Current Issues in Archaeological Computing

edited by

M. A. Cooper and
J. D. Richards

BAR International Series 271
1985

Wilcock,J.D. 1974. "The facilities of the PLUTARCH system" Science and Archaeology, 11,16-24.

- 1981. "Information Retrieval Applications for Archaeology in Britain" in Gaines,S.W. (ed.), 1981, 100-122.

Voorrips,A. 1984. "Data catchment analysis and computer carrying capacity", in Martlew,R. (ed.), 1984, 45-53.

- 1985. (ed.) To Pattern the Past: Mathematical Methods in Archaeology, volume 11. P.A.C.T., Amsterdam.

EXPERT SYSTEMS IN ARCHAEOLOGY

Jeremy Huggett

INTRODUCTION

The aim of this paper is to examine selected aspects of artificial intelligence and to describe some of the problems and the potential that the use of such techniques may offer archaeologists. Discussion will be restricted primarily to the subject of 'problem-solving', since this facility is becoming increasingly available in the form of commercial expert system shells and is therefore of more immediate interest than those areas which are still strictly experimental. A number of archaeological expert systems have recently appeared, but there has been little discussion concerning their application in the discipline.

A standard dictionary definition of intelligence is "an ability to reason and understand" (Fink 1968, 209), but to suggest that machines can be capable of such a thing tends to raise a storm of protest and philosophical discussion since there is an inherent implication in this that people are themselves no more than machines. However, artificial intelligence may be defined as "the science of making machines do things which would require intelligence if done by man" (Minsky 1968, v). Minsky's definition of machine intelligence requires only that the machine should simulate intelligent action, not be intelligent itself. This behavioural definition seems more appropriate since, when the program of an 'intelligent' machine is examined, it is clear that the intelligence is only apparent, achieved by means of clever programming devices.

McCarthy and Hayes (1969, 466) see intelligence as consisting of two parts: the epistemological part, a representation of information or knowledge; and the heuristic part, a mechanism that solves problems within the represented domain on the basis of that knowledge. Boden (1977, 17) defines artificial intelligence as:

> the use of programs as tools in the study of intelligent processes, tools that help in the discovery of the thinking procedures and epistemological structures employed by intelligent creatures.

Thus artificial intelligence requires both a knowledge (or representation) of a particular subject or domain and the ability to reason within that model world. Decision-making within this domain is seen as being fundamental to intelligence and crucial to problem-solving activities (Millar 1971).

The simulation of these "intelligent processes" requires that computers should be able to learn from experience, to organise information in such a way that it can be manipulated in solving problems, and presumably to perform at something approaching a human level of ability (see Fink 1968, 209,225). In other words, a machine displaying artificial intelligence should be capable of solving problems by deploying relevant knowledge from its knowledge base, or if the knowledge is not there it should be able to recognise this and 'learn' by adding to its knowledge base. Whether this constitutes 'intelligence' remains open to question. A dedicated machine with domain-specific knowledge may appear to be intelligent within its own sphere, but outside it, it is comparable to an 'idiot savant', "capable of performing mental miracles within a single narrow category ... yet ... otherwise of subnormal intelligence" (Michie 1982, 135). Any discussion of machine intelligence is clearly fraught with difficulty and perhaps the safest attitude is that:

> if the special task which a program performs is a very difficult one by human standards .... we may call the program 'powerful', but we still do not call it intelligent (Michie 1982, 135).

It may therefore be useful if emotive terms such as 'clever', 'intelligent', or 'understanding' which may creep into discussion of such machines and their programs are qualified by the use of inverted commas or "scare-quotes" (Boden 1977, 17).

Classic examples of early machine problem-solvers are those programs applied to problems in chess (see Michie 1982) and draughts (see Samuel 1963), for instance. These employed 'heuristic programming' techniques in order to assist with solving the problem, rather than use a 'brute force' method. The 'brute force' approach gives "a bigger bang for the buck" (Michie 1982, 180) in that it is simpler (and therefore cheaper) than the heuristic approach; however, in attempting to solve the problem it will evaluate all the possibilities open to it. Such an approach is clearly impractical: in chess, for example, there are estimated to be 10**120 possible board positions, and 10**40 in draughts (Boden 1977, 512). Theoretically, such a search algorithm will <u>eventually</u> produce the answer, although it may take a considerable time to do so. Without the use of heuristic programming techniques, a system would pursue approaches long after a human problem-solver would have realised that a solution was not in sight. This brute force technique is otherwise known as the 'British Museum Algorithm' (Newell, Shaw and Simon 1957; Raphael 1976, 59), so called because it seems as sensible as providing monkeys with typewriters in order to reproduce all the books in the British Museum.

A human problem-solver does not attempt to carry out an exhaustive search of all available options but limits evaluation of the problem by intuition, experience or 'common-sense' to the area where a solution seems most

likely. Such human heuristics tend to be 'rules-of-thumb', intuitive guesses, or, as Boden says:

> human reasoning - even of the most 'explicit' or 'rigorous' type, as in science and mathematics - employs integrative principles of tacit inference or global knowledge of which one is not introspectively aware, but which usually determine the nature of the thought contents of which one is focally aware (1977, 435; quoting M. Polanyi).

Heuristic programming attempts to emulate this kind of behaviour by greatly reducing the area searched, but as a consequence risking the chance of overlooking the solution altogether or finding one that is less than satisfactory. Putting these thought-processes into words is a notoriously difficult task since, as Sloman (1971, 272) says:

> Many persons can recognise and use valid inferences even though they have learnt no logic and become incoherent when asked to explain why one thing follows from another...

The inability of a human expert to express clearly, if at all, this kind of tacit inference poses a major problem in attempting to understand the means by which a human problem-solver arrives at a solution.

In terms of artificial intelligence, a heuristic device is simply:

> a method that directs thinking along the paths most likely to lead to the goal, less promising avenues being left unexplored (Boden 1977, 347)

and which "offers solutions that are good enough most of the time" (Feigenbaum and Feldman 1963, 6). Two main types of heuristic may be isolated - general heuristics, common to most types of problem, and special-purpose heuristics, restricted to a specific problem area (Feigenbaum and Feldman 1963, 6). The former tend to define the type of search strategy, while the latter define the actual route that the search takes, usually by calculating a crude weighting for each decision in an attempt to simulate the human ability to assess approaches and situations. Search strategies and heuristic weighting will be discussed in more detail below.

EXPERT SYSTEMS

A criticism of work in artificial intelligence during the late 1960's and early 1970's (for example, Sammet 1971) was that it was on the whole applied to relatively trivial problems, particularly games such as chess and draughts, for instance, albeit with some success. As a result, the 1970's saw the development of a number of computer systems devoted to solving practical problems in the real world.

Such systems attempted to formalise the knowledge of human specialists (or domain-experts) in order to model their reasoning and thus be able to reach conclusions that would correspond with those of the human experts given similar circumstances. Examples of such 'expert systems' include DENDRAL, a system which analyses the structure of organic compounds and can outperform human experts in the field (see, among others, Buchanan and Feigenbaum 1978); PUFF, which interprets lung function tests (for example, see Hayes-Roth, Waterman and Lenat 1983, 32-5); MACSYMA, which performs symbolic differential and integral calculus (for example, Barr and Feigenbaum 1982 vol.2:143-49); INTERNIST, which carries out internal medical diagnoses (see Barr and Feigenbaum 1982 vol 2:197-202); MYCIN, which identifies bacteria in blood and urine and recommends therapy (Shortliffe 1976); PROSPECTOR, designed to aid mineral exploration (Duda et al 1979); and a number of others (see, for examples and references, Barr and Feigenbaum 1982 vol.2; Michie 1982 fig.14.4; Weiss and Kulikowski 1984 table 1.1; Winston 1984).

Expert systems that have been specifically applied to archaeological problems include a system for ageing domestic animals from their teeth (Brough and Parfitt 1984); a classification system for Beaker pottery (Bishop and Thomas 1984); a simulation program for teaching purposes on the excavation of a burial mound (Dean and Nichol 1984); and a classification system for wine amphorae (Bourrelly and Chouraqui 1984).

Before discussing the application of expert systems to archaeological problems, a brief examination of their characteristics and operation will be attempted.

An expert system is a 'problem-solver' program which has knowledge of a narrowly-defined area built into it, together with an internal model of a domain-expert's reasoning and problem-solving know-how, and uses these features to simulate the performance of the human specialist in solving problems in the same field. The system is presented with a problem or situation by the user, and is not only able to request further information if required, but it can also explain its steps in reasoning. Thus the system may be interrogated about questions that it has asked the user, or about the stage it has reached in its investigation, or about the conclusions it has offered. In addition, the expert system should also have a corrective facility so that new knowledge may be added or faulty knowledge altered by the human expert. Expert systems are seen as 'models of competence' (Hartley 1981); they are not intended to act as models of human reasoning (as a number of earlier 'problem-solvers' were), but are primarily performance-orientated, designed to produce competent results in the domain of expertise. They:

> are not concerned with similarities between the resulting systems and human performance (except insofar as the latter may provide a possible hint about ways to structure the domain or to approach the problem, or as a yardstick for success...) (Davis and King 1977, 306).

Ultimately, an expert system is seen as acting as a consultant or advisor. Weiss and Kulikowski see expert systems as serving as "interactive intelligent problem-solving and advisory systems that augment the capabilities of the user" (1984, 10). Michie (1982, 197) divides this consultancy function of expert systems in three ways, which he defines in terms of "user-modes". The expert system may be approached with the user in the role of a client or pupil, obtaining information and learning about the specialism of the expert system through a series of questions and answers. Alternatively roles may be reversed and the user acts as tutor: refining, modifying or enlarging the system's knowledge.

An expert system may be seen as having three main parts - first, the knowledge base which contains the rules and facts about the domain; secondly, a control system or inference engine which applies the rules; and thirdly, a global data base consisting of a store of information concerning the consultation at hand (Weiss and Kulikowski 1984, 41). A knowledge-based system is fundamentally different from a conventional program in that the knowledge is made explicit in the form of facts and rules, rather than being implicit in the coding of the procedures and modules found in a standard system (Sowa 1984, 278). The system is rule-driven, in that problems are solved by the application of rules from the knowledge base, although the actual representation of these rules may vary.

KNOWLEDGE REPRESENTATION

The most common form of knowledge representation, and arguably the most successful in terms of results to date, is the 'production system' (see, among others, Davis and King 1977; Weiss and Kulikowski 1984). A production system knowledge base consists of a series of production rules and facts such as are found in MYCIN and DENDRAL and are also the basis of a variety of commercially available expert system 'shells' (empty expert systems to which a knowledge base may be added by the user). A production rule is basically of the type "If [A], then [B]", also known as a 'situation-action pair'. An example of such a production rule taken from an archaeological expert system for ageing animal teeth is:

```
<tooth-n> evidence-for <age-range>
     if   <tooth-n> identified-as <type-of-tooth>
     and  <type-of-tooth> present-from <age-range>
                         (Brough and Parfitt 1984, 53)
```

In this case, the knowledge base contains a list of types of teeth which are associated with a particular age range, and this rule links the facts known about these age ranges with the identity of tooth-n. An alternative example (after Doran 1977, 437) is:

RULE A: if [two graves intersect]
then [they are likely to be well-separated in time]

RULE B: if [two graves are adjacent]
and [they do not intersect]
then [they are likely to be similar in date]

Note that both halves of the 'if' statement in rule B must be true if the conclusion is to be reached. Another rule might use an 'or' construct, in which case at least one of the parts of the 'if' statement must be true.

It should be pointed out here that the 'rules' in an expert system are not necessarily related to archaeological 'laws' as characterised by, for example, Stickel and Chartkoff (1973), Schiffer (1976) and Salmon (1982). An archaeological expert system might well employ rules that could be classified as 'laws' or 'law-like statements', but an expert system 'rule' may equally well be seen as a 'rule-of-thumb' or a heuristic device. Such rules may appear to be trivial, but this triviality is largely derived from these often tacit inferences being made explicit. Both 'rules' and 'facts' are found in the knowledge base, the distinction between the two being that facts are merely simple cases of rules, having no antecedents and therefore nothing to prove.

Production rules are fundamental to most expert systems in one shape or another. Their format is rigid but even so they can carry different nuances of meaning. Rules A and B above work from observations to a hypothesis, but rules can also work from hypothesis to hypothesis - for example:

RULE C: if [the two graves are likely to be similar in date] (from rule B)
then [any grave goods in the graves are likely to be contemporary]

Here, an inference can be made between two hypotheses, where the antecedent of rule C may be matched with the consequent of rule B. In addition, rules can be used to structure or direct the reasoning process: if the system has been told by the operator that the occupant of the grave in question was male, then there is little point in the system asking questions which were designed to determine the sex of the burial. An ability to trap pointless questioning will also enhance the apparent intelligence of the system in the eyes of the user.

The most important feature of a knowledge base is that

it should be well-structured, that is, rules should not be scattered at random throughout it. This simplifies the task of altering the knowledge base at some future date, particularly if the knowledge base is large. A problem that can arise is that a number of different rules may be found that have matching antecedents or consequents, and some decision has to be made by the system as to which are less relevant, or less suitable. Whether a rule is appropriate or not is dependent on the given situation. Some means of resolving the conflict and selecting the 'best' rule is therefore needed. One way of achieving this is to group the rules into 'procedures', in which all the rules work towards a single goal (Davis and King 1977, 315; Lenat et.al. 1983). These procedures are selected and triggered by 'meta-rules' which direct the system's reasoning to the relevant procedures for a given situation. This clearly involves a larger rule base - not only are there additional rules in the form of the strategic 'meta-rules', but some duplication of rules between procedures may also be necessary.

A development of this kind of technique is the form of knowledge representation known as 'frames', (otherwise known as scripts or prototypes) (see Minsky 1975; Winston 1984; Barr and Feigenbaum 1982). Frames represent knowledge about specific situations, and provide a structure or framework consisting of a number of 'slots' which carry information. The advantage of frames over procedures is that they group the knowledge more explicitly into relevant contexts. Both frames and 'meta-rules' act as an index to the knowledge base and focus the reasoning process. Differences between the two types of representation should perhaps be seen more in terms of organisation than any suggestion that one or the other is necessarily better or more effective. Given the proven success of production systems in problem-solving, these will consequently be the type considered in the remainder of this paper.

What is clear from a brief examination of knowledge representation in expert systems is that they are quite rigid in format and need to be well-structured. This may raise difficulties when such systems are applied to an archaeological problem domain. The effect of formulating 'rules' and 'facts' in a relatively unstructured domain such as archaeology is to make the subject appear to be more objective or 'scientific', whereas in fact the process may simply reinforce the subjective nature of the inferences and concepts employed.

REASONING IN EXPERT SYSTEMS

In a production system, rules are evaluated by a series of pattern-matching operations (see, for example, Michie 1982,199-201) which are controlled by the inference engine. This inference engine controls the reasoning of the system and applies the rules found in the knowledge base by logically matching the premise of an argument with a rule. The inference engine does not change when new knowledge is added to the system since it is held separately from the knowledge base. In the example given above, presented with the premise 'Two graves, X and Y, intersect', the system might find a match with the the antecedent of rule A which would trigger its consequent, 'The two graves, X and Y, are likely to be well separated in time.'

A language such as PROLOG has this associative or pattern-matching facility built into it, while expert system 'shells' generally consist of an inference engine stripped from a well-tried expert system to which any knowledge base can be added. For example, EMYCIN ('Essential' or 'Empty' MYCIN) is an expert system shell consisting of the inference engine that has been removed from MYCIN and has subsequently been used as the basis of other systems, such as PUFF.

As well as controlling the logical pattern-matching process, the inference engine also handles the overall search strategy. The type of search conducted by the inference engine is characterised as either being 'forward chaining' or 'backward chaining', though the two types may be combined in a single system.

A forward chaining search strategy will start from known facts and work towards a conclusion by matching a premise with the antecedents of a rule and triggering off the consequents of that rule (a process which was described above using rule A as an example). The problem with such a system is that it is difficult to control; rules used in this way are known as demons, and may be fired off by the addition of a single piece of information to the knowledge base. The SOLCEM-D system proposed by Doran (1977) for cemetery analysis employed 'recognition demons', for example.

Backward chaining works from the consequents of rules to their antecedents. The system searches for a rule with a consequent which matches the premise and attempts to prove the antecedent by treating it as a new premise. Using the example shown above from Brough and Parfitt's system for ageing animal teeth (1984, 53), if the consequent '<tooth-n> evidence-for <age-range>' is matched by the premise, the system then takes this consequent as its new premise and attempts to prove the antecedents - '<tooth-n> identified-as <type-of-tooth> <u>and</u> <type-of-tooth> present-from <age-range>'. The language PROLOG, and the expert system shell APES (an augmented PROLOG with a query facility) both operate using a backward chaining mechanism. In other words, a forward chaining inference engine works from known facts towards a solution, while a backward chaining inference engine works from a conclusion back to the facts - it effectively 'guesses' at a likely solution and attempts to prove it, moving on to another solution if the proof fails. Both types of search may be found in a single

expert system: MYCIN has a backward chaining inference engine yet employs antecedent rules which work by forward chaining, triggering demons in order to avoid asking the user pointless questions.

Search strategy may also be controlled in a more immediate way: by applying heuristic weightings to rules in order to guide the system's reasoning by associating weightings with potential solutions. Such heuristic weightings may not only guide the reasoning of the expert system but can also indicate to the user the level of confidence that can be placed in the system's conclusions. In some cases, for example the expert system R1 (Sowa 1984, 284), the handling of such weightings has been largely replaced by exact logical reasoning, but it seems highly improbable that an archaeological expert system could operate successfully without some means of handling uncertainty. Archaeology is very much an intuitive discipline and is not as formalised as the scientific and engineering disciplines to which expert systems have usually been applied. As a result, the ability of an expert system to handle uncertainty is particularly relevant when considering an archaeological problem domain. A generalisation in archaeology, such as 'adjacent graves in a cemetery are likely to be similar in date', is simply a rule-of-thumb: it may be patently incorrect in some circumstances, and not all archaeologists would agree that such a statement was valid. This kind of implicit heuristic in a rule may work adequately in a well-structured area, but for most applications, and particularly archaeological problems, a more flexible means of handling uncertainty is required of an expert system. Archaeologists have to deal with large amounts of inexact data, and no knowledge base could ever be considered to be exhaustive or complete.

HANDLING UNCERTAINTY

To date, expert systems have handled uncertainty through a variety of quasi-probabilistic methods using calculated heuristic weights. These are variously known as probabilities, degrees of belief or confidence factors (see Weiss and Kulikowski 1984 for examples). The most common representation of confidence factors is to associate a number between -1 (or 0) and +1 to a particular statement, where -1 (or 0) indicates no confidence in the statement, or lowest probability, and +1 equals highest probability, or certainty. Confidence factors associated with rules and facts that are employed during a consultation are combined using a number of techniques, such as variations on the theme of Bayes' Theorem (Shortliffe's Model of Inexact Reasoning for instance) or introducing elements of fuzzy logic (see Stefik et al. 1982 for examples), in order to give an overall indication of the degree of confidence that may be permitted in the conclusions. Thus, for example, rule B might appear as follows:

if [two graves are adjacent]
and [they do not intersect]
then [they are likely to be similar in date (with a confidence factor of 0.7)]

In this case the confidence factor was assigned arbitrarily and this is a major weakness of the expert system's ability to handle uncertainty. However reliable the manipulation of the confidence factors may be, their initial values, and any values subsequently assigned to information provided by the user during the course of a consultation, are provided by an individual. Worse still, they may be provided by several individuals, separated in time and place. The values given and their consistency will depend ultimately upon the the confidence that the user feels at that time. Initial values for expressing degrees of certainty in a hypothesis are only approximations: the level of belief that can be displayed in the confidence factors therefore may not be certain, compounding the problem that they were in fact intended to alleviate. In addition, once a confidence factor has been subjectively assigned, there remains the difficulty of the probability that is implicitly assigned to the opposite condition - in rule B, the unstated implication is that there is a 30% chance that these graves are <u>not</u> similar in date.

> Many studies have shown that experts do not easily carry probabilities ... in mind, and even when they do report them, the numbers do not turn out to be true probabilities, because the reporting specialist will often deny agreement with the hidden implicit probability of the opposite conclusion (Weiss and Kulikowski 1984, 28).

A warning about such probabilities was voiced by McCarthy and Hayes:

> The information necessary to assign numerical probabilities is not ordinarily available. Therefore, a formalism that required numerical probabilities would be epistemologically inadequate (1969, 490).

Inadequate or not, the fact remains that not only are these <u>ad hoc</u> probabilities used, they are used with some success in systems like MYCIN, DENDRAL and PROSPECTOR. However, given the problems associated with the prior assignment of confidence factors in an expert system, it may seem to some that the handling of such probabilities is something of a juggling act with numbers. For example, Cendrowska and Bramer criticise MYCIN on the grounds that:

> there are interdependence restrictions that need to be applied to the estimation of certain parameters (measure of belief and measure of disbelief in a hypothesis, supplied by the physician) ... which are not included in the MYCIN model. In addition, the use of certainty factors as a means of ranking hypotheses is also suspect since examples can be given of cases where, of two hypotheses, the one with the lower probability would have the higher certainty factor .... It is interesting that such a flawed method should give results that are apparantly acceptable in practice (1984, 489).

If the prior probabilities are open to question, then however sound the statistical procedures may be, the resulting probabilities cannot be trusted.

An alternative means of expressing uncertainty is to use 'fuzzy logic' (see Zadeh 1979). Using fuzzy logic, imprecise terms such as 'few', 'several' or 'many' may be used, each term having been previously defined as having a set of values. For example, 'several' is likely to imply more than two, but less than ten, therefore a low probability would be assigned to the ranges 1-2 and >10 with a high probability for the range 3-9. Quite clearly, however, the same criticisms made of the prior assignment of confidence factors may also be levelled at these fuzzy logic sets. Indeed, while they may provide greater flexibility, they also tend to conceal the assumptions that were made in their formulation more effectively.

EXPERT SYSTEMS AND ARCHAEOLOGY

Having discussed the general characteristics of expert systems and touched on some of the problems that archaeologists may encounter in knowledge representation and in dealing with uncertainty, one question remains to be considered: why (or whether) archaeologists should be interested in these systems. This is particularly relevant since a number of the most frequently cited expert systems are not widely employed in their problem domain. Of those mentioned earlier, INTERNIST, MYCIN and PROSPECTOR are not "extensively used" (Weiss and Kulikowski 1984 table 1.1) in spite of having a good track record in terms of reliable results. Indeed, of all the medical expert systems only PUFF is in routine use, the reason apparently being that:

> ... they have yet to satisfy the indispensability criterion: They are not indispensible to the practice of medicine, and physicians perform adequately without them (Barr and Feigenbaum 1982, 183).

Thus it need not be seen as imperative that archaeologists adopt these systems wholesale, indeed in many situations they may be totally inappropriate, either for practical or ethical reasons.

One of the major reasons put forward for the adoption of expert systems is that they can be used as a means of communicating expertise between specialists in different subject areas. As Bishop and Thomas (1984, 56) point out, specialisation in archaeology is now the rule. This is clearly the raison d'etre of the system for ageing domestic animals (Brough and Parfitt 1984), a field where incorrect interpretations may result from the lack of immediately available expertise (op.cit., 49). A similar argument is used by Ennals and Brough (1982) who see expert systems as a kind of reference work that may be consulted by the non-specialist. Weiss and Kulikowski believe that this dissemination of expertise is a major motivation for building expert systems, and claim that they can save both time and money and may be more accurate and certainly more efficient than human experts, who:

> are in short supply, and when available, have little time at their disposal. While they may be very proficient at their work, the possibility of distraction by many different conflicting demands makes them more vulnerable to errors than a computer -based system (Weiss and Kulikowski 1984,9).

Associated with this is the belief that expert systems may be useful not only in modelling the expertise of one specialist (for example, Clarke's classification of Beaker pottery - see Bishop and Thomas 1984) but that of several, thereby enabling comparisons to be made between the variety of approaches used by experts within a particular specialism. One of the trends in expert systems outside archaeology is to attempt to synthesise the experience of a number of experts in order to enhance the expertise offered by the system.

This formalisation of knowledge is claimed to be one of the most important reasons for constructing an expert system, and it is considered that such systematisation is beneficial to the host discipline. Weiss and Kulikowski, for example, argue that:

> in formalising the knowledge of how an expert human solves difficult problems with today's best knowledge, we are laying out explicitly how future alternatives may be sought. As long as the expert states his reasoning only informally and imprecisely, it is impossible to pin down the alternatives; but as soon as there are formalised statements that enable a computer to reproduce the outcome of human reasoning, we can proceed to experiment and see under what circumstances these statements are applicable (Weiss and Kulikowski 1984, 10).

A number of major assumptions are made in this statement which raise questions about the use of expert systems in archaeology, largely associated with the problem of the formalisation of knowledge that is implicit in an expert system. The problem of formalising archaeological knowledge is not emphasised by the proponents of archaeological expert systems to date, but it is there nevertheless. There remains considerable disagreement within the profession as to whether formalisation across the board is desirable or

even possible. Those archaeologists who do not subscribe to the 'scientific' approach to archaeology, as typified by attempts to discover universal generalisations, will consequently object to the formalisation necessary for an expert system. Conversely, those who believe that the scientific approach is valid will presumably be more willing to see expert systems as having a role to play in the systematisation of archaeological knowledge.

However, the degree of formalisation necessary to construct an expert system is a form of reductionism, in that the translation of knowledge from the implicit to the explicit will inevitably involve the loss of elements in the process. The formal, fundamentally deductive reasoning of an expert system may be inapplicable when dealing with some types of archaeological problem and the reduction resulting from formalising a body of knowledge can have profound consequences:

> Science can proceed only by simplifying reality. The first step in its process of simplification is abstraction. And abstraction means leaving out of account all those empirical data which do not fit the particular conceptual framework within which science at the moment happens to be working (Weizenbaum 1976, 127).

Even stating the means by which a solution was arrived at is not a straightforward process since:

> when we explain to someone how we solved a problem, we often invoke 20-20 hindsight and leave out the mistakes that we made along the way. Our explanation makes it appear that we followed a very direct and reasonable route from beginning to end (Stefik et al. 1982, 162).

As if the problem of abstracting archaeological knowledge is not severe enough, expert systems potentially present the additional hazard of fossilising the "particular conceptual framework" that was current at the time of abstraction. In archaeology, however, to use Hurst's words, "the hypotheses of yesterday become the beliefs of today and the untruths of tomorrow" (1964, 149). Archaeological knowledge is in a state of constant flux: systematising that knowledge could have serious implications for its future development. It may be that this over-states the problem, but even though expert systems can 'learn', there is a danger that once an expert system has been constructed, its area of expertise might be seen as 'complete'.

This fossilisation of knowledge in an expert system is not helped by the ultimate inability of the system to justify all its rules and assumptions (see Clancy 1983). The system's response to a user's request for information or 'reasons why' is to quote rules from its knowledge base as justification for making a particular assumption or inference. However, as the explanation moves up from the particular to the general the rules that can be invoked as an explanation begin to run out, until at last the machine has to respond with something along the lines of "because I was told it was the case". Some, but not all, rules will be justified by facts held in the knowledge base, but assumptions normally have to be made in dealing with an archaeological problem, and such assumptions may not be made explicit in an expert system. Even if they are stated, the explanation may be inadequate and the system unable to provide further information. This would perhaps be acceptable if the system was being operated by an expert in the same field, but could cause considerable problems if the system was being used by a pupil or non-specialist when incorrect or inapplicable advice may be the result. In other words, the system does not have 'complete' knowledge about the conceptual framework within which it was constructed. While the knowledge base could be extended, perhaps by the inclusion of bibliographical references for example, the ultimate inability of the expert system to justify itself will always be present. This problem also exists in dealing with the confidence factors associated with hypotheses - these weights may be assigned by different people at different times for different reasons, but the expert system, other than checking them for internal consistency, will be unable to explain why certain levels of confidence were assigned to particular hypotheses.

THE SUITABILITY OF ARCHAEOLOGY AS A HOST DISCIPLINE

Discussion of the motivations behind building expert systems, and their associated problems, raises a further question - to what extent is archaeology a suitable medium for the application of these expert systems? Perhaps not surprisingly, those who have worked on archaeological expert systems claim that this is not in doubt. Doran, for example, considers that "archaeology has clear attractions as a problem domain for artificial intelligence research" (1977, 433), though he adds that archaeological problems have their own special characteristics and need special treatment. Most expert system specialists appear to believe that there is no area to which expert systems cannot be eventually applied. Probably the most useful approach to this question is to examine the particular characteristics of a problem domain that specialists in expert systems look for, and to see how far archaeology conforms to these requirements.

Artificial intelligence experts are generally agreed on the requirements of a suitable problem domain. For example, Stefik et al. (1982, 142) isolate three fundamental characteristics of a problem domain: the data and knowledge should be reliable; they should be static; and the number of solutions should be small. It is possible to relax these requirements - for example, unreliable data can be handled to some extent using confidence factors or fuzzy

logic - and indeed there are few, if any, problem areas that can be said to meet these requirements in full. However, it may be argued that the nearer an application is to these three ideal requirements the more reliable that expert system will be. Sowa defines three common features of application areas:

> first, there exist recognised experts in these fields; second, the knowledge that the experts have is quantifiable; third, the knowledge can be expressed in declarative rules instead of procedures. If a problem is so difficult that no human being knows where to begin, no computer system will be able to solve it either. If the problem requires intuitive judgements about novel situations .... then the intuition cannot be formalised in an expert system. The best applications for an expert system are ones that require a large amount of well-defined, formalisable knowledge (Sowa 1984, 285).

Barr and Feigenbaum conclude from an examination of medical expert systems that:

> The domain must be narrow and relatively self-contained, the computer should provide substantive assistance to the physician, and the task should be one that the physician either cannot do or is willing to let a computer do. (Barr and Feigenbaum 1982, 183).

It is therefore apparent that archaeology is very much on the fringe as far as the application of expert systems is concerned. Archaeology may have its experts, but it is doubtful whether their knowledge is quantifiable, and even if it were, archaeological knowledge is not complete, reliable or static. This must cast some doubt upon the utility of expert systems for archaeology in general.

Expert systems that have been successfully applied have, on the whole, been used in areas where the problem and the theory associated with it is well-structured. That is, they have been applied to problems which are capable of being effectively reduced to a logical system of relationships. This point may be made concerning the archaeological expert system used to age domestic animals (Brough and Parfitt 1984). However, most archaeological theory and knowledge is as yet insufficiently well-structured or defined to sustain expert systems in their current form. Weizenbaum for example, compares the interpretation of the structure of molecules from mass spectrometer output (which DENDRAL was successfully designed to carry out) to an archaeological situation; a problem which:

> ... is somewhat analagous to that of reconstructing the life of a prehistoric village from the remains recovered by archaeologists. There is, however, an important difference between the two problems: there

> exists a theory of mass spectrometry .... the analyst is in a better position than the archaeologist, who has no strong methods for verifying his hypotheses (1976, 229-30).

It is probably this feature which makes archaeology attractive to proponents of artificial intelligence and expert systems in particular. Archaeology presents them with a challenge, but the results of that challenge may be of little use to archaeologists (Hunt 1975, 442). Those in favour of expert systems suggest that they will assist progress towards formalisation, but this will be seen as beneficial only by those archaeologists who wish to travel along that road. Archaeology does not seem to be the natural candidate for the application of expert systems that some would appear to believe (for example, Ennals and Brough 1982; Bishop and Thomas 1984).

However, it may be suggested that there are some areas within archaeology which are capable of sustaining the systematisation required for an expert system, and in which such a system might actually be of some use. The prerequisite of such an area will be that there already exists a body of formal and systematic knowledge which may be encapsulated in an expert system. To paraphrase Barr and Feigenbaum (1982, 183): the domain must be narrow and relatively self-contained, the computer should provide substantive assistance to the archaeologist, and the task should be one that the archaeologist either cannot do or is willing to let a computer do. Certain specialist areas, such as stratigraphical analysis, may be considered to be suitable candidates. Here, for example, is a subject which has well-defined and generally accepted 'rules', but which is a complex and repetitive task for a human. Stratigraphical sequences can already be organised by a variety of computer programs and arguably the addition of a degree of 'expertise' to these would greatly enhance their performance. The example set by DENDRAL could suggest that those highly scientific areas of archaeology, such as radiocarbon dating and thermoluminescence, might be suitable areas for the application of expert systems. Other areas may develop with time, but it would seem reasonable to require that such areas achieve a degree of stability before an expert system is applied, rather than during the process of its application. In other words, we should not be in the situation of having expert systems looking for an application; the subject should dictate the system and not the reverse. Expert systems that are intended for general use should only be applied to well-defined and 'static' problem areas in archaeology; they are probably not suitable for use in fields that are highly theoretical and subject to the vagaries of new, often mutually exclusive, hypotheses.

Current work is examining the possibility of applying expert systems to the more traditional form of archaeological database in order to assist with the manipulation

of data in a much more limited way. Since archaeological data sets are fundamentally incomplete and inexact, an expert system could be of some use in partially overcoming some of these problems for the purpose of a particular analysis. In the analysis of a cemetery, for example, although the sex of many of the bodies may be unknown, in some cases it may be possible to infer the likely sex from, say, the grave goods and an expert system could be used to indicate likely male and female burials using a variety of techniques. This kind of 'intelligent database' system would have to be used with extreme caution since it may justifiably be claimed that it would be 'creating' data and thereby contaminating the database proper. However, this could be offset to some extent by the ability to 'switch off' the expert system and in addition, the data that it 'creates' during the consultation or analysis need not be saved in the database at the end of the session. Caution must still be exercised, however, since these systems can only be as good as as the information they contain. While this information may be added to, the system cannot of itself improve a bad hypothesis or concept (Weizenbaum 1976, 35).

It is suggested, therefore, that while there may be areas in which archaeological expert systems may be usefully applied, they should for the present be seen as research tools rather than be immediately placed on general release in the field. The widespread and uncritical adoption of expert systems could result in the stifling of archaeological theory, since the encapsulation of archaeological knowledge may in fact lead to its stagnation. Discussion and development of new ideas may be seen as a sign of a healthy discipline: the gracing of a tool with the title of 'expert' could forestall this kind of activity. Leach emphasises the role of human intuition in archaeology:

> Computers and similar gadgetry have their proper place in archaeological method, but do not forget that in the past real progress in your subject .... has always originated in an inspired guess (1973, 771).

REFERENCES

Barr,A. and Feigenbaum,E.A. 1982. The Handbook of Artificial Intelligence vols. 1 & 2. Pitman.

Bishop,M.C. and Thomas,J. 1984. '"BEAKER" - An Expert System for the BBC Micro', Computer Applications in Archaeology, 56-62.

Boden,M. 1977. Artificial Intelligence and Natural Man. Harvester Press.

Bourrelly,E. and Chouraqui,E. 1984. 'Systeme Expert et Simulation d'un Raisonnement en Archeologie', Informatique et Gestion 151,46-51.

Brough,D.R. and Parfitt,N. 1984. 'An Expert System for the Ageing of a Domestic Animal', Computer Applications in Archaeology, 49-55.

Buchanan,B.G. and Feigenbaum, E.A. 1978. 'DENDRAL and MetaDENDRAL; Their Applications Dimension', Artificial Intelligence 11(1), 5-24.

Cendrowska,J. and Bramer, M. 1984. 'Inside an Expert System: A Rational Reconstruction of the MYCIN Consultation System' in O'Shea,T. and Eisenstadt,M. (eds.), 1984, 453-97.

Clancy,W.J. 1983. 'The Epistemology of a Rule-Based Expert System - A Framework for Explanation', Artificial Intelligence 20 (3), 215-51.

Davis,R. and King,J. 1977. 'An Overview of Production Systems' in Elcock,E.W. and Michie,D. (eds.), 1977, 300-32.

Dean,J. and Nichol,J. 1984. 'Burial Mound - a Simulation of an Archaeological Dig', Computer Applications in Archaeology, 38-48.

Doran,J. 1977. 'Knowledge Representation for Archaeological Inference' in Elcock,E.W. and Michie,D. (eds.), 1977, 433-54.

Duda,R., Gaschnig,J. and Hart,P. 1979. 'Model Design in the Prospector Consultant System for Mineral Exploration' in Michie,D. (ed.), 1979, 153-67.

Elcock,E.W. and Michie,D. (eds.) 1977. Machine Intelligence 8. John Wiley.

Ennals,R. and Brough,D.R. 1982. 'Representing the Knowledge of the Expert Archaeologist', Computer Applications in Archaeology, 56-62.

Feigenbaum,E.A. and Feldman,J. (eds.) 1963. Computers and Thought. McGraw-Hill.

Fink,D.G. 1968. Computers and the Human Mind. Heinemann.

Hartley,R.T. 1981. 'How Expert Should an Expert System Be?', Proceedings of the Seventh International Joint Conference on Artificial Intelligence vol II, 862-67.

Hayes-Roth,F., Waterman,D.A. and Lenat,D.B. (eds.) 1983. Building Expert Systems. Addison-Wesley.

Hayes,J.E., Michie,D. and Mikulich,L.I. (eds.) 1979. Machine Intelligence 9. John Wiley.

Hunt,E.A. 1975. Artificial Intelligence. Academic Press.

Hurst,J.G. 1964. 'White Castle and the Dating of Medieval Pottery', Medieval Archaeology 6-7, 135-55.

Leach,E. 1973. 'Concluding Address' in Renfrew,C. (ed.), 1973, 761-71.

Lenat,D.B., Davis,R., Doyle,J., Genesereth,M., Goldstein, I., and Schrobe.H. 1983. 'Reasoning about Reasoning' in Hayes-Roth,F., Waterman,D.B., and Lenat,D.B. (eds.), 1983, 219-39.

McCarthy,J. and Hayes,P.J. 1969. 'Some Philosophical Problems from the Standpoint of Artificial Intelligence' in Meltzer,B. and Michie,D. (eds.), 1969, 463-502.

Meltzer,E.W. and Michie,D. (eds.) 1969. Machine Intelligence 4. John Wiley.

Michie,D. (ed.) 1979. Expert Systems in the Microelectronic Age. Edinburgh University Press.

- 1982. Machine Intelligence and Related Topics. Gordon and Breach.

Millar,P.H. 1971. 'On Defining the Intelligence of Behaviour and Machines', Second International Joint Conference on Artificial Intelligence, 279-86. The British Computer Society.

Minsky,M.L. (ed.) 1968. Semantic Information Processing. MIT Press, Mass..

- 1975. 'A Framework for Representing Knowledge' in Winston,P.H. (ed) 1975, 213-77.

Newell,A., Shaw,J.C. and Simon,H.A. 1963. 'Empirical Explorations with the Logic Theory Machine. A case study in heuristics' in Feigenbaum,E.A. and Feldman,J. (eds.), 1963, 109-33.

O'Shea,T. and Eisenstadt,M. (eds.) 1984. Artificial Intelligence: Tools, Techniques and Applications. Harper and Row.

Raphael,B. 1976. The Thinking Computer- Mind Inside Matter. Freeman.

Renfrew,C. (ed.) 1973. Explanation of Culture Change: Models in Prehistory. Duckworth.

Salmon,M.H. 1982. Philosophy and Archaeology. Academic Press.

Sammet,J.E. 1971. 'Challenge to Artificial Intelligence: Programming Problems to be Solved', Second International Joint Conference on Artificial Intelligence, 59-65. The British Computer Society.

Samuel,A.L. 1963. 'Some Studies in Machine Learning using the game of Chequers' in Feigenbaum,E.A. and Feldman,J. (eds.), 1963, 71-105.

Schiffer,M.B. 1976. Behavioural Archaeology. Academic Press.

Shortliffe,E.H. 1976. Computer-Based Medical Consultations: MYCIN. Elsevier, New York.

Sloman,A. 1971. 'Interactions Between Philosophy and Artificial Intelligence: The Role of Intuition and Non-Logical Reasoning in Intelligence', Second International Joint Conference on Artificial Intelligence, 270-78. The British Computer Society.

Sowa,J.F. 1984. Conceptual Structures: Information Processing in Mind and Machine. Addison-Wesley.

Stefik,M., Aikens,J., Balzer,R., Bendit,J., Birnbaum,L., Hayes-Roth,F. and Sacerdot,E. 1982. 'The Organisation of Expert Systems, A Tutorial', Artificial Intelligence 18 (2), 135-73.

Stickel,E.G. and Chartkoff,J.L. 1973. 'The Nature of Scientific Laws and their Relation to Law-Building in Archaeology' in Renfrew,C. (ed.), 1973, 633-71.

Weiss,S.M. and Kulikowski,C.A. 1984. A Practical Guide to Designing Expert Systems. Chapman and Hall.

Weizenbaum,J. 1976. Computer Power and Human Reason: from judgement to calculation. Freeman.

Winston,P.H. (ed.) 1975. The Psychology of Computer Vision. McGraw-Hill.

- 1984. Artificial Intelligence. Addison-Wesley. (2nd edition).

Zadeh,L.A. 1979. 'A Theory of Approximate Reasoning' in Hayes,J.E., Michie,D. and Mikulich,L.I. (eds.), 1979, 249-94.