

'BEAKER' - AN EXPERT SYSTEM FOR THE BBC MICRO

M.C. Bishop, Dept. of Ancient History & Classical Archaeology, The University, Sheffield, S10 2TN

J. Thomas, Dept. of Prehistory & Archaeology, The University, Sheffield, S10 2TN

This paper describes an attempt to implement a simple expert system on the BBC micro, using a knowledge base derived from Clarke's classification of Beaker pottery, which is particularly well-suited to the task. The proposed system uses a combination of textual and graphics prompts to enable the user to identify and ultimately classify items, with provisions for new types being 'learned' by the program and added to the knowledge base.

The implications of the use of expert systems in archaeology are considered and possible future developments are briefly examined. It is concluded that the main value of expert systems lies in their ability to help us re-appraise our attitudes towards the classification of artefacts.

As the information recovered by the process of excavation becomes ever more detailed, the excavator himself perforce becomes increasingly more divorced from the minutiae associated with every site - gone are the days of James Curle, when the excavator could not only dig Newstead himself, but also write the specialist reports (with one or two exceptions) and produce the overall synthesis of the evidence.

It is for this reason that the archaeological specialist has appeared over the years; the most obvious category being the pottery specialists, since the study of ceramics soon passed the point at which it was possible for the excavator to keep up with all material from his particular period, let alone other periods. Specialists have proliferated and become increasingly more esoteric.

Everybody involved in the profession of archaeology has, at some point, encountered 'the need to specialise' and moved into a chosen field, usually selected with an eye to making them as indispensable as possible, without being too eccentric. It is interesting to speculate where this trend towards specialisation is leading us.

Specialisation is not, of course, limited to archaeology: many professions, particularly the more 'scientific', have their specialists - people who have concentrated on one, fairly narrow, aspect of their subject,

which few others will know in detail, but many will understand in a more or less comprehensive fashion. In many professions, much of this knowledge is only of minor interest to most researchers, but in some - such as medicine - all aspects of this 'knowledge base' may, occasionally, be of interest to a particular person not normally cognisant with its refinements.

To cut a long story short, this tendency towards specialisation has led to the development of a number of computer programs aimed at making knowledge bases, such as specialised areas of medicine, available to other people - the so-called expert systems.

Before proceeding any further, it will be as well to distinguish between a specialist and an expert: for the purposes of this paper, a specialist is human, whilst the term expert is used as shorthand for 'expert system'; besides, specialist sounds rather less pretentious than expert.

Most of the best-known experts, such as PROSPECTOR, run on mainframe computers, simply because of the large amounts of information they need to handle. Systems such as PROSPECTOR have achieved a degree of notoriety by virtue of their successes (Webster 1983).

Experts are available for micros, but they naturally lack the sophistication of their larger brothers. A number of do-it-yourself medical diagnosis programs are available for home micros (Naylor 1984b), whilst details have been published of an expert on fungi which runs on a tape-based MZ-80K (MYCIX - Pratt, 1983), and the recent publication of 'Build Your Own Expert System' by Chris Naylor (1983) proved extremely popular with owners of home computers. Commercially available systems include EXPERT-EASE for the IBM PC and HULK for the BBC micro (Naylor 1984a).

Experts do not have to be sophisticated to be useful to the archaeologist - what matters is that they should be used intelligently. Most tasks which an archaeologist can set an expert are fundamentally simple in nature: if we ask someone to decide whether a pot is like three others, then they will probably answer that it is either like one or more of them, or it is not like any of them. This is an oversimplification, but it might help to demonstrate that an expert need not be overly elaborate to be of use.

The present exercise - the construction of an artefact-orientated expert system for a popular micro - has begun out of curiosity: one of the writers is a 'specialist' and a potential candidate for computerisation, so it was thought worthwhile investigating the feasibility of this sort of system.

At the outset, it is as well to state that the system, in the form of the original specification, is not functioning at the time of writing, but that does not, in our view, detract from the value of the exercise, for the feasibility of the system is such that it is only a matter

of a few weeks' programming away from being a reality.

The specification for BEAKER was slightly unusual for a present-day expert, in that it required the use of graphics, as well as text, although this sort of system will undoubtedly be widespread in only a few years' time (and will probably incorporate video-disc technology). The reason for this is obvious: archaeology is rich in illustrative material, much of which is far more manageable, from the user's point of view, than lengthy descriptions. A similar idea is the use of 'icons' in some of the latest office software, where the selection of certain images (e.g. a filing cabinet) will access the appropriate part of the program (the database management section). Graphics-linked experts are more sophisticated than icons, but the principle remains the same.

The number of subjects suitable for incorporation in BEAKER seemed infinite, but, as its name implies, the classification of Beaker pottery was selected, using Clarke's seminal publication (1970) as the basis for the construction of the knowledge base. BEAKER was designed so that, given a Beaker, a relatively inexperienced user would be able to progress towards a positive identification of the category to which the artefact belonged - the process is conveniently summarised in Table 1.6 of Appendix 1 in Clarke's book (1970, 430).

Briefly, the program would have to isolate a series of characteristics that would allow it to state a probability that a given Beaker belonged to a certain category; further refinement, where necessary, would depend upon the use of minor characteristics to isolate the vessel. However, this is an opportune moment to look in greater detail at Clarke's system and how it can be adapted for our purposes.

Beaker pottery appears in both domestic and sepulchral contexts in Britain from the end of the third millennium b.c. onwards, and is taken to be of continental ancestry. Beakers show a considerable range of variation in decoration and form over the six or seven hundred years of their use, and thus their classification is of considerable importance in terms of both chronology and cultural affiliation over an important period of British prehistory. Furthermore, the existing classification of Beakers is of considerable interest in the context of expert systems and computer use in archaeology in general. For, in producing the classification as his PhD thesis, David Clarke explicitly attempted to develop a system which would minimise the intuitive element in archaeological typology. A specialist will, in the course of his or her career, examine perhaps thousands of potsherds, flints, or bronzes. While their opinion upon an individual artefact will be beyond reproach, they might not easily be able to articulate exactly what criteria led them to that decision.

For David Clarke, then, the development of a means of classification which could be made explicit and reproduced between one worker and another was a polemical point in the

assault against archaeology as 'an undisciplined empirical discipline' (Clarke 1978, xv).

Beaker pottery in Britain was also an excellent showcase for Clarke's theory of material culture being organised as a nested hierarchy of traits, trait-groups, artefacts, and so on, up to culture-groups. Under these conditions, an individual artefact could be seen as existing within a structure of many levels of classification. Not only this, but the definition of groups of artefacts for classificatory purposes was seen as necessarily polythetic. Entities were defined simultaneously by as many different criteria as the archaeologist considered important: 'By being selective, arbitrary, but consistent', Clarke argued, 'we can hope to learn about very complex systems' (Clarke 1978, 471). He was well aware that such a system was an arbitrary way of splitting up the world in order to study it, and bore no necessary relation with the way in which the artefacts concerned were looked upon by the people who made and used them.

With Beakers, one is concerned with a class of artefacts quite separate from all other pottery, yet representing within themselves a continuous range of variation. The groups which Clarke set up to classify Beakers are the peaks within this range of variation: the fuzzy-edged groups which seem to define a series of ideal types. Given this situation, what our expert system must do is not so much merely to drop an artefact into a conveniently-labelled box, as to assess which of these ideals it most closely resembles, and to what degree. Beyond this, it may be necessary to be able to accept that a new specimen may not fit closely enough with any of these types.

Clarke's system is simple in concept, yet complicated in operation. The replacement of the long process of poring over pages and tables, and cross-referring from one part of his monumental work to another, with a simple series of graphically-assisted question-and-answer routines thus seems most worthwhile in the context of relatively unversed staff on site or in the museum. Not only this, but the system is so constructed as to be ideal for this structure of enquiry. The three major attributes selected by Clarke are shape, chosen from nine ideal forms which he illustrates; style, meaning the way in which bands of decoration and blank space are arranged on the Beaker; and which of the minor groups of motifs are employed. Even before this stage, two of the major Beaker groups, 'All-Over-Corded' and 'Barbed-Wire', would be separated out by the computer, with the aid of simple and specific questions relating to the very specific decorations of those groups. These are respectively the use of a single spiral cord impression decorating the entire outer surface of the Beaker, and the use of a cord-wound stamp.

The combination of shape, style, and motif groups will, in most cases, give a fairly definite answer as to which group a vessel belongs. However, in some cases, such as distinguishing between Wessex/Middle Rhine and

Northern/Middle Rhine Beakers, a further level of questioning is needed. This will also act as a check on the identification of other Beakers. One might use either the occurrence of specific motifs or the positioning of motifs in specific zones within the Beaker, but in practice either of these options might prove difficult to operationalise. Instead, one might choose to employ several of these 'minor characteristics' which Clarke lists, all of which show considerable variation from one group to another. Amongst these we can list the external colour of the Beaker, the substance used as filler in the pottery, and the type of comb used to apply the decoration.

Implementation of BEAKER on a BBC micro does not present many problems. Many micro-based experts make use of arrays for the storage of data, but this consumes large amounts of memory and would be impractical if graphics facilities were to be available. The scarcity of second processors for this particular computer meant that a different approach had to be taken: 'virtual memory' employs discs normally used for program or data storage to act as additional memory space for the computer. Index sequential direct-access filing effectively allows the computer to store 'arrays' on the magnetic medium, accessing them by a series of simple formulae. This method is comparatively slow, being dependent upon the access-time of the mechanical disc-drive, but this is more than compensated for by the size of knowledge base thus made possible - this is, of course, dependent upon the size of disc used, so it can range from 100K upwards.

The knowledge base itself is created from the Knowledge Base Creation Suite (KBCS) and it is at this stage that the various menus, assertions, and rules are encoded. It is then a relatively simple matter to switch to the Knowledge Base Interrogation Suite (KBIS) - the real 'expert' - which simply follows the appropriate course, depending upon the user's responses. A 'profile' array is kept in the computer (cf. Pratt 1983, 187) to check on the progress of the expert, and if a previously unknown type is encountered, this profile is used to add that type to the knowledge base by means of KBCS.

This is all completely transparent to the user - they need only answer the questions posed by the computer, which will then take care of the details of classification and rule formation.

What our expert system will do, then, is classify artefacts within a known typology or, if they do not yet appear in that typology, create a new category according to predetermined criteria. In effect, this is an electronic typology: the task of constructing the typology has been taken out of the hands of the archaeologist and given to the computer; all that is needed is for the archaeologist to specify the ground rules, by means of which the typology will be formed.

It is thus possible to envisage an expert system available on-site that will not only classify finds that are known, but also those that are previously unknown; every evening, the site expert could communicate with a central computer, running a main expert, which would then incorporate all that was new from the various sites in operation into its existing knowledge base - hence creating an automatic typology. This would not only provide units working within a clearly-defined geographical zone with instant access to the current typology status, but could also communicate its findings to a large national expert system, thus obviating the need for archaeologists to go through the laborious task of constructing typologies by hand. Meanwhile, the site expert could be selectively updated with the latest discoveries from relevant local (or even national) sites.

An electronic typology like this would, naturally, make the process of publication considerably easier. Most obviously, such an intercommunicating network of experts allows virtually instantaneous 'publication' of material in a form easily accessible to the field archaeologist. Furthermore, the publication of a site could easily make use of a report-writing facility to select those artefacts that would have to be drawn and those which warranted comment.

This is, of course, sheer fantasy. Whilst, on a theoretical level at least, this electronic typology would work, the fundamental flaw lies in the initial concept of the typology - somebody has to make up the rules by which the expert operates, and if these are wrong (and most typologies have their opponents) then the whole system will merely compound the errors. What is needed is a self-testing expert that will regularly check the validity of its typologies and possibly ask the opinion of specialists or other expert systems.

It is well known that Japanese manufacturers are working on a 'fifth generation' of computers, the main feature of which will be their use of artificial intelligence (AI) software. As yet, archaeology has had little to do with AI, but the predicted revolution in both computers and software will almost certainly change this. Expert systems could well be the harbingers of a new era in archaeological computing.

What this study would seem to suggest is that the implementation of artefact-orientated expert systems in archaeology is a valuable process - if only because it will make us question the nature of typologies and the role of the human specialist in their construction. The true role of expert systems in archaeology must lie in advising specialists when making decisions about their particular field; the expert system presents an extremely objective method of executing the specialist's current ideas on typology, and the possibility of comparing the effect of different strategies of classification upon the same group of artefacts could prove particularly interesting.

Clarke, D.L. Beaker Pottery of Great Britain and Ireland (Cambridge 1970)

Clarke, D.L. Analytical Archaeology ed.2 (London 1978)

Naylor, C. Build Your Own Expert System (Wilmslow, 1983)

Naylor, C. 'Discriminating experts' Practical Computing 7:3, 1984a, 108-12

Naylor, C. 'Health and efficiency' Practical Computing 7:4, 1984b, 74-8

Pratt, C. 'Expert advice. An expert system for the MZ-80K' Personal Computer World 6:9, 1983, 184-7

Webster, R. 'Expert systems' Personal Computer World 6:1, 1983, 118-9