28 The potentials of hybrid neural network models for archaeofaunal ageing and interpretation

Paul M. Gibson

28.1 INTRODUCTION

This paper outlines the limitations of expert systems in archaeology and other disciplines and suggests why neural networks can offer a superior alternative. It examines their potential use in archaeology and explains their functionality. Finally, the paper reports on the preliminary results of a research study into the use of hybrid neural network models for archaeofaunal ageing and interpretation and describes future work.

28.2 THE FUNCTIONAL DISABILITY AND SOCIAL ACCEPTABILITY OF RULE–BASED APPROACHES

A survey of the literature, not only in archaeology but in other disciplines too, reveals that the active research in expert systems has dwindled in recent years. This demise in expert systems popularity is due perhaps to the limited potential of expert systems in host disciplines. Before an alternative can be sought we must first appreciate the inadequacies of the rule–based approaches in general terms, and more specifically in archaeology. To do this, we can categorise the inadequacies of expert systems into two groups, functional disability and social acceptability.

28.2.1 Functional disability of expert systems

Rule–based systems contain chains of IF..THEN statements that can be used to infer a conclusion from a set of premises. The fact that knowledge is encoded as these chains of reasoning, along with other factors, leads to self–imposed limitations. In general, the issues that limit the functional abilities of expert systems can be outlined as follows:

1) It is often difficult to clearly identify an "expert" since individual's views and tactics are often diametrically opposed.
2) Using a single source of knowledge means that the expertise embodied by a system is, at best, no greater than that of the original expert.
3) Knowledge acquisition is an arduous task in that human experts often cannot coherently express their inference as a chain of rules. Indeed, the expert may not be using rules to deduce the solution to a problem.
4) Expert systems show little tolerance to unforeseen circumstances since their rule–base has been explicitly programmed to cope with predictable situations.
5) By accommodating rules that can cope with exceptional circumstances an expert system can grow into a framework containing tens of thousands of rules. This may lead to a system that is unmanageable, unmaintainable and even unpredictable.
6) Few safe strategies exist automatically to inject more rules to cope with new problems. In other words, it lacks the ability to learn from its failings without human intervention.
7) The choice of certainty factor values on rule–chains is usually an arbitrary decision. Consequently, the combination of such certainties along a line of inference may be unrepresentative of the situation.
8) The performance of rule–based systems is high within their area of expertise, but if the problem is even slightly out of the defined domain their performance becomes almost zero. This trait is known as the mesa effect.

The rule–based approach works well when the problem has a reasonably accessible set of rules.

We have clearly witnessed success in archaeology of such systems as (Brough & Parfitt 1984; Bishop & Thomas 1984), and more recently (Patel & Stutt 1988) with KIVA. However, when rule–based systems are applied to more ambitious cognitive tasks they become difficult to build, if not impossible.

Archaeology doesn't lend itself well to expert systems because it is a subjective and intuitive discipline. This lack of formalism is contrary to the requirements of expert system approaches. Doran states that archaeological theoretical ideas and interpretation aren't quantifiable and the domain of archaeology is "not fully analysed". This is reiterated by Vitali and Lagrange.

> «Archaeology may have its experts, but it is doubtful whether their knowledge in 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» (1988)

Doran states that expert systems in the interpretation domain fail because of the lack of socio–cultural theory. For an archaeologist to make an interpretation he/she must rely on everyday anthropological knowledge, which is difficult to foresee and program symbolically in a set of rules.

The overwhelming observation is that the human cognitive process can't easily be broken down into a set of ordered symbolic representations. Clearly, expert systems have only a limited scope in archaeology, the role of which has been studied by a few including (Wilcock 1985; Baker 1988; Huggett 1984; Huggett & Baker 1985; Doran 1988; Gardin 1988; Vitali & Lagrange 1988) justifying the criticisms made above.

28.2.2 Social acceptability

Whether it works or not, an expert system will face the problem of social acceptability. People tend to fear technology when it is professed to have qualities that humans have. Baker, Doran and others have outlined a number of likely fears and apprehensions concerning new computing techniques in archaeology.

The merit of expert systems is that they provide a basis for the development of new archaeological techniques and lead to the formulation of knowledge and theory which wouldn't otherwise be addressed.

28.3 THE NEURAL NETWORK MACHINE

The limitations of expert systems in host disciplines could explain why we are now witnessing in the field of computer science a revival of more tolerant and adaptive paradigms. Such methods have been derived from observations of nature and go some way to overcoming the pitfalls of rule–based and logical approaches.

One such discipline can be labelled alternatively as parallel distributed processing, connectionism or neural networks. Whilst these labels have their own distinctions their common principle is to model the cognitive phenomena of the human brain. In doing this they exhibit a number of brain–like characteristics such as learning, generalisation and abstraction.

The neural network approach is dramatically different to the concept of the Von Neumann machine which forms the basis of most modern day computers. A Von Neumann machine has an architecture based around the powerful and sophisticated CPU (Central Processing Unit), its memory and peripherals. The fundamental implication of this architecture is that it always needs a set of instructions to control the CPU and its components. This means that the computer programmer has to make a great effort to tell the machine **how** to get to a solution rather than concentrating on **what** is an acceptable solution to a problem.

The architecture of the Von Neumann machine can be classed as a "how to" model, where a set of instructions are given to the machine stating how a set of premises can lead to a solution. This task of instructing the computer has been made easier in recent years with the evolution of high–level languages. The principles of expert systems has been to run on these types of machines, hence the "how to" nature of their inference engine.

The architecture of the neural network machine can be classed as a "what is" model, where the programmer supplies an example set of problems and states what the ideal solution is. The machine itself creates the inference chain from the given problem to the solution without the programmer explicitly stating any rules.

For the purposes of this discussion we can view the neural network machine as a black box, as in Figure 28.1. The black box contains an internal representation of the problem domain which is constructed by the machine itself on the basis of example problems and solutions. The neural network has two modes of operation, the learning phase and the evaluation phase. In the learning phase, the machine is repeatedly supplied with a training set of problems from the domain. A learning algorithms is applied to the network which merely adapts the internal representation until the desired output pattern or pattern of excitation is achieved.

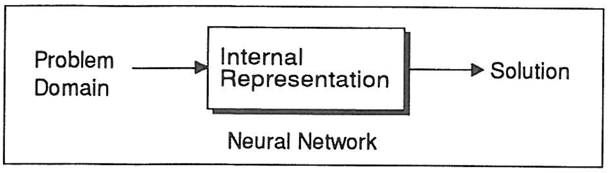

Figure 28.1: The black box analogy

The choice of learning algorithm depends on the architecture of the network and the nature of the problem. Broadly, we can class these algorithms as either supervised or unsupervised. A supervised learning algorithm needs the solution to a problem and uses this to determine the desired excitation pattern. One such algorithm is known as the back–propagation algorithm (Rulmehart *et al.* 1986). The training input example is passed to the network (propagated), the resulting output is compared to the target output and the difference is used to adjust the internal representation. The algorithm is reapplied until all output excitations are the same as the training target outputs. An unsupervised learning algorithm needs no training target outputs, instead it takes the common features of the training input and classifies them into groups. One such algorithm is known as the Kohenen self–organising map (Kohenen 1974).

Once the network has been trained to respond to input problems, unseen problems can be presented for evaluation.

Neural networks eliminate the processor/ memory concept and the need for explicit symbolic instruction sets. Instead, they provide a network of computing modules in which different patterns of excitation are observed as a function of the interconnections between modules. Processing and memory, no longer exist as one homogeneous unit, but instead are distributed uniformly through the network architecture.

In addition to learning, connectionist approaches also provide secondary additional benefits such as fault tolerance, default assignment, abstraction and spontaneous generalisation.

A number of good books describe the functionality of neural networks one of which is Wasserman (1989).

28.4 WHY NEURAL NETWORKS ARE SUPERIOR TO EXPERT SYSTEMS?

Neural networks overcome the major functional disabilities outlined in section 28.2.1. There is no longer a need to sit down with an expert in order explicitly to define his/her knowledge as a set of rules. Neural networks have the capacity to formulate their own representations of the apparent cognitive processes of the human from a set of representative problems and solutions. The responsibility for rule production is no longer that of the designer, but instead that of the neural network. This has a further advantage in that the designer need only have knowledge of how best to present input and output to the network.

Since neural networks generalise from examples they can show adaptability to varying circumstances and gracefully degrade in the face of ambiguous input. They work well with incomplete data and their performance at the edge of their knowledge is far superior to that of the expert system. This generalisation from the training set makes them less likely to suffer from the mesa effect.

Furthermore, they can learn from their failure by simply re–adjusting their internal representation without having to grow to accommodate new rules. This means that once the architecture of a neural network has been established it remains static.

However, it is worth stating three points about the performance of neural networks:

1) Neural networks have long learning times which are computationally expensive. This is due to the small adjustments made to the internal representation after each example presentation.
2) It is important to choose representative examples to ensure that the internal representation is stable across all training sets.
3) Their internal performance is not well–understood. It has been observed that they can respond to desired solutions but it is difficult to say how they got to a solution. Consequently, neural networks lack the ability to explain their reasoning.

28.5 HYBRID MODELS

Hybrid models combine the feature of certain artificial intelligence tools into one homogeneous architecture. By constructing such mutualistic architectures the advantages of the tools used are compounded.

A hybrid expert network is a combination of neural networks and expert systems. More often than not, the expert system drives the model performing the more basic and simple tasks, when the problem becomes too difficult it pulls in a neural network to solve a sub–problem. In this way, the advantages of both are gained.

Another hybrid model uses Genetic Algorithms (Holland 1975) which are evolutionary strings of information. The neural network drives the evolution of the genetic algorithms to the desired solution. An example of this type of architecture has been used for musical composition (Gibson & Byrne 1991)

Finally, a classifier system combines expert systems with genetic algorithms resulting in a rule–based system that can evolve to cope with new circumstances. An explanation of genetic algorithms and classifier systems can be found in (Goldberg 1989).

28.6 APPLICATIONS OF NEURAL NETWORKS AND THEIR POTENTIAL USAGE IN ARCHAEOLOGY

Neural networks have been applied in many application areas across a diverse set of disciplines. Foremost of these application areas is in image processing and pattern recognition, where they are commonly used as feature detectors. Networks have been taught to recognise written characters, analyse finger–prints, identify bombing targets from aerial photographs and other such image processing tasks. Other applications areas include monitoring and control situations such as nuclear reactor monitoring, robot control, musical composition and ECG monitoring to name but a few.

As yet archaeology has remained neural network free. However, neural network models may have an impact on three areas of archaeology: image processing, simulation and interpretation.

Image processing techniques are now common in archaeology and neural networks provide an easy solution to some vision problems. Neural networks could be used for pottery analysis, analysis of aerial photographs, estimating missing measurements on animal bones and perhaps automatic classification of artifacts from photographs. Moreover, neural networks could be used for the analysis of micro–wear on flints. Claxton has also suggested that neural networks could be used for GIS applications, by extracting features from maps (Claxton, personal communication).

Neural networks could play an important part in the interpretation of archaeological data, in that the user doesn't need to explain the reasoning behind an interpretation. The network simply acts as a bridge from data to interpretation. Neural networks should be able to cope with the informal structure of archaeological knowledge provided the networks are structured correctly.

Clearly, neural network models could be beneficial to archaeology.

28.7 THE RESEARCH STUDY

This section deals with the research study into the potential use of neural network models in archaeology, and more particularly for the ageing and interpretation of archaeofaunal material. The task of ageing and interpretation is initially described followed by the preliminary success and failure of using hybrid neural networks models. Finally, the future work is outlined.

28.7.1 The archaeological problem/task

The age at death of a set of animals from a site can reveal a great deal about the economy and exploitation of the livestock on that site. One approach to age–estimation is based on the eruption and wear stages of an animal's teeth. As a mammal matures the dentition progresses through a set of distinguishable stages.

A young mammal will initially grow a set of deciduous teeth which are progressively replaced by permanent teeth in a predictable sequence and at a certain stage in its life. Once the mammal has its set of permanent teeth the amount of attrition (wear) on each tooth's occlusal (biting) surface can be used to distinguish between individuals of different ages.

A number of important papers have been published in recent years that establish criteria for the classification of a given mandible to a particular age group on the basis of this tooth eruption and occlusal attrition.

A main contributor to this method are Deniz and Payne (1982) who have studied the wear stages of Anatolian Sheep and goats. A more wide spread study has been carried out by Grant (1978, 1982) who has studied the common ungulates of pig, sheep/goat and cattle.

The two methods concentrate on the fourth deciduous pre–molar (DP4), the fourth permanent pre–molar (P4) and the three permanent molars (M1, M2 and M3) since these teeth survive better in antiquity. Incisors are poorly recovered from excavations.

The methods have diagrams for each tooth that represents the idealised wear stages. A tooth in a given mandible is compared against the diagram and the diagram that best matches with the tooth suggests the wear stage. Each tooth is treated in the same way, and the wear stage of each tooth is used to categorise the mandibles into a relative age group.

These methods can't give an exact age because the wear of the tooth is dependent on the coarseness of the foliage and the amount of sand in the soil. However, the fauna within a site can be compared for relative ageing.

Because these methods use the judgment of the human eye to compare the diagram against the actual tooth, they tend to be subjective and not very accurate. Furthermore, the diagrams are general representations of the wear stage and exact matches can't be performed. More often than not, a tooth's wear may fall between two stages and it is left up to the bone specialist to make an arbitrary decision as to which one to pick. Even worse, incomplete and missing teeth complicate the proceedings. Therefore, the estimation of age from these methods must be viewed with some caution.

Once the age of each mandible has been estimated, the number of individuals within each age group can be used to determine the kill–off strategy of the livestock, which in turn suggests the type of economy. For example, if the livestock is kept solely for wool production, we would see that there would be little infant mortality, whereas if the livestock is kept for meat production we may see more young animals being killed. This method of interpretation has been discussed in detail by (Payne 1973) and an attempt to simulate it has been undertaken by (Cribb 1984)

It is likely that the livestock has been used for a variety of reasons and again simply looking at the resulting graph can lead to a subjective interpretation.

28.7.2 The aim and method of the study

The overall objective of this research study is to devise a computer system that can estimate the ages of archaeofaunal mandibles and form an economic interpretation using the principles outlined above. The aim is to analyse, aid and eventually improve the current techniques for age estimation and interpretation. Such a system would eliminate the subjectivity and standardise the procedure across data sets. Furthermore, it would release the archaeologist from the labour and tedium of the task.

The working title of this research study is «Applied Hybrid Neural Network Modelling and Artificial Intelligence techniques to analyse, aid and improve archaeofaunal ageing and interpretation».

Before such a system can be built a number of feasibility studies have to be carried out to prove that such an ambitious task can be achieved. In the early stages of this research study, a number of general tasks were identified, and a critical research path devised:

- Selection of the adequate training set
- Data analysis of the supporting database
- Analysis of suitable digitisation approaches
- Feasibility study of tooth identification
- Feasibility study of attritional estimation
- Feasibility study of interpretation
- Selection of appropriate hardware and software

To reduce the complexity of the overall problem, it was decided initially to concentrate on sheep/goat mandibles for a number of reasons. Sheep/goats tend to be commonly found on archaeological excavations, and there is a panoply of bone reports to use for interpretation case studies. Furthermore, analysis of sheep dentition and interpretation has been studied in detail. Moreover, pigs teeth are far more complex in structure than those of sheep/goat and cattle.

Coy (1978) and others have suggested that a modern–day set of assemblages with known age at death should be used for comparison with archaeofaunal data. Consequently, Dr. Dobney of the Environmental Archaeology Unit at York University kindly lent a set of modern sheep mandibles with varying stages of attrition and a number of tooth abnormalities. Furthermore, Dr. O'Connor of Bradford University lent some archaeofaunal mandibles for system testing.

28.7.3 Hardware and software

The platform chosen for the initial set of feasibility studies was the standalone PC (80486 processor). The PC tends to be the more portable and readily available computer system, making it easier to use on site or in the laboratory.

C^{++} was chosen because of its object–orientated features, which provide an adequate framework for the generation of multiple neural network models. For speed, assembly language has also been selected.

28.7.4 Analysis of suitable digitisation approaches

The first question to consider is how "best to present" the mandibles to a computer for analysis. A number of approaches were considered including ultrasonic scanning, spot ranging, strip scanning using laser light, still camera and video camera. Considering that the wear analysis requires recognition of different patterns the video camera was chosen. A video digitiser was acquired that had a frame buffer of 640 × 480 pixels with 16 intensities. This approach allows the application of existing computer vision techniques and principles.

Experiments were then carried out to determine the most appropriate way of presenting the mandibles to the video camera. Five physical factors had to be overcome

- lighting to reduce shadowing and reflection
- orientation of mandibles
- orientation of camera
- noise and background filtration
- normalisation across data sets

From the results of the experiments a box was constructed with a black cloth interior and two light sources mounted in the top corners. This reduced the effects of shadowing and reflection.

Experiments suggested that a camera placed towards the intersection of the occlusal and buccal surfaces of the tooth yield an information rich image. Occlucal views of a tooth, as illustrated by Grant and Deniz and Payne, reduced emphasis of prominent features, such as cusps.

A soft clamp covered in black cloth has been used to hold all mandibles in the required orientation.

Software written in C^{++} and assembly language had to be devised to perform all image capturing functions. Noise is average out by using a Gaussian filter on captured images.

28.7.5 Feasibility study into the potentials of neural networks for tooth separation and identification

From the tooth row image the software must identify which teeth are present as an indictor for the application of wear stage models. This means identifying the borders of the tooth and mapping their outline to one of the tooth descriptors (e.g. DP4, M3, etc).

This task is very complex because the shape of the teeth are not that regular across mandible samples. Furthermore, broken teeth and deformed teeth vastly complicate any effort to identify teeth descriptors. Even a simple rotation of the tooth in the view frame completely changes its appearance as far as the machine is concerned. Therefore, the development of a tooth identification algorithm must be tolerant to all such circumstances.

To simplify the feasibility study and computing effort it was decided that any subsequent analysis would be performed solely on complete M2 teeth. An algorithm was created that highlighted the approximate border of the teeth on the basis of background information, as in Figure 28.2. A neural network architecture based on the back propagation model was constructed and trained to distinguish between M2 and non–M2 silhouettes. The EAU sample set was used for training and the archaeological set used for testing. The network responded well to the test data.

Although, due to the limited training and test data, the experiment doesn't prove that neural networks can perform under all circumstances, it shows that neural networks have a potential for tooth recognition. It would have been difficult to use an expert system under such circumstances, because the data are just a series of pixel points. Effort would have had to be made to get the raw data into some descriptive representation before any rules could be applied.

28.7.6 Feasibility study into the potentials of hybrid neural network models for archaeofaunal ageing

The main emphasis of this research study has been to assess the feasibility of using hybrid neural network models for archaeofaunal age estimation. Not only was an investigation carried out to show that hybrid models could perform archaeofaunal ageing, but effort was made into investigating the use of rule–based systems to see how they compared.

The ultimate product from the study was to gain a wear stage value from the image representation of the tooth. Any image processing task of this magnitude is complex, since we are trying to make sense out of a series of pixels. Whereas the human eye readily perceives objects in an image, a machine has to be told how to cluster the pixels into objects.

Four separate attempts were made to devise a system that could analyse the attrition of M2 teeth.

28.7.6.1 Approach 1 : using traditional image processing and artificial intelligence tools

The initial approach was to follow the ideas of Marr's theory of vision (Marr 1982) and produce a representative sketch of the tooth that could then be interpreted by an expert system.

Many techniques were employed in an attempt to highlight all the features of the tooth. However, the colouration of enamel and dentine prevented the individual identification of these features using such non–semantic techniques.

The failure even to identify features within a tooth meant that no comparative assessment of traditional and neural network approaches could be undertaken.

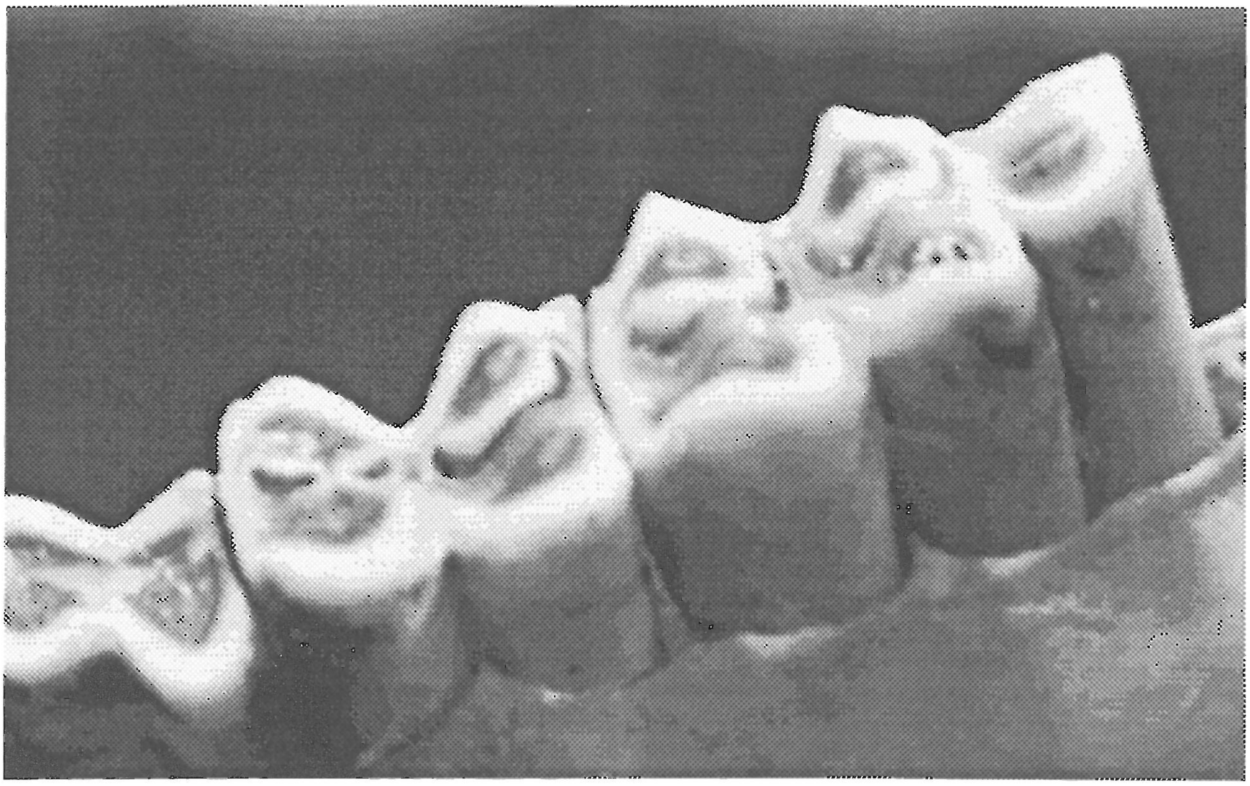

Figure 28.2: Example image, showing M1 to M3 under optimum lighting and camera alignment.

28.7.6.2 Approach 2 : brute–force approach using singular neural network models

The second approach involved building a large neural network model based on the back–propagation algorithm that covered the whole image array of 120 × 80 pixels (9600 neurons in the input layer). Each output from the network represented all possible wear stages. The network was taught the stages of wear of the EAU training set. For example, for a tooth at wear stage A the network was taught to produce a 1 output from the first neuron in the output layer and 0 in the rest.

The archaeofaunal set was then used to test the network. This approach failed, but justifiably so. The number of training examples given to it were limited and the volume of input neurons was too extensive. This approach could be made to work using perhaps 1,000 training examples.

A self–organising map was also constructed with the hope that it would find the common wear stage classifications within the training set. However, due to the limited set of training data this was not a valid study.

28.7.6.3 Approach 3: using cooperative neural networks

The third approach involved building seven neural network models that were cooperatively trained to identify features of the tooth. Each network could identify a particular component of the tooth, such as dentine, enamel and infundibulum.

After training, the networks cooperatively searched the image frame for their particular feature and passed the results on to another network that could establish a wear stage.

This approach performed well on the set of teeth it was trained on, but degraded when presented with the archaeofaunal mandibles.

28.7.6.4 Approach 4 : hybrid neural network models for region attritional assessment.

Up until this point the emphasis was to highlight features of the tooth as a means to establishing a wear stage. However, it was observed from Grant's sketches and from real mandibles that certain individual characteristics of a tooth relate to tooth wear, meaning that the features of the tooth, such as dentine and enamel, were not that important.

Each Grant wear stage has a new region on the tooth that comes into wear. By identifying the regions on a tooth we can determine which of them is in wear and consequently ascertain a wear stage. Fifteen regions were identified as important to wear assessment. The conceptual schema of the computer model is shown in Figure 28.3.

The raw image is simplified by using an algorithm based on Robert's and Sobel's operators. After which, an expert system is used to identify the likely wear regions and pull in the appropriate neural network. A second expert system combines the results from the networks to produce a wear stage value. Note that the simple tasks are performed by the expert system and the more complex and intuitive tasks are performed by individual networks.

To prove that the approach would work one region was manually selected for the feasibility study. A neural network was built to determine whether that region was in wear or not. The EAU set was again used for training and the archaeological set used for testing.

Even from the limited training set the network successfully identified whether the region was in wear or not from unseen teeth. This preliminary result shows that neural networks can be used for archaeofaunal age estimation.

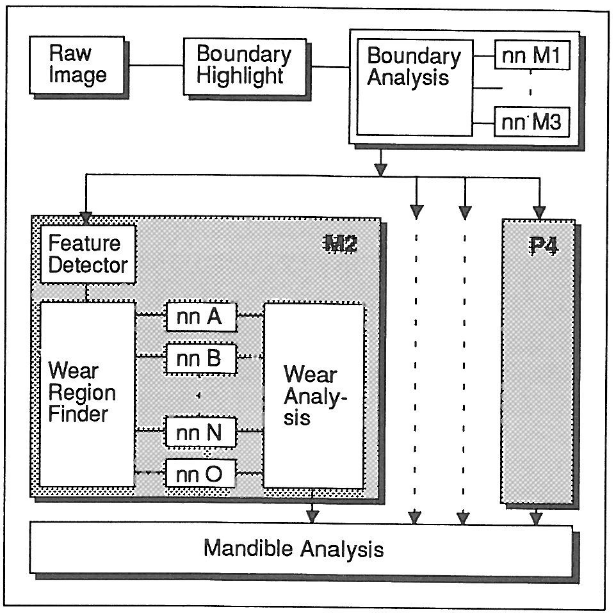

Figure 28.3 : The conceptual schema

28.7.7 Future work

Further work has to be carried out to prove that the whole mechanism of the fourth approach will produce wear stages. The other fourteen network models have to be trained and tested. An algorithm has to be devised to locate the initial regions in the tooth image. Moreover, the expert network at the end of the chain has to be developed to provide the ultimate wear value. More tolerance training is required to cope with broken and incomplete teeth. Thereafter, other teeth in the mandible have to be considered. After which the resulting wear stages and states of eruption must be assessed by another network to produce a final age– group for the mandible.

A study must be carried out to determine the feasibility of using neural networks for the development of an economic model from the age at death data.

Eventually, the components of the feasibility studies can be combined to assess mandibles of varying species with accuracy.

Once the model has been completed it will be easy to take its structure and apply it to other species and possibly other areas such as micro–wear.

28.8 CONCLUSION

It has been seen that neural networks have a greater potential in archaeology than expert systems. The research study shows that combining the features of expert systems and neural networks provide a robust architecture.

Future work has to be carried out to analyse the full potential of neural networks for archaeofaunal ageing and interpretation, but the initial results are very promising.

Acknowledgements

The author is grateful to Dr. Keith Dobney of the Environment Archaeology Unit (EAU), University of York for the loan of a sample set of sheep mandibles. The help and enthusiasm shown by the author's supervisors Dr. Julian Richards, University of York and Dr. T O'Connor, University of Bradford is gratefully acknowledged. The initial advice given be Dr. J. Austin, Computer Science Dept., University of York, is much appreciated. And lastly, the author is grateful to Nick Willson for his ideas and advise on using traditional image processing techniques. The research is partly funded by the Nestle Rowntree Division of Nestle UK.

References

Baker, K.G.
1988 Towards an archaeological methodology for Expert Systems, in C.L.N Ruggles & S.P.Q Rahtz (eds.) Computer and Quantitative Methods in Archaeology, BAR International Series 393, Oxfoord, pp 229–236.

Bishop, M.C. & Thomas, J.
1984 BEAKER, an expert system for the BBC Micro, CAA Conference Proceedings 1984, pp.56–62.

Brough, D.R. & Parfitt
1984 An Expert System for the ageing of a domestic animal, CAA Conference Proceedings 1984, pp. 49–55.

Coy, J.
1978 Comparative collections for zooarchaeology, in Brothwell, Thomas & Clutton–Brock (eds.) *Research problems in zooarchaeology*, Institute of archaeology, Occasional publication No. 3. pp. 143–145

Cribb R.
1984 Computer simulation of herding systems as an interpretive and heuristic device in the study of kill–off strategies. in J. Clutton–Brock & C. Grigson (eds.) *Animals and Archaeology 3: Early Herders and their Flocks*. BAR International series 202. pp. 169–170.

Deniz, E. & S. Payne
1982 Eruption and wear in the mandibular dentition as a guide to ageing Turkish Angora Goats, in B. Wilson, C. Grigson & Sebastian Payne (eds.) *Ageing and Sexing Animal Bones from Archaeological Sites*. BAR British Series 109, pp. 155–205.

Doran, J.
1988 Expert Systems and Archaeology: what lies ahead ?. in C.L.N. Ruggles & S.P.Q. Rahtz (eds.) Computer and Quantitative Methods in Archaeology, BAR International Series 393, pp 237–241.

Gardin, J.C.
1988 *Artificial Intelligence and Expert Systems, case studies in the knowledge domain of archaeology*. Halsted Press, New York.

Gibson, P.M. & J.A. Byrne
1991 NEUROGEN: Musical composition using Genetic Algorithms and Cooperating Neural Networks. *IEE Second International Conference on Artificial Neural Networks, Institution of Electrical Engineers*, Publication No. 349, pp. 309–313.

Goldberg, D.E.
1989 *Genetic Algorithms in search and optimization and machine learning*. Addison–Wesley Publishing Company, Inc.

Grant, A.
1978 Variation in Dental Attrition in Mammals and its relevance to age estimation. in D.R. Brothwell, K.D. Thomas & Juliet Clutton–Block (ed.) *Research Problems in Zooarchaeology*, Institute of Archaeology. Occasional Publication No. 3, University of London pp.103–106.
1982 The use of tooth wear as a guide to the age of domestic Ungulates. in B. Wilson, C. Grigson & Sebastian Payne (eds.) *Ageing and Sexing Animal Bones from Archaeological Sites*. BAR British Series 109. pp. 123–142.

Holland, J.H.
1985 *Adaptation in natural and artificial systems*. Ann Arbor, The university of Michigan Press.

Huggett, J.
1985 Expert Systems in Archaeology, in Cooper, M.A., & Richards J.D. (eds.) *Current Issues in Archaeological Computing*. BAR International Series 271, pp. 123–142.

Huggett, J & K. Baker
1985 The computerised Archaeologist: The Development of Expert Systems. *Science and Archaeology* 27: 3–7

Kohenen, T.
1974 An adaptive associative memory principle. *IEEE Transactions C–23*, pp. 444–445.

Marr, D.
1982 *Vision*. Freeman, San Francisco.

Patel, J. & A. Stutt
1988 *KIVA: an archaeological interpreter*. Technical Report 35, Human Cognition Research Laboratory, Open University, Milton Keynes.

Payne, S.
1973 Kill–off patterns in sheep and goats: the mandibles from Asvan Kale. *Anatolian Studies* 23:281–303.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J.
1986 Learning internal representation by error propagation. in D.E. Rumelhart, & J.L. McClelland, *Parallel Distributed Processing, Explorations in the Micro–structure of Cognition Volume I: Foundations'*, MIT Press, Bradford Books. pp. 318–364

Rumelhart, D., & McClelland, J.
1986 *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol I*, MIT Press, Cambridge, MA.

Vitali, V & M.S. Lagrange
1988 VANDAL: an Expert System for the Provenance Determination of Archaeological Ceramics based on INAA Data, in S.P.Q Rahtz (ed) *CAA Conference Proceedings* 1988, pp. 369–375.

Wasserman, P.D.
1989 *Neural Computing — Theory and Practice*. Van Nostrand Reinhold.

Wilcock, J.
1985 A Review of Expert Systems: Their shortcomings and Possible Applications in Archaeology. in E. Webb (ed.) *Computer Applications in Archaeology 1985*, Institute of Archaeology, University of London, London pp. 139–144,

Author's address
Paul M. Gibson
35 Otterwood Lane
Foxwood Hill, Acomb
GB–YO2 3JP York