

Knowledge Representation for Archaeological Inference

James Doran

Department of Computer Science
University of Essex

INTRODUCTION

In this paper I shall be concerned with the problems of interpretation which archaeologists encounter when they study excavation data from prehistoric cemetery sites, and I shall consider these problems in the light of the evolving artificial intelligence theories of specialist knowledge representation.

Prehistoric archaeology is a subject dominated by the need to collect and piece together fragmentary and diverse evidence to form a coherent overall picture. In this task archaeologists have for some years now made increasing use of various mathematical models and techniques of statistical data analysis (Doran and Hodson, 1975). But these techniques are greatly limited by their inability to embody the specialist knowledge required in typical archaeological inference contexts. It is therefore natural to ask whether formal techniques of interpretation can be developed which 'know' something about archaeology. As I discuss this question, and the form which a possible interpretive system might take, I shall try to present both the broad archaeological background and sufficient actual archaeological detail for the issues to come to life.

Archaeology has clear attractions as a problem-domain for artificial intelligence research. Many of the problems of recognition and interpretation encountered in archaeology have close parallels with classic artificial intelligence problems, notably those of scene analysis. But archaeological problems have their own characteristics and need special treatment. Further, any development of 'knowledgeable' interpreters of archaeological data seems bound to force an instructive investigation of the logical relationship between such systems and existing archaeological tools such as factor analysis and computer simulation modelling.

PREHISTORIC ARCHAEOLOGY

The fundamental objectives of prehistoric archaeology are to reconstruct and

to explain the evolution of human society up to the time when written records begin to be available. Archaeology is thus one of the social sciences. But compared with 'mainstream' social sciences such as political science or sociology it has two noteworthy characteristics: the enormous time-span that is its domain, and the relatively 'concrete' nature of its evidence. The great time-span covered by archaeology means that it has a potentially crucial contribution to make to social theory, which is often criticised for being too concerned with structure and not enough with long-term dynamics. Archaeological evidence is 'concrete' by comparison with that of, say, sociology since it is typically formed of directly observable objects or traces rather than of statements by or about people. This is clearly both a strength and a weakness. It makes archaeology somewhat more of a hard science (and archaeologists are often aided by physicists, chemists and biologists) but removes it further from its ultimate social goals.

Paradoxically archaeological data are both embarrassingly abundant and typically insufficient. The evidence collected from field surveys and excavations of prehistoric hunting camps, settlements, fortifications, burial mounds and the like, elaborated by specialist studies (carbon-14 dating, pollen analysis, dendrochronology, magnetometer surveys, bone studies) is typically too extensive to be adequately processed, whence current interest in computer-based methods of data capture and handling. But even so it is unusual to be able to work out with any confidence just what was happening and why in some particular locality at some particular time. It may well be possible to decide that over a period of some five hundred years there was a sequence of four successive houses on some spot with such and such individual characteristics. But it is likely to be much more difficult to work out the economics and trading links of the settlement of which the house, in its successive incarnations, was a part. And social organization is probably just a matter for speculation. The source of the difficulty is partly that the evidence available (there may be much still in the ground) is simply insufficient, partly that the complexity of the inference problem becomes too great, and partly the general weakness of archaeological theory.

Strictly, the task of unravelling the past of human society cannot be split into independent subtasks. Every act of archaeological interpretation is influenced by the results of every other act. But, of course, archaeologists simplify: the subject is broken down into specialities by period or region or type of evidence. Analogously I shall now concentrate on the interpretation of prehistoric cemetery sites, while recognising that this is to some extent a dangerous and potentially misleading thing to do.

CEMETERY SITES AND THEIR INTERPRETATION

Societies almost always dispose of their dead with care and elaborate ritual. Burial, and therefore extended survival of remains, is common. It follows that there are two reasons why funerary evidence is particularly important in archaeology: it tends to reflect social organisation, and there is a lot of it. Look in any archaeological museum! To elaborate the former reason a little, not only

is it likely that burial practice reflects the beliefs of a society, but differential treatment of the dead within a cemetery probably closely corresponds to differential status among the living. These assumptions can be misleading but they certainly have general validity. When comparing cemeteries, the regular assumption made by archaeologists is that clearly different burial practices probably correspond to different societies or to the same society at widely different times. Even making only this weak assumption, such evidence has been of enormous importance in archaeological work.

By a (flat) cemetery I mean a collection of relatively simple graves, dug as pits or shafts, without mounds or other surviving surface monument. Even the simplest grave will, however, be capable of extended description. It will, for example, be possible to specify its location, orientation, size and depth. The burial may or may not be multiple, and may be an inhumation or a cremation. If the former, the posture of the skeleton may be recoverable and possibly its age and sex. There will probably be "grave goods"—remains of clothing, jewellery, weapons, pottery vessels and the like. This list is by no means exhaustive. Commonly only part of the data potentially available from a grave will have been recorded. Similarly it is common for only a fraction of a cemetery to have actually been excavated.

It will be helpful to distinguish between the *structural* and the *social* interpretation of a cemetery. By a structural interpretation I mean the process of moving from the descriptions of the individual graves to a description of the cemetery as a whole expressed in terms of appropriate concepts (for example, age and sex distribution, spatial organisation, internal chronology, typical combinations of grave goods). By social interpretation I mean carrying this process of inference further to statements about the society from which the cemetery is derived. Some examples of typical inferences will make this distinction clearer:

Recognition of coexisting but consistently different burial styles (a structural interpretation) may suggest an established caste structure in the population (social interpretation)

Recurring orientations of the graves in a cemetery (structural interpretation) may prove to be linked to the position of the sun at the time of burial and may even indicate at what seasons of the year an associated settlement was occupied (social interpretation)

A full reconstruction of the growth of a cemetery, including an estimate of its total lifetime (structural interpretation), leads to an estimate of the size and demographic structure of the corresponding population (social interpretation).

An early and mistaken structural interpretation of the famous Hallstatt cemetery (European Early Iron Age) was that it consisted almost entirely of male graves. The equally mistaken social interpretation was that this reflected an all male work force for the nearby salt mines (Kromer, 1959).

Actual archaeological studies of cemetery sites come, of course, from a diversity of geographical areas and chronological periods. Nevertheless the process of interpretation has a certain identity which is independent of any particular cemetery. In particular it is obvious that when interpreting a cemetery, at either structural or social level, the archaeologist brings to the task certain beliefs and assumptions based on knowledge about cemeteries and their relationship to people and societies. Thus he knows or believes that in most populations the sexes are about equally balanced; that cemeteries are laid down over a period of time, not all at once; that it is most unlikely that a woman would be buried with a sword; that an empty grave may be the result either of drowning (with the body lost) or of poor excavation technique; that two quite different burial practices can coexist in the same society at the same time. Table I lists some of the more commonly used beliefs and assumptions and will indicate their diversity.

Each belief or assumption allows the archaeologist to recognise certain properties of the cemetery as significant and to draw immediately useful conclusions. Thus the process of sexing graves, typically based on osteological and social assumptions, provides data for the next step in the inference chain. Inference terminates, of course, with conclusions that have no usable implications. To infer from grave orientation that a cemetery is Christian leads nowhere unless we can infer something (other than grave orientation!) from the fact that a cemetery is Christian, for example that it is relatively late in date.

Cemeteries are rarely studied entirely out of context. There is almost always some parallel evidence available though often surprisingly little. For example, there may be similar cemeteries nearby so that it is really a matter of interpreting the cemetery group rather than any one of them. Or a pottery or metalwork typology may be available based on material from dwelling sites, although it is unusual for a dwelling site and its burial ground to be located and excavated together. Naturally archaeologists take such evidence into account whenever it exists. In a fairly obvious sense it forms a task specific extension of the archaeologist's more general knowledge of cemeteries and their properties.

STATISTICAL AND COMPUTER METHODS

The Hallstatt cemetery mentioned above contains well over a thousand graves (excavated, incidentally, as early as the 1850's). Although this number is unusually large, it is common for a cemetery excavation to uncover several hundred graves. Since a grave is itself a complex structure, and may obtain a score or more artifacts each of which is in turn of considerable complexity (for example, an intricately decorated glass bracelet) there is a great deal of data to be considered. The archaeologist is therefore obliged either to be very selective and subjective or to make use of formal methods of analysis.

The latter option has been more and more exercised in recent years (see, for example; Peebles, 1971; Saxe, 1971; Moberg, 1972; Randsborg, 1973; Hodson, 1975). Commonly what has happened is that archaeologists have resorted to

TABLE I

Skeletons are the remnants of people
 Age and sex determinations are often unreliable
 A predominance of male or female graves requires explanation
 Cemeteries are built up over a period of time
 Weapons imply men, spindle whorls imply women
 Burial grounds are associated with habitations
 Men sometimes die in battle
 Design changes, usually continuously and usually for the better.
 Very different burial practices can coexist
 Deviance is often reflected in burial—an isolated grave, without grave goods, with the skull
 trepanned, makes sense.
 Cemeteries are rarely fully excavated
 Some excavators are unreliable
 Family structure is often reflected in cemetery organisation
 Grave goods may indicate social roles
 Specialist craftsmen are rare
 People have limited memories—the locations of graves may be forgotten
 Artifacts decay at different rates depending upon material and conditions
 Metal work may be traded over long distances
 A sword might have a lifetime of a hundred or more years
 Graves may be oriented by reference to the sun, or to the settlement, or to the “land of the
 dead”
 Exhumation and reburial happens in some rituals.
 Secondary burials may be associated with an important primary
 The level of the ground surface may change dramatically over millenia
 If two graves intersect, they are likely to be well separated in time.
 Women bear children and sometimes die in childbirth
 A randomly generated set of locations may appear to be patterned
 Grave diggers avoid difficult ground
 Charcoal indicates a fire
 People wear clothes
 Adjacent graves are more likely than not to be similar in date.
 A young rich woman is likely to wear the latest fashions
 Graves are sometimes disturbed or robbed
 Most, but not all, people are right handed.

A few of the facts and beliefs commonly employed by archaeologists when they interpret cemetery excavation data. The ordering of the list has no significance. It is an instructive exercise to attempt to recast each of these propositions as one or more “recognition demons” (observing X indicates the existence of a Y with implications Z) of the kind discussed in the text.

relatively straightforward tests of statistical significance at the level of structural interpretation. A typical example is the use of a significance test to check for systematic grave alignment. Often tested too is the null hypothesis that a particular type of artifact is equally common in different sections of a cemetery. Much more elaborate computer based methods of data analysis, notably forms of cluster analysis, of non-metric scaling, and of combinatorial seriation, have also been used. One particular cemetery, that of the la Tène ‘culture’ at Münsingen-Rain near Berne in Switzerland (Hodson, 1968), has provided the

data for a whole series of experiments and a very brief description of it and what has been done should serve to give the flavour of this kind of work.

The Münsingen cemetery dates from about 500 BC to about 50 BC, and was discovered and excavated in the first decade of this century. It turned out to contain about 200 graves, almost all of them inhumations, of which 77 provided skeletal remains which could be sexed (see Figure 1). Also in the graves was a great deal of sophisticated metalwork in iron, bronze and silver, including notably a hundred or more decorated bronze fibulae (brooches). For examples of the grave goods see Figure 2. The importance of the Münsingen cemetery lies in its unique size (for the late Iron Age) and, more specifically, in the possibility of reconstructing the stylistic evolution of the la Tène metalwork designs by recovering the internal chronology of the cemetery, that is, the order in which the graves were laid down. If the evolution of metalwork design can be established in detail, then it can be used to refine the dating of la Tène material found elsewhere.

An initial guide to the chronology of the cemetery is provided by its roughly linear layout (see Figure 1) together with background knowledge of la Tène metalwork which indicates that the northern end of the cemetery is relatively early and the southern end late. This is the starting point for the most important of the conventional studies of the cemetery, that of Hodson (Hodson, 1968). Hodson developed a full artifact typology for the cemetery (that is, the set of artifacts found in the cemetery is divided into classes and sub-classes ('types') first on functional and then on morphological and stylistic grounds). He then used this typology to arrange the graves into their most likely true chronological sequence ('chronological seriation'), invoking the principle that graves are roughly contemporaneous to the extent that they contain artifacts of the same type.

A very different study is that of Schaaff (Schaaff, 1966) who investigated the spatial organisation of the Münsingen cemetery and of other smaller cemeteries of the same type (notably that at Andelfingen—see below). He detected a tendency for graves of a particular age/sex class (male, female, or child) to be grouped together. And Martin-Kilcher (Martin-Kilcher, 1973) has studied in detail the ways in which customs of personal adornment changed during the lifetime of the Münsingen cemetery. For example, women wore neck torcs in the earliest phase of the cemetery, but never thereafter. She succeeded in distinguishing changes in adornment custom from changes in burial custom and from changes in the structure and styling of jewellery. Martin-Kilcher also suggested that the Münsingen cemetery corresponded to a small farming community of not more than three or four households each of about ten persons.

Each of the foregoing studies features the construction of careful arguments based on a mass of detailed evidence. The Hodson study in particular required much laborious description and comparison of artifacts, followed by an equally laborious search for an optimal artifact classification and grave seriation. And this study was, in fact, the starting point for the numerous computer studies

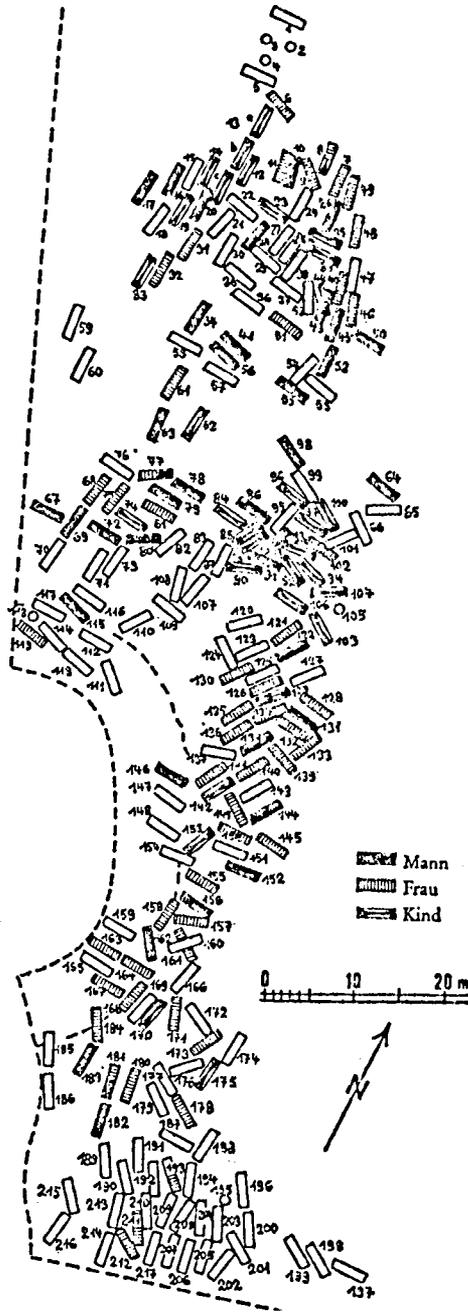


FIG. 1. Plan of the la Tène cemetery at Münsingen-Rain. Graves are specified to be male, female or child wherever determined (Schaaff, 1966, fig. 7).



FIG. 2i. A drawing of the skeleton and the objects associated with it

Grave W-E, 190 cms. deep. Stones in half circle by head. Female (?), 40-60 yrs.

Grave goods:

round neck: 373 a-b

on breastbone: 369-372

on r. forearm: 367

on r. hand: 368 a, 368 b (bent ring fragments)

on ankles: 374 a-b, 375 a-b, in pairs

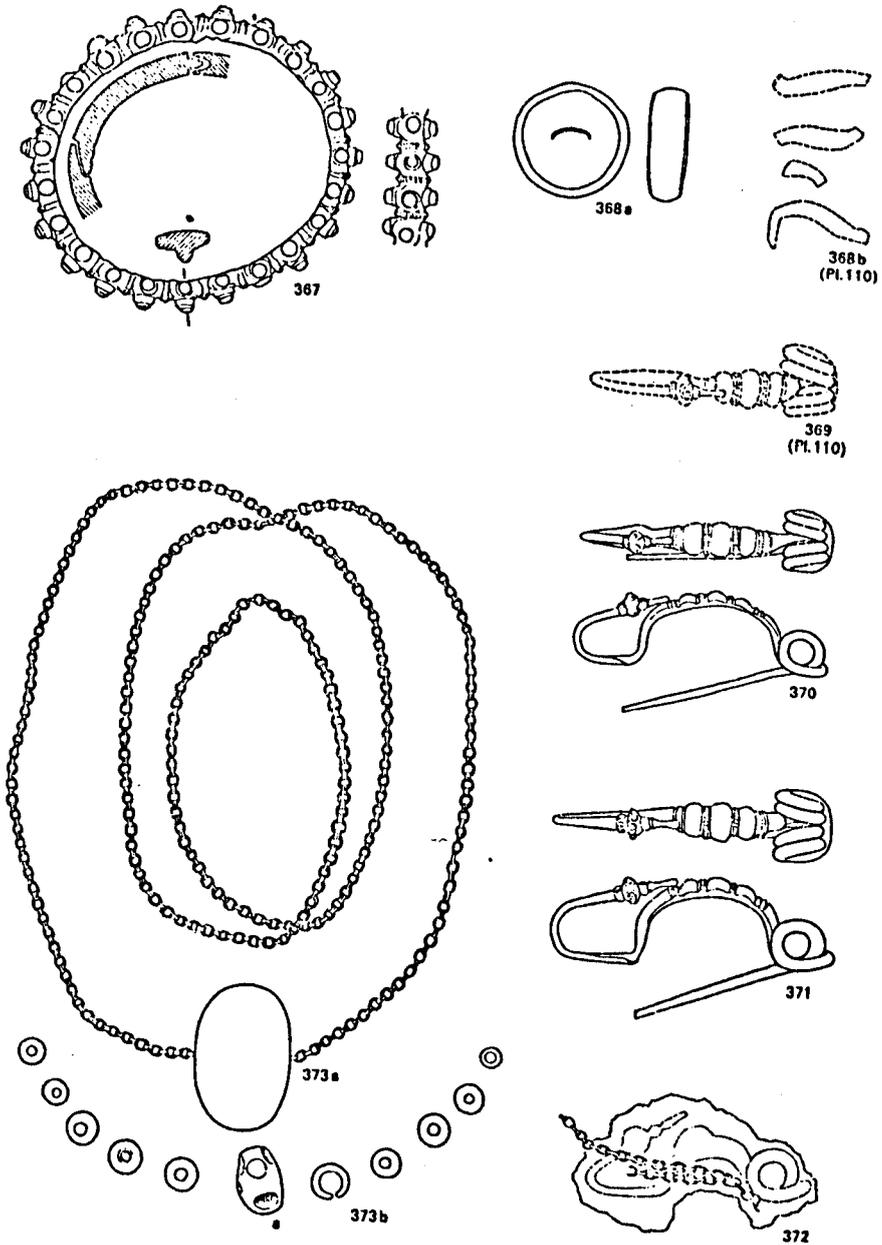
FIG. 2ii. The published description of the grave and its contents

FIG. 2. A typical Münsingen grave [Hodson, 1968, pp57,130]

which have been made of the Münsingen data. These fall into three broad categories:

- (a) cluster analyses (notably K-means analysis) of sets of brooches drawn from the cemetery, with the objective of refining typology
- (b) seriation studies of (a subset of) the graves. That ordering of the graves is sought which most consistently puts similar graves close to one another. As noted above, the assumption is that similar graves are chronologically close, so that an ordering derived in this way must reflect the true chronological ordering.
- (c) non-metric scaling studies of both the fibulae and the graves with objectives similar to those specified in (a) and (b).

Details of, or references to, virtually all of this work will be found in (Doran and Hodson, 1975). It is worth noting that each of these methods of multivariate



368 a-b silver, 372 iron, 373 a bronze and amber, 373 b (a) amber and blue glass, the rest bronze.
Scale 3:4.

FIG. 2iii. A drawing of some of the artifacts found in the grave

FIG. 2(cont'd). A typical Münsingen grave [Hodson, 1968, pp57,130]

data analysis involves a discrete function optimisation problem which is solved by iterative search. The search techniques involved are simple forms of those explored by artificial intelligence scientists within what Nilsson has called the "state-space" paradigm (Nilsson, 1971). There is also a sense in which each of these techniques *interprets* the data. This aspect is clearest in the case of the seriation studies where the techniques can reasonably be said to propose a particular chronological sequence for the graves.

SOLCEM — A CEMETERY INTERPRETATION PROGRAM

The statistical techniques mentioned in the preceding section, from tests of significance to complex clustering or seriation procedures, have in common that they are applicable only to small and artificially isolated parts of the total problem of inferring from a cemetery excavation record to plausible structural and social conclusions. Typically, their use implies that interactions between different categories of data and between different categories of conclusions are ignored. And this is substantially true of conventional archaeological studies. For example, the Hodson and Schaaff studies of the Münsingen cemetery sketched above treat internal chronology and spatial organisation respectively. But each of these two aspects of the cemetery's structure is unravelled largely as if it were independent of the other, which is manifestly not the case. Similarly, Martin-Kilcher's study of personal adornment at Münsingen accepts the Hodson chronology as a fixed framework, although logically each has implications for the other. Obviously the assumption that interpretive subgoals are independent and their conclusions additive is a highly convenient one. But equally obviously it can be highly misleading.

A few years ago I wrote, in Algol-60, a program called SOLCEM (Doran and Hodson, 1975, pp. 309-315). My purpose was to try to gain a better understanding of just what is involved in the cemetery interpretation problem. SOLCEM is capable of considering the bulk of the evidence in a (small) cemetery excavation record, and of generating an integrated overall interpretation.

Table II lists the main features of the SOLCEM program and Figure 3 is an outline flowchart. The program possesses a range of data analytic capabilities, embodied as a set of *data analysers*. Each data analyser is, roughly, a specialist procedure which may be selected and applied by an executive to a partially developed cemetery interpretation (initially the 'null' interpretation). A data analyser is applicable only when the categories of data and prior results which it requires for its action (typically including results returned by other analysers) are available, and when applied it adds its own conclusions to the specific interpretation of the cemetery being developed.

Since data analysers can and often will return sets of alternative conclusions, and since there is usually more than one analyser applicable to any particular partial interpretation, it is sometimes appropriate for SOLCEM to explore alternative interpretations in parallel and it typically does so. The program's executive chooses which partial interpretation of those currently being developed to

TABLE II

The cemetery data used by SOLCEM comprises:

- any prior artifact classification
- artifact descriptions
- artifact-grave associations
- grave locations
- age/sex determinations for graves
- any established chronological relationships between pairs of graves

The capabilities of the eight data analysers are:

- (i) classify artifacts from their descriptions
- (ii) detect associations between artifact classes and age/sex grave categories
- (iii) infer age/sex of indeterminate graves from artifact class associations
- (iv) analyse spatial organisation of cemetery
- (v) infer age/sex of indeterminate graves from spatial organisation
- (vi) infer chronological sequence of graves from their contents
- (vii) integrate findings of analyser (iv) with those of analyser (vi)
- (viii) reconstruct the evolution of the artifact types during the lifetime of the cemetery.

A complete SOLCEM cemetery interpretation involves:

- the chronological sequence of the graves
 - the spatial organisation of the cemetery
 - age/sex determinations for previously indeterminate graves
 - patterns in the assignment of grave-goods to graves
 - the evolution of the artifact types in antiquity
-

A summary of the capabilities of the SOLCEM cemetery interpretation program. The program generates an integrated interpretation of the cemetery taking into account the interactions between different types of evidence. But in any particular respect, for example the analysis of spatial organisation, its capabilities are very limited.

pursue next by reference to estimates of plausibility and internal consistency calculated by the data analysers themselves. The choice of which analyser actually to apply to the chosen interpretation, from those currently applicable, is made by reference to a simple priority list.

In this way the different analytic capabilities of SOLCEM are interwoven. Effectively the program searches through the space of partial interpretations of the given cemetery data looking for a complete interpretation which is sufficiently plausible to be finally accepted. In form a complete interpretation is a structural description of the cemetery's generation, covering the order of deposition of the graves, an estimate of the rules by which they were located and grave goods assigned to them, and an estimate of artifact type evolution. It roughly corresponds to a stochastic simulation model embodied in an auxiliary program, SIMCEM, used in experimentation to simulate the generation of cemeteries and hence to provide test data for SOLCEM.

Although SOLCEM copes satisfactorily with cemetery test data extending to

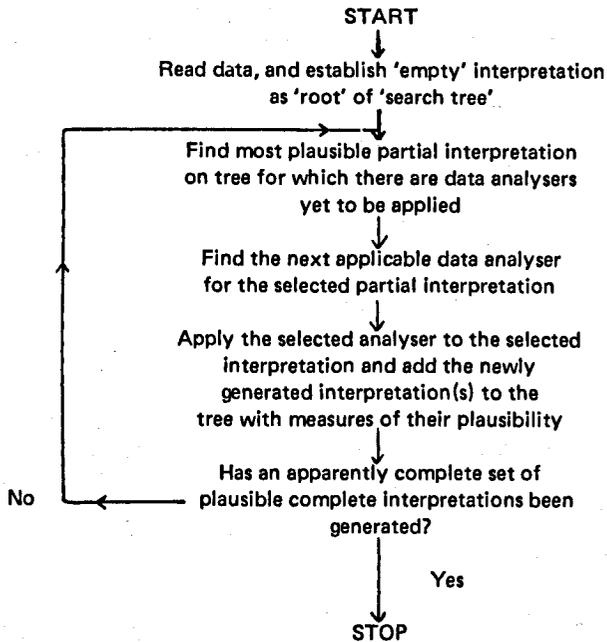


FIG. 3. Outline flowchart of the SOLCEM program.

about a dozen graves, it has never been applied to any significant body of real cemetery data. This is because it is impossible to justify the substantial effort needed to prepare an actual data set sufficiently large and detailed to be realistic, when it is clear that the program in its existing form would not produce sensible answers. Experience with SOLCEM has demonstrated that its performance is limited by the following two facts:

- (i) for any realistically sized cemetery, the total inference problem is computationally extremely demanding (which is, of course, why archaeologists do not attempt it, but always simplify), and
- (ii) alongside statistical techniques the problem requires the application of substantial specialist and general knowledge.

There are thus two directions in which improvements to SOLCEM might be sought. The problem of computational overload might be attacked by seeking an efficient solution strategy involving, no doubt, a trade-off between accuracy and effort. This would certainly raise deep issues of statistical estimation and general problem-solving theory. Secondly, an attempt might be made to render SOLCEM much more knowledgeable in a general sense. This is the possibility that I shall pursue here.

The point is that although SOLCEM is statistically quite intricate, it has all too little actual knowledge of cemeteries (see Table I) and would certainly be

blind to important aspects of real cemetery data. Worse, such knowledge as it has is irregularly scattered through the data analysers and the executive. On the other hand, efficient utilisation of knowledge should not only eliminate SOLCEM's "blind spots" but also go far to control the combinatorial explosion.

Initially one might suggest that there should be not eight, but eighty or even eight hundred data analysers. But this is a recipe for computational collapse unless coherent principles of organisation and knowledge representation are followed, and it is far from clear what these principles should be. It is natural to look for guidance to current artificial intelligence work on knowledge representation for recognition and inference.

ARTIFICIAL INTELLIGENCE TECHNIQUES FOR KNOWLEDGE REPRESENTATION

A variety of knowledge representation schemes have been explored in artificial intelligence work including predicate logic assertions, semantic networks, procedures, production systems, and frames. But there exists no hard theory comparing these representation schemes and capable of indicating which will be the most useful in any particular application. So all that we can do here is to look at the most prominent and recent of the schemes which have been proposed, and to try to extract concepts and methods which seem likely to be useful in designing a successor to SOLCEM.

Production systems have been used successfully in a number of artificial intelligence applications projects, notably in the Heuristic DENRAL project (Buchanan and Lederberg, 1971) in analytic organic chemistry, and in the MYCIN project (Davis, Buchanan, Shortliffe, 1975) in medical diagnosis. Further, the basic idea of a production is that it is a rule of the form "if A is encountered, then do X." In a general way, therefore, a system or list of productions already bears at least a passing resemblance to the SOLCEM set of data analysers. But this promising line of thought founders on the observation that all the SOLCEM data analysers are vastly more elaborate than the productions in any implemented production system. Short of breaking down such archaeologically fundamental activities as artifact typology or chronological seriation into sets of productions, which even if possible would hardly be sensible, production systems look too simple to be useful.

A working artificial intelligence system which is a major and effective elaboration of the production system philosophy is the HEARSAY II speech understanding system (Erman and Lesser, 1975). Committed to the 'hypothesize-and-test' view of inference, HEARSAY II is structured as a number of independent 'knowledge sources' each of which can act upon a uniform, multilevel, hypothesis structure, the 'blackboard', which is the system's understanding of the speech data put to it.

Each knowledge source is a kind of 'super-production' which, if and when its preconditions are satisfied in the blackboard, will be activated and will create, modify or delete hypotheses as appropriate in context. It is fundamental to the

philosophy of HEARSAY-II that knowledge sources should be quite independent, should communicate only through the blackboard and should in no sense 'call' one another with the rigidity of control that would imply. Predictably this stand leads to problems in achieving efficient and goal-directed behaviour, and HEARSAY-II has additional and not entirely elegant mechanisms directed to overcoming these problems.

The currently most debated proposal for knowledge representation is the suggestion that knowledge should be structured as 'frames'. The debate was initiated by Minsky (Minsky, 1975) who suggested that:

"Whenever one encounters a new situation. . . . he selects from memory a structure called a *frame*: a remembered framework to be adapted to fit reality by changing details as necessary.

A frame is a data-structure for representing a stereotyped situation, like being in a certain kind of living room, or going to a child's birthday party"

Minsky conceived of frame-selection as an essentially "top-down" process and of frames themselves as having a great deal of procedural knowledge attached to them concerned with, for example, what should be done if a selected frame turns out to be inappropriate.

Minsky's ideas have been developed and elaborated in recognition contexts by, among others, Fahlman (Fahlmann, 1973) and Kuipers (Kuipers, 1975). Bobrow and Winograd (Bobrow and Winograd, 1975) are implementing a new programming language (Knowledge Representation Language - KRL) which is largely based on frames ideas and which should therefore facilitate the development of functioning frames based systems. But in spite of this activity there is as yet nothing very concrete which might be adopted as the basis for a successor to SOLCEM. Currently frames ideas provide no more than useful guidelines.

Of these three approaches to knowledge representation-production systems, the HEARSAY II organisation, and frames-it is the HEARSAY II organisation which looks most useful from the archaeological standpoint and which I shall therefore pursue. But before I do so it is worth noting that all three approaches represent knowledge as a set of independent procedural or semi-procedural units. This facilitates identification and modification of the knowledge base. In both HEARSAY-II and frames systems the pattern of activation or utilisation of the knowledge units is fully data-driven, in contrast to the partially predetermined activation pattern of production systems and, indeed, of SOLCEM.

SOLCEM D - A PROJECTED SUCCESSOR TO SOLCEM

Let SOLCEM-D be the name of the proposed successor to SOLCEM. Following the HEARSAY-II system, I shall suppose that SOLCEM-D is structured as a set of knowledge bearing units, the successors to the SOLCEM data analysers, which independently act upon an evolving interpretation structure. However I

shall refer to the knowledge bearing units in SOLCEM-D as *recognition demons* rather than 'knowledge sources' since they have something in common with the demons proposed by Selfridge and Charniak. Further they correspond more to specific concepts (interpreting the word "concept" in a very broad sense), than to broad sources of knowledge. Admittedly this distinction is far from precise. I shall refer to the evolving interpretation structure simply as the *data-base*. The task of each recognition demon is continuously to scan the data base looking for examples of the concept which it embodies. When a demon discovers an example of its concept (in *its* judgment) then it generates and adds to the data-base an *instance*. An instance is, essentially, a description of the specific properties of the concept example found. For example, a demon might specialise in recognising isolated groups of graves. Then each instance that it generated would describe a particular grave group and its immediate properties, such as its location and the number of graves comprising it. Since instances once generated may be inspected and utilised by any other demon, the consequences of a single act of recognition can be propagated throughout the growing interpretation structure.

This is no more than a first sketch for SOLCEM-D. Any attempt to fill in details soon encounters difficulties familiar from the frames debate and from the HEARSAY II work. To see just what form these difficulties take in the context of cemetery interpretation we can usefully look at more specific examples.

ANDELFINGEN AND NEBRINGEN

The la Tène cemeteries at Andelfingen in Switzerland (Viollier, 1912) and at Nebringen in South Germany (Krämer, 1964) are smaller and simpler than that at Münsingen, but they are of the same general period and type. Figure 4 is a sketch plan of Andelfingen showing the locations of the graves together with their orientations and sex (male, female, child) where determined. It is clear by inspection that:

- the graves are almost uniformly aligned N-S or NNE-SSW
- the graves fall into two spatially separated groups, labelled A and B on the plan
- it is possible to divide the graves into four clusters (F1, F2, M and C on the plan) so that each cluster is made up entirely of graves of one sex (with one child's grave somewhat anomalously placed)

Figure 5 is a similar plan of the Nebringen cemetery. Here it is clear that

- the prevailing grave orientation is again NNE-SSW
- the graves fall into five spatially distinct groups (A,B,C,D,E).

These plans are somewhat misleading because they present the sexes of the graves as given. In fact sex determinations were made both from skeletal remains and from a study of the grave goods associated with the burials. Where these two

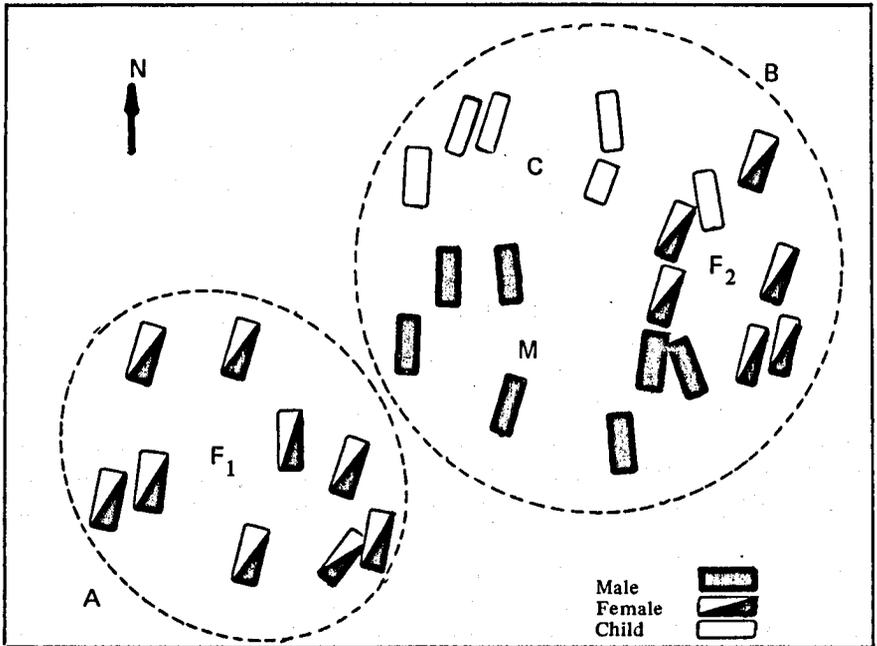


FIG. 4. Sketch plan of the Andelfingen cemetery indicating age/sex determinations and the grave groupings mentioned in the text. (redrawn from Schaaff, 1966, fig. 1).

sources of evidence contradicted one another, then a choice had to be made between them. This choice might reasonably have been influenced by the perceived overall structure of the cemetery, but neither of the authors of the studies concerned admit to having reasoned 'top-down' in this way.

At Nebringen the spatial groups were taken by Krämer to reflect in some way the family structure of the population. Each group contains at least one adult male and one adult female and no group contains more than three adults *provided* that spatial group A is treated as two adjacent family groups (see Figure 5). This is again an example of high-level organisation, once recognised, influencing the interpretation of lower-level detail.

At neither cemetery has it seemed reasonable to try to work out the chronological sequence of the graves. Nor have any significant social conclusions been drawn. For example, no interpretation has been offered for the prevailing NNE-SSW grave alignment. However Schaaff (Schaaff, 1966) has stressed the potential significance of the contrast between the organisation of the Andelfingen and the Nebringen cemeteries. He has suggested that the existence of two different burial customs among the Celts in the Early Iron Age may help in the determination of their later migration routes.

What we see, even from cemeteries as simple as those at Andelfingen and Nebringen, is that SOLCEM-D must be able not only to establish relevant structural features of a cemetery from imperfect evidence, but must also be able

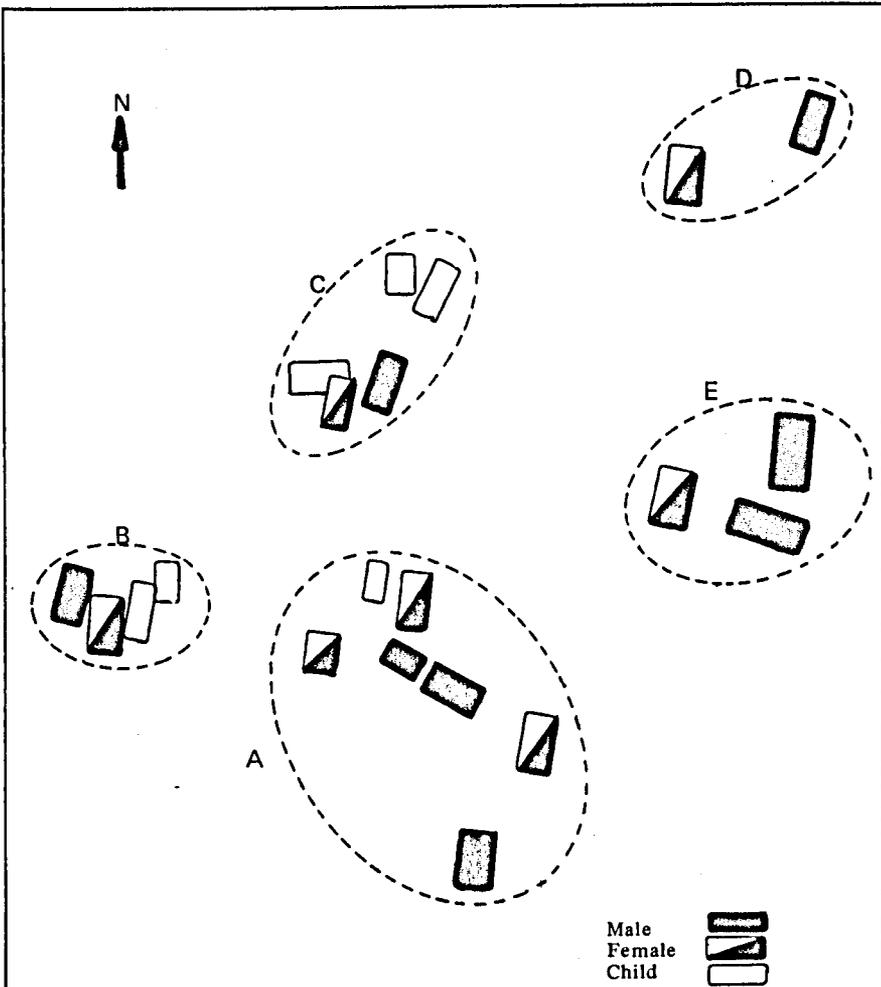


FIG. 5. Sketch plan of the Nebringen cemetery indicating age/sex determinations and the grave groupings mentioned in the text. (redrawn from Krämer, 1964, plan 2).

to reason “top-down” where appropriate. More generally, it must clearly be able to compare alternative interpretations opting for that with the greater internal consistency and plausibility. Notice that to a strictly limited extent each of these capabilities is already possessed by the existing SOLCEM program.

ISSUES IN THE DESIGN OF SOLCEM-D

Imagine now that SOLCEM-D, as sketched earlier, is equipped with a repertoire of recognition demons variously able to recognise spatially isolated groups of graves, to recognise that a particular grave group has a prevailing alignment, to recognise that a set of grave groups all have a common structure itself recog-

nisably a 'family' structure, and so on. Assuming that the data base is initialised with, say, the details of the Andelfingen cemetery (a set of grave instances and perhaps a cemetery instance) then it is plausible enough that the demons would in due course generate instances corresponding to the features of the cemetery noted earlier. But between this observation and an effective working system there stand (at least) four major and interrelated clusters of problems.

Firstly there is the problem of *mapping knowledge into demons*. By this I mean that the recognition and inference capabilities which the system is to have must somehow be allocated to a finite set of demons. Each demon must be able to 'see' its concept given only imperfect and fragmentary evidence. Each must compute the 'right' set of properties for the instances it generates.

The precise range of demons available in the system might seem immaterial provided that all the knowledge needed is embodied somewhere. On the other hand, it is a commonplace that using the right concepts can be crucial to understanding. The internal structure of demons must surely include multiple representations of the ideal concepts which they embody.

Related to the foregoing problem is that of how best to handle *instance properties*. As demons generate instances they associate with them their immediate properties. Exactly what form do these properties take, and what is the difference between a property and a concept with its own demon? Are relations properties or concepts or both? Where instances are hierarchially organised how much redundancy is permitted? For example, a common grave alignment throughout a cemetery implies that the same common alignment exists in any subset of the graves. How should such properties be handled?

The problem of *control* is fundamental. The idea that all the demons can be continuously active is computationally quite unrealistic. How, therefore, should the system be structured so that at any time the "right" demons are active and are investigating the "right" structures in the data base. Working top-down is a special case of this: it implies that a high-level demon has activated demons to look for the evidence it requires. Another facet of the control problem is that the generation of a new instance may well imply the recomputation of part, but not all, of one or more existing instances. How can this be handled efficiently?

A possible general strategy is to enable one demon to activate another direct, passing relatively complex 'instructions'. This seems to require that each demon must be knowledgeable about a range of other demons and their capabilities, and so runs counter to the idea that demons should be independent of one another. An alternative, essentially the PLANNER philosophy of invocation by pattern matching, is to enable demons to specify objectives without specifying which demons are to achieve them. Either way, control may or may not be automatically returned to the initiating demon. A further possibility, reminiscent of human thought processes, is that a quite simple network of "relevancy" links might be improved upon both the demons and their instances (possibly utilizing links already existing for other purposes) which would determine patterns of activation of a relatively unstructured but nonetheless effective variety.

Finally there is the problem of creating, evaluating and choosing between *alternative interpretation structures*. The problem ranges from deciding whether or not a particular grave should be included in a recognised grave group, to elaborating and then choosing between radically different complete cemetery interpretations each of which is consistent and reasonably plausible.

It is natural to suggest that each instance generated must have its own credibility rating. It also seems plausible that some explicit mechanism or inherent tendency must be devised whereby elaborate interpretation structures (which in this context might be better called /"hypothesis structures") persist whilst fragmentary structures are automatically discarded. One way of achieving such an effect might be by adding to the system's basic repertoire of demons others which recognise not patterns of data but patterns of inference. Human beings have, for example, the concept of a logical implication or of a *reductio ad absurdum* argument and are able to reason *about* particular arguments. A similar capability in SOLCEM-D might be the right way to introduce more sophisticated reasoning processes.

Few of the problems and suggestions sketched above are particularly new. Most have been considered in some detail either in the context of the knowledge representation schemes cited earlier or in connection with new A.I. programming languages such as PLASMA and KRL. Unfortunately only piecemeal progress has been made, and anything like a comprehensive understanding of all these issues still seems very far off.

KNOWLEDGE REPRESENTATION, DATA ANALYSIS, AND SIMULATION

Although sufficiently complex to motivate discussion of current artificial intelligence research issues, the Andelfingen and Nebringen cemeteries are relatively simple by archaeological standards. An archaeologist would not feel the need of statistical or computer aid when studying them, as he might well do when confronted with cemeteries of the size of Hallstatt or even Münsingen. This raises the question, mentioned in the introduction to this paper, of the exact relationship between the ideas of knowledge representation and use which artificial intelligence offers to archaeologists, and the ideas behind the statistical and computing techniques which archaeologists have already begun to use. These apparently very different, even hostile, clusters of ideas must somehow be brought into logically coherent relationship.

Social scientists generally, and archaeologists in particular, typically look to multivariate data analysis (cluster analysis, factor analysis, non-metric scaling) when they are confronted with a great deal of relatively homogeneous data which exhibits no immediately recognisable 'structure.' The recourse to multivariate analysis is partly an attempt to *uncover* meaningful structure, partly an attempt to *impose* structure, meaningful or not, in order to render the data tractable.

If this is an accurate characterisation of the use of multivariate analysis in archaeological work, then it seems that particular techniques can and should be

embodied as individual demons in the hypothetical SOLCEM-D system, but that they will make rather unusual demons—demons which are computationally very demanding (typically involving major iterative searches), which have a very wide range of application within the data base and which generate correspondingly non-specific instances. How important a role such demons would play in typical problems of archaeological inference is difficult to say. The use of multivariate techniques in archaeology is still highly experimental. Everything suggests that they are a kind of 'last resort' for use when knowledgeable interpretation of data is not possible.

The use by archaeologists of computer simulation as an investigatory tool is an even more recent development than their use of multivariate analysis. While discussing the SOLCEM program I mentioned my own use of a kind of stochastic simulation model to generate artificial but realistic cemetery excavation records. But much more significant than this is the use by American archaeologists of such models to infer from field and excavation data to social processes in antiquity (Doran and Hodson, 1975, chapter 11). The growing use of computer simulation is a methodological reflection of the current emphasis on process in prehistoric archaeology, with a corresponding rejection of the traditional, if usually unstated, view of the past as a series of static situations.

It is clear that SOLCEM-D (and SOLCEM for that matter) *can* construct an interpretation which embraces the time dimension. This was implicit in the earlier suggestion that a chronological seriation algorithm might be embodied in a recognition demon. As the interpretation is built, time relationships are built into it. But interpreting a particular body of data by fitting a computer simulation model to it is a little different. It corresponds to the validation stage in a standard simulation study. The simulation is run repeatedly, with successive modifications and varying parameter settings, until a convincing match to the data is achieved. This 'tuned' model then *is* the interpretation sought.

A simulation model is clearly itself a repository of knowledge, but in a much more structured and rigid representation than any of those mentioned earlier. What is the relationship between a computer simulation study and the kind of analysis which SOLCEM-D would perform? It is tempting merely to suggest that a simulation model should be assigned to a single demon as was proposed earlier for statistical techniques. But this is surely absurd. Rather, the event and process generators out of which the simulation model is constructed must individually be represented as demons. But if this is so, how are the demons concerned to be persuaded to cooperate within the loose control structure proposed for SOLCEM-D? There is no obvious answer, although it is not too implausible that a crude form of dynamic modelling might emerge as the temporal aspect of the requirement that newly generated constraints must be compatible with those already existing.

Dynamic modelling is of deeper archaeological significance than may at first appear. It is a commonplace to assert the weakness of archaeological theory in particular and of social theory in general. It is very likely that this weakness has

much to do with our inability to develop sufficiently complex process models and to infer from observations to particular instances of them. The drive towards a "systems approach" throughout the social sciences is an attempt to overcome this fundamental weakness, and is as prominent in the interpretation of pre-historic cemetery excavation data as anywhere else [see, for example, (Binford, 1971) and (Rathje, 1973)]. Thus if real insight can be achieved into the relationship between dynamic modelling and knowledgeable reasoning systems, then the impact on archaeology is likely to be as much at the conceptual and theoretical level as at the level of practical interpretation.

REFERENCES

- Binford, L.R. (1971) Mortuary practices: their study and their potential. *Memoirs of the Society for American Archaeology*, No. 25 pp. 6-29. (*American Antiquity* 36, 3 part 2, July 1971).
- Bobrow, D. and Winograd, T. (1975) Unpublished working papers.
- Buchanan, B.G. and Lederberg, J. (1971) The Heuristic DENDRAL program for explaining empirical data. *Information Processing 71* (proceedings of the IFIP Congress 1971; ed. C.V. Freiman) 1, pp. 179-88, Amsterdam: North Holland.
- Davis, R., Buchanan, B., Shortliffe, E. (1975) Production rules as a representation for a knowledge-based consultation program. *AIM 266*, Department of Computer Science, Stanford University.
- Doran, J.E. and Hodson, F.R. (1975) *Mathematics and Computers in Archaeology*. Edinburgh University Press.
- Erman, L.D., and Lesser, V.R. (1975) A multi-level organisation for problem solving using many, diverse cooperating sources of knowledge. *Internal research report*, Department of Computer Science, Carnegie-Mellon University.
- Fahlman, S. (1973) A hypothesis-frame system for recognition problems. *Working Paper 57*, A.I. Laboratory, M.I.T.
- Hodson, F.R. (1968) *The La Tène Cemetery at Münsingen-Rain*. Berne: Stämpfli.
- Hodson, F.R. (1975) *Hallstatt Inaugural lecture in the Institute of Archaeology, University of London*.
- Krämer, W. (1964) *Das keltische Gräberfeld von Nebringen (Kreis Böblingen)* (Reihe A, Heft 8, Veröffentlichungen des staatlichen Amtes für Denkmalpflege Stuttgart) Verlag Silberburg, Stuttgart.
- Kromer, K. (1959) *Das Gräberfeld von Hallstatt* Florence: Sansoni.
- Kuipers, B.J. (1975) A frame for frames: representing knowledge for recognition. *Representation and Understanding: Studies in Cognition Science* (eds. D. Bobrow and A. Collins) Academic Press.
- Martin-Kilcher, S. (1973) Zur Tracht- und Beigabensitte im keltischen Gräberfeld von Münsingen-Rain (Kt. Bern) *Anseiger für Schweizerische Archäologie und Kunstgeschichte*.
- Minsky, M. (1975) A framework for representing knowledge. *The Psychology of Computer Vision* (ed. P.H. Winston) McGraw-Hill.
- Moberg, C.A. (1972) Questions sur l'outillage mathématique d'une archéologie sociologique. *Les Methodes Mathematiques de l'Archéologie* C.A.D.A. - C.N.R.S. Marseille, pp. 230-248.
- Nilsson, N.J. (1971) *Problem-Solving Methods in Artificial Intelligence*. McGraw-Hill.
- Peebles, C.S. (1971) Moundville and surrounding sites: some structural considerations of mortuary practices II. *Memoirs of the Society for American Archaeology*, No. 25, pp. 68-91. (*American Antiquity* 36, 3 part 2, July 1971).

CASE STUDIES IN EMPIRICAL KNOWLEDGE

- Randsborg (1973) Wealth and social structure as reflected in bronze age burials—a quantitative approach. *The Explanation of Culture Change* (ed. C. Renfrew), pp. 565-570, London: Duckworth.
- Rathje, W.L. (1973) Models for mobile Maya: a variety of constraints. *The Explanation of Culture Change*. (ed. C. Renfrew), pp. 731-757, London: Duckworth.
- Saxe, A.A. (1971) Social dimensions of mortuary practices in a mesolithic population from Wadi Halfa, Sudan *Memoirs of the Society for American Archaeology*, No. 25, pp. 39-57, (*American Antiquity*, 36, 3 part 2, July 1971).
- Schaaff, V. (1966) Zur Belegung latènezeitlicher Friedhöfe der Schweiz. *Jahrbuch des Römisch-Germanischen Zentralmuseums, Mainz*, 13, pp. 49-59.
- Viollier, D. (1912) Andelfingen, Kt. Zürich *Anz. Schweiz, Altkde.* NF14.