciating its assumptions, effects, or implications. In the case of expert systems the problem is even worse because the successful application of the technology is relatively new even in its host discipline.

Without thought and management we may lose out on this technology as misuse causes disillusionment. Use of a borrowed technique may follow a path on which a phase of over-enthusiastic application, or often mis-application, is followed by a period of disillusionment and possible neglect, together with the loss of any benefits the new techniques had to offer. Archaeologists have not yet had enough to do with expert systems for the disillusionment stage to set in. If we can link an awareness of the problems together with the careful management of scarce archaeological computing resources, it might be enough to avoid the negative parts of this path.

21.2 Problems to be addressed

Not only are expert systems the theoretical product of another discipline; the practical application of expert systems is likely to be in business and commerce, for business and commercial ends. This means that the practical problems of applying these systems are going to be ironed out in that same atmosphere. This gives us the advantage of access to a technology which has been developed to maximise efficiency and return on investments—the disadvantage, of course, being that our aims are not those of a necessarily commercial concern.

At the moment our expectations of these systems are confused and we need to think hard about what we need as archaeologists, not just to jump onto a fashionable bandwagon. Archaeologists have been using computers for many years particularly for database and number-crunching work and expert systems should be seen as another tool to help in the handling of our data, rather than as a magic wand or the arrival of the substitute human. The original specification of one of the current projects at North Staffordshire Polytechnic (Huggett 1985) suggested that expert system techniques might be useful in the analysis of Anglo-Saxon inhumation cemeteries. This has

proved to be both less feasible and less desirable than was thought at that time. Interestingly, although it is carried out by 'experts', this is an area of archaeological activity for which there is no accepted methodology. It has been observed (Lagrange & Renaud 1985) that 'we are not yet quite sure that an archaeological 'artificial expert' would prove useful for archaeological research, properly speaking', and this case does appear to support that view. It is important to differentiate between the use of an expert system as another kind of tool to aid research and treating it as some kind of magic box which will produce answers all on its own!

The names and language associated with these systems in both the commercial and research fields imply that the systems associated with them are much, much more than the ordinary computer systems which we are already using. The words used include 'expert', 'intelligent', 'thinking', and 'knowledge'. Perhaps we should get rid of the commercial names as far as possible and redefine the systems in our own terms as tools, pure and simple. Expectations that these systems will solve all our archaeological problems cannot be fulfilled, especially as we probably could not define what those problems are. Unfortunately the more that seems to be promised, the greater the disappointment will be and the more likely that we will get onto the negative path of disillusionment and neglect. End-user expectations of expert systems are also increased by work on natural language interfaces, opaque systems, and user friendliness and tolerance. A system with a good front-end may well appear to be intelligent, and a greater reliance placed on its judgement than is warranted. It is very easy to produce a visually impressive package which is not solidly founded in the domain subject. It is very important that we note the differences between the theoretical concepts appropriate to expert systems, current commercial and research practice, and the present state of possibilities. Negative views and subsequent problems often seem to arise from an attempt to work with the image of the sort of super-genius machine which turns up in science-fiction.

In theory archaeologists could commission individually-designed systems for specific applications but there is still the problem of differences in views and expectations between the end user and the designer, especially if one of the partners is not an archaeologist. Careful systems analysis can minimise but probably not eliminate the problems which may lead to the non-acceptance or misuse of systems. In practice archaeologists probably cannot afford machines, different machines, packages, shells, and languages and almost certainly we cannot afford to commission tailor-made systems. Bespoke computing tends to be as expensive as bespoke anything else. If we are to use these systems we must therefore try to develop a relationship with the industry which supplies them—which must be symbiotic as they are not usually in the charity business. We must also move towards the use of bought tools which have as wide an archaeological application as possible.

If the possibility of links with suppliers does occur we should not undersell ourselves. When we are trying to get commercial help we could exploit both our novelty value and the fact that the rôle of the domain expert is a subject attracting a great deal of attention at the moment. Lagrange and Renaud note that their work on iconographic problems with the SNARK system has proved interesting to expert system specialists as well as to archaeologists (Lagrange & Renaud 1985). It should be emphasised that archaeology can prove useful to people to other disciplines as well as vice-versa. Piecemeal development, with several versions of the same problems may also mean that the technology is not used to our best advantage. Expert systems offer several benefits one of which is the opportunity to standardise interpretations of agreed systems—samian identification for example. If there are several versions of a system to perform this identification then the situation is as it is at present. If, however, one system could be built, verified, and agreed upon then an advance has been made.

It is important that we clarify our expectations, but we are hampered by the very nature of artificial intelligence and hence expert systems. It is very hard to define them: how can we have realistic expectations when they are all things to all people? Pure artificial intelligence/expert systems work is often difficult to pin down as it tends to file itself under another title which is easier to deal with, such as natural language interface, planning systems, or data storage. It is a well recognised phenomenon that once artificial intelligence/expert systems are applied they tends to stop being artificial intelligence/expert systems and to start being, or rather to start being perceived as, part of the theory or techniques of the subject domain. Hart 1982, quoting Nilsson 1974: 'AI is a no-win field. It exports all its winning ideas.' This does tend to support the view that it is really, however glamorous and exciting its presentation, just a collection of techniques, rather than a practical domain in its own right.

In the past expert system researchers have descended onto a domain, such as medicine, produced a system which 'works' in some way and have then left, written up their project and received lots of praise and admiration, probably mostly from other computer people. Recent figures on the actual take-up and use of expert systems are worrying: in the real world very few of the large impressive systems are in use. More are being adopted now but, underneath the trimmings, they are small systems working on a limited problem area within a restricted domain. As archaeologists we must look at the reasons for both these facts and adapt our strategies accordingly—one of the advantages of adopting the techniques and technologies of other disciplines is that we can learn by their mistakes.

The commercial and research expert systems literature suggests reasons for the non-adoption of systems (for example Morris 1986), to which we can add some archaeological ones. The interaction of expert system researchers and the domain experts may not always have been totally satisfactory on either side, however immediately appealing the results. The needs of many of the domains used for research have not been analysed fully, and the niche into which the expert system is to fit, or the frontier through which it is going to burst have not been defined. There is, however, a move away from the old situation and towards application-driven expert system research, with the views, needs, and expectations of the eventual end-users being taken into account. There is now, for example, an annual computing conference which is structured from the domain point of view. This trend is something which we should take heart from and, if possible, encourage.

21.3 Computing problems accepting expert systems

Threat or aid

Human experts are rightly concerned that the systems may be seen as means of replacing them rather than assisting them. While you do not always need a human expert to build a system unless its performance satisfies the humans working in that field it will probably not be accepted. If you are working in an area where you need a human expert you will not succeed without their support. Some areas in archaeology are greatly understaffed, but there are political and moral decisions to be made here.

Responsibility

One of the major problems is the responsibility for building, maintaining, and distributing a system. If it makes an error, whose fault is it? After all, human experts sometimes make mistakes so we must expect the same from a machine which tries to mimic their performance.

One response in the business world has been the foundation of groups of people within companies who are jointly responsible for updating and maintaining the system. This kind of organisation might be harder to achieve in archaeology but, if expert systems are to be widely used it is a problem we must tackle.

Integration

The outside world is worried that it cannot fit expert systems into its present computing set ups. I do not think that archaeology is that computerised yet—but interfacing to existing facilities and the human elements of using the systems should be addressed in any design and analysis for an expert system.

Resources

There are very few people with good experience in constructing these systems. In archaeology we suffer from the usual financial limitations as well. We really need a strategy for computerisation rather than a piecemeal approach.

Immaturity

If the outside world is worried that the technology is immature, we, without the experience and the resources, should be doubly careful. Do we really want to be at the sharp edge of other disciplines' research projects?

Costs v. benefits

While our aims are not purely commercial we too have to work out exactly what benefits we are going to get for our outlay. In the area of expert systems this is made more difficult by the intangible nature of some of the possible benefits. How do you put a value on an expert's expertise? What if that person leaves and there is no time to train a successor? What extra could the expert achieve if the machine helped with the time consuming aspects of the task?

21.4 Archaeological problems accepting expert systems

Academic subject

Archaeology is a subject with an academic base. This means that in many cases the reasons for an answer will be as important to the archaeologists as the answer itself. Tackling archaeological problems does not have the possibilities for disaster which are present in, say, medical or defence applications but, at least to the archaeologists, the results and the processes by which they are achieved are important. Any system which could not provide an adequate explanation of its answers would be unlikely to be accepted in many archaeological situations. Of course, systems which provide 'right' answers for the 'wrong' reasons might be extremely useful, or at least thought-provoking.

Snobbery

There is still the attitude that archaeology is uncomputerisable, that archaeological data sets and problems are somehow beyond reduction to a form which can be handled by a computer. There is still the sneaking feeling that it is prestigious to have uncomputerisable data. This attitude may well be even more pronounced in the case of areas of expertise than it is in areas

of straightforward data (Richards & Ryan 1985). It should not be assumed that this attitude is incorrect. If the change in form is indeed a reduction then we should be wary of how we use it.

Fossilisation

If we select one version of something to be encapsulated within an expert system, there is the possibility that we are suppressing the rise of other versions, different but equally valid. Stagnation could well be a result—and the end of archaeology as a lively discipline the next step.

Suitability

How much archaeological expertise is suited to expert systems? If the benefit of, say, a more standard identification is to be gained then the subjects should be found on a national rather than a local level. Much of what we see as archaeological expertise may well be too fuzzy and unstructured for the present state of the technology. We should not stop using the technology for those things it can do, however.

21.5 Why should we bother?

Even in an ideal world with ample resources for archaeological projects, we would still face many of the problems noted above. Those defined as being particular to archaeology are the most urgent if these systems are ever to be more than very basic administrative assistants. There are various benefits which archaeologists might gain from the use of expert systems. The following is not a complete list and, at the present state of play, benefits are conjectural rather than proven. It should indicate, however, that there are reasons for trying to deal with the problems mentioned above.

- They can speed up processes, perhaps cutting down the time taken to produce a report and encouraging rapid feedback to site.

- Using an expert system is an improvement on the present system where a novice in the particular field uses unhelpful books, assuming that such books exist. With a book there is no external check on what is done. What if a vital footnote has been missed somewhere? Not only should a machine be consistent in its application of the rules, but you the human can check what it is doing. There is also the possibility that expert systems can provide a teaching aid but, unless they have been designed with this as an aim they tend not to be terribly effective. There is also the question of how far down the line learning from a copy is acceptable—at what point do we need to talk things out with a human being?

- These systems should free the human experts to perform 'real' work: less mechanical identification and more analysis and synthesis. Archaeology is becoming more and more compartmented: each expert has a greater understanding of a narrower field and synthesis becomes more difficult. In areas with a strong practical, classificatory element it is easy to envisage systems which will take on this area of 'expertise'. A properly supervised expert system can free a human to take on analysis and synthesis. In practice, however there is the danger of becoming distanced from the data.

- While there is a real danger of fossilisation, the deification of certain ways of doing things together with the theory which underlies them, there is also the opportunity to reassess some of our practices from first principles. Of course, we should not then fossilise this new version!
- Expertise often remains unrecorded, even in the form of a book. At the point when the expert dies, resigns, or changes jobs the expertise is no longer available. It may be possible to encapsulate some aspects of expertise against disaster by creating an expert system.

21.6 Suggested actions

I suspect that the chances of radically reorganising archaeology to exploit these systems in an ideal way are remote—and really why should we do so? They are not, after all, the philosophers' stone. They are simply a new technology which seems to have archaeological potential in terms of developmental benefits if not yet in terms of output utility. Their use can be seen as part of the continuing shift from qualitative to quantitative, from humanity to soft science. (On a practical level I also hesitate to suggest we should have yet more committees.) I think that the answer, if there is one, is based on the 'mighty oaks from little acorns' principle. If we act in concert at a low level then we can at least limit any negative results until we can see what is happening. It is perhaps the attitude and methods of individuals building individual systems which should concern us most.

The archaeological use of expert systems has similarities with the earlier introduction of techniques such as computer simulation and systems theory. In these systems a model of the situation must be constructed. Attempting to build a model is likely to encourage a basic rethinking and hopefully to lead to further insights. These are major benefits, especially in areas where existing methods seem to have stopped producing new results. The side benefits of building models are also emphasised by Lagrange & Renaud 1985: each creative attempt is likely to lead to a progressively better model. Aldendefer 1981 states that 'simulation models based on fuzzy laws cannot (and should not) be used to gain output utility.' In terms of expert systems this is not so. At the very least we can gain Aldenderfer's conceptual and developmental benefits while advances in knowledge elicitation, knowledge representation, and inference have meant that we can, and should, aspire to 'gain output utility' 'based on fuzzy laws'. The danger is that we should aspire to more than is reasonable and, in the effort to demonstrate our technological wizardry, we might ignore the pitfalls: we must accept that 'experimental' means just that. We face a 'Catch 22' in that if we do not have practical and intellectual problems implementing one of these systems then we have probably implemented something else! Or we have worryingly and totally misunderstood the implications.

We must be able to assure the human experts that these things are aids and tools and not rivals. Practically, in most systems you need a human expert if you are to do more than produce a simple type of classification system. If all you do is take your data and rules from a book and perform a text animation process, then you have innovated very little and, I would suggest, the resultant system is at most a very simple form of 'expert'. It is that part of expertise, however defined, which cannot be, and has not been, successfully committed to paper which should concern us. These systems are most interesting when offering the possibility of doing or understanding some thing new, rather than doing better something we already do. It is my personal view that to use these systems as replacements for humans in order to cut costs is immoral; there can be little doubt that it is also very short sighted. At the present state of

technology and legal development we need the involvement of humans. We must keep sight of the limitations of the technology: at the moment a machine cannot replace a human expert in any real sense. They cannot, for example, grow and develop without substantial help from human beings.

Just as we have found that many commercial products in other areas of computing can be used for an archaeological purpose, I suspect that we will be able to exploit expert systems technology largely as it stands. The difficulty lies in accepting and interpreting its limitations and possibilities within an archaeological framework. We can imagine a situation in which all the problems mentioned above have been solved, or at least come to terms with. Or we can accept that expert systems technology should be used at a low level of intensity—that the computing problems we tackle should be at a trivial level, the implications as well understood as possible. Perhaps, in our own interests, we should stick to applying those parts of the technology which are relatively well understood and developed, for example classifications. To take advantage of the new and exciting parts of artificial intelligence/expert systems research archaeologists must accept that experimental projects must be treated as such, and that we must not let the computer sophistication of these systems carry the archaeological theory beyond the point at which it is self-supporting. It is dangerous to try to develop archaeological theory when we do not fully understand the implications of the computing involved.

Given that we cannot fund massive research purely into expert systems on our own account we are reliant on a few researchers and field archaeologists working on their own projects in semi-isolation. The nature of expert systems suggests, however, that the more cohesive our efforts the better. The acceptability of expert systems depends to a large degree on their 'safety', however we define it. As each domain has its own limits and tolerances on what can be considered 'safe' it is up to us as archaeologists to try to formulate our own standards. Just as in testing, where a system can be said to succeed if it meets its design specification, we need external standards to help us in judging these new systems. We must work on developing standards for assessment, testing, validation and acceptance before the system is built and farmed out. We may well produce a perfectly valid and useful small scale not-really-expert systems which we unfortunately fail to appreciate because we do not know what we want.

Many of the most important problems stem from the fact that we are an academic discipline. People tend not to stay in archaeology unless they actually care about the results, albeit to a greater or lesser degree. For these reasons we cannot just hand over our expertise to a pack of machines as we might hand over our business administration. We are not in a position simply to accept the hard copy which the machines produce. What happens to us if they are wrong? One of the most striking things about expert systems is that they need areas to work in. What is more, the relationship between an expert system and the domain in which it performs is more complicated than that between, say, a database package and the data held within it. Given the notorious difficulty of using database packages, whether off-the-shelf or custom-built, with our archaeological data, the mutual impact of these expert systems and archaeological data should be pondered very deeply.

References

ALDENDEFER, M. S. 1981. 'Computer simulation for archaeology: an introductory essay', *in* J. A. Sabloff, (ed.), *Simulations in Archaeology,* University of New Mexico Press, Albuquerque.

HART, P. E. 1982. 'Directions for AI in the eighties', *Sigart Newsletter,* 79, pp. 11–16.

HUGGETT, J. W. 1985. Research proposal for CNAA. Unpublished.

LAGRANGE, M. & M. RENAUD 1985. 'Intelligent knowledge-based systems in archaeology: a computerised simulation of reasoning by means of an expert system', *Computers and the Humanities*, 19, pp. 37–52.

MORRIS, A. J. 1986. 'Why people aren't using expert systems', *in Second International Expert Systems Conference, London, 30th September–2nd October 1986*, Learned Information, Oxford.

RICHARDS, J. D. & N. S. RYAN 1985. *Data Processing in Archaeology*, Cambridge University Press, Cambridge.