

Beyond Classification: Automatic Learning in Archaeology and Cultural Studies

Juan A. Barceló

Universitat Autònoma de Barcelona
jbarcelo@seneca.uab.es

KNOWLEDGE ACQUISITION AS LEARNING

Any dictionary definition suggests that learning involves the acquisition of knowledge or behaviour. Chris Thornton, in a thought provoking essay suggests that learning tasks must be presented in the form of example association pairs (Thornton 2000). Consequently, he defines all learning tasks in terms of *mapping*. A learner agent can always read this mapping in two ways: as saying what actions should be produced in a given situation or as *predicting* which actions should be produced in a given situation. The idea of prediction comes from the fact that a task will be considered to have been successfully learned if we are able to produce correct associations for test items in the majority of cases.

Prediction tasks are usually named *supervised learning tasks*, on the grounds that the examples are like information given to the learner by a teacher or supervisor. The thing that is learned is generally termed the *target function*, and the inputs and associated outputs are treated as the arguments and values (respectively) of an unknown function. Concept learning problems have the same form, except that target outputs are either “yes” or “no” (or “true”=1 and “false”=0). The learning is then conceptualised along computational lines.

CLUSTERING AND CLASSIFICATION

Learning and *classification* are seen as procedures for knowledge acquisition. A *classification problem* is just like the concept learning problem except for the fact that we now have a number of output targets that are the labels of classes. The cases illustrate what sort of object belongs in which class.

The goal in a classification problem is to develop an algorithm which will assign any artefact, represented by a vector $\mathbf{x}$, to one of c classes (chronology, function, origin, etc). The problem is to find the best mapping from the input patterns (descriptive features) to the desired response (classes). Some finite or infinite set of patterns (binary or real valued vectors) is to be partitioned into classes, and a particular problem is specified by a set of selected training patterns, which are given together with their corresponding class names, and the goal is to classify *all* patterns as correctly as possible.

The problem is of dividing the set of possible input vectors into two sets, one for which 1 is the output, and the other for which 0 is the output (“Yes”/“No”). The task of learning consists in separating the input vectors, which are assigned to the output 1, from the input vectors which are assigned to the value 0 (Zeidenberg 1990, p. 103). When a few examples may be incorrectly classified, but the separation of the patterns is much better than if no decision rule was used, the classes are said to be **linearly separable**. In the other case, when it seems there is no clear decision rule to separate examples from counterexamples, we say that classes are not separable. This is a classical non-linear problem.

The purpose of the classification problem is to estimate the probability of membership of the case in each class. We are forced to treat the measured variables as random quantities, and to accept that perfect classification of new examples may not always be possible. Instead we could aim to build a classifier which has the smallest probability of making a mistake. The objective is to build a model with significant predictive power. It is not enough just to find which relationships are statistically significant.

It is important to understand the difference between clustering and classification. Because of the differences between (explicit) commonalities and (implicit) relationships, classification as a learning task with a predictive form cannot be reduced to clustering.

Clustering is the process of grouping input samples in similarity classes. From a theoretical viewpoint, clustering is a form of non parametric density estimation. In the absence of a supervised learning approach, the best we can do for categorization is to use information about the input data distribution to separate inputs into groups that share the same region in data space. The basic idea of clustering is to seek regions of high sample density –*data clusters*– and represent their centres in the network’s output layer. Most clustering mechanisms are based on some distance measure. Each object is represented as an ordered set (vector) of features. “Similar” objects are those that have nearly the same values for different features. Thus, one would like to group samples so as to minimize intra-cluster distances while

maximizing inter-cluster distances, subject to the constraints on the number of clusters that can be formed. *Vector quantization* is the process of dividing up space into several connected regions, using spatial neighbourhood as an analogue of similarity. Every point in the input space belongs to one of these regions, and is mapped to the corresponding nearest vector. For example, the attributes for "object A" are mapped to a particular output unit or region, such that it yields the highest result value and is associated with that object, while the attributes for "object B" etc. are mapped to different regions.

Classification is a predictive task, whose result is new knowledge added to a set of previously unknown data. The new knowledge added is just a label, the fact that the case belongs to a class. But in a classification system, a class is much more than a set of similar things. Because classify is not to cluster similar data, the fact of assigning a label to a case means to *explain* that case, by linking it to a concept, for which it is an *instance*.

A CASE STUDY: ARCHAEOLOGY

Archaeologists have been erroneously using clustering as a learning mechanism. The underlying idea seems to be that everything which is similar, was produced, distributed and used in the same way. For instance, some tools have different use-wear texture, because they have been used to cut different materials; some vases have different shapes because they have been produced in different ways; graves have different compositions, because social objects circulated unequally between members of a society and were accumulated differentially by elites...

But we have already seen that clustering is not classification, and we cannot learn knowledge nor predict concept assignments on the basis only of perceived similarities. We need a method for labelling perceptual information (shape, size, texture, composition, location) via some external criterion. Therefore, in a supervised-learning task, the identification of classes (solutions to the problem) is known beforehand. In clustering problems, on the other hand, all that is available is a set of samples and similarity measures that can be derived from the sample descriptions. Since clustering is determined only by the natural grouping of the input data, and it is unsupervised, it cannot be used directly for predicting, and the archaeological problem cannot be solved. However, in many practical applications, data from each class tends to be dense and there is a natural separation between classes. In such cases, clustering can be a preprocessor for classification.

NEURAL NETWORKS FOR RELATIONAL LEARNING

The power of neural computation comes from the massive interconnection among the artificial neurons, which share the load of the overall processing task, and from the adaptive nature of the parameters (weights) that interconnect the neurons. Artificial Neural Networks should be described as distributed, adaptive, generally nonlinear learning machines built from many different individual processing elements or "neurons".

For a computing system to be called "Artificial neural network" or "connectionist system", it is necessary to have a labelled directed graph structure where nodes perform some simple computations (Zeidenberg 1990, Mehrotra et al. 1997). A "directed graph" consists of a set of nodes (vertices) and a set of connections (edges/links/arcs) connecting pairs of nodes. It is the pattern of interconnections which is represented mathematically as a weighted, directed graph. The vertices or nodes represent basic computing elements (neurons), the links or edges represent the connections between elements, the weights represent the strengths of these connections, and the directions establish the flow of information and more specifically define inputs and outputs of nodes and of the network (Fine 1999). Each node's activation is based on the activations of the nodes that have connections directed at it, and the weights on those connections. To complete specification of the network, we need to declare how the nodes process information arriving at the incoming links and disseminate the information on the outgoing links. The nodes of the network are either input variables, computational elements, or output variables (Zeidenberg 1990, Mehrotra et al. 1997, Fine 1999).

Essentially the idea is that nodes (artificial neurons) communicate by sending numbers along the lines that connect them (Rumelhart 1989). In a neural network, each node performs some simple computations, and each connection conveys a signal from one node to another, labelled by a number called "weight" or "connection strength". That means that without variable weights amplifying and/or diminishing signals through connections, there is not any neural network. The way nodes are connected determines how computations proceed.

Getting the networks to do something is a matter of giving them inputs, and letting the information travel through the topology by simulating each artificial neuron. The role of neural networks is to provide general parameterised non-linear mappings between a set of input variables and a set of output variables. The neural network builds discriminant functions from its neurons or processing elements.

SUPERVISED NEURAL NETWORKS IN ARCHAEOLOGY

Backpropagation is a special case of *supervised* learning algorithm. For all patterns in the training set, the actual output of the network is compared to the desired “target” vector, and the difference is used to adapt the weights of the network, in a backward way from output to input, and by recursively using the corrections already made. Here you have a list of archaeological applications.

SHAPE ANALYSIS: Since neural networks can express highly non-linear decision surfaces, they are especially appropriate to classify objects presenting high degree of shape variability. The idea of identifying a specific archaeological object type based on an image of the object is that the outer frame of the object is enough to specify the archaeological type it belongs to, and then to calculate the nonlinear mapping between shape and production form, past use, distribution, origin, chronology, etc. To accomplish that, it is necessary to identify this outer frame exactly. The inputs for this neuronal network are the individual tokens of a archaeological image. The number of output neurons is normally specified by the amount of different types. The Network helps to learn the nonlinear mapping between geometry (shape, size) and some conceptual labels (for instance use categories).

TEXTURE ANALYSIS: Archaeologists studying lithic remains usually wish to determine whether or not these stones have been used as tools and how they were used. The best way to do this is through the analysis of macro- and microscopic traces of wear generated by the use of the tool. The texture of different exemplars should allow us to discriminate between them. Our main assumption is that different artefacts have different textures because they have been altered by different work activities. M.H. van den Dries (1998) and Barceló and Pijoan-López (2003) have tried to classify lithic tools according to function using texture descriptors.

ROCK-ART STUDIES: One of the first Neural Networks applied to archaeological data is a classification of Late Bronze Age warrior decorated *stelae*, dated between 1100 BC and 7th. Century BC in south-western Iberian peninsula (Barceló 1995a). The goal was to build a working definition of the concept “Colonisation”, using as raw information the similarity among all *stelae* in the same region, and with the same chronology. In other words, whether all rock engravings are similar according to an explanatory concept, and not only whether they are similar in iconographic terms. This is a typical supervised learning example. For each *stelae* we know its chronology, geographical situation and the number of engraved figurations of imported items. In archaeological terms, the network has to generalize the relationship between iconography, geography, chronology and colonisation. Are the iconographically most complex *stelae* situated in the regions and chronological phases where Phoenician colonisation was stronger? (see also Barceló 1995b, Barceló 1996) .

ARCHAEOZOOLOGICAL APPLICATIONS: Paul Gibson (1996) has used neural networks and backpropagation learning to create a system for the ageing and interpretation of archaeofaunal material. From the tooth image the software must identify which teeth are present as an indicator 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. An algorithm was created that highlighted the approximate border of the teeth on the basis of background information. The neural network was trained to distinguish between different teeth silhouettes.

ARCHAEOMETRIC APPLICATIONS: Bell and C. Croson (1998) review some applications of backpropagation learning to archaeometry. For example, if archaeological material were analysed for chemical composition, each sample would be represented by a vector, the individual components corresponding to chemical concentrations found in that particular sample. The backpropagation training algorithm would associate each chemical composition of iron slags with a provenance or probable location of origin. John Fulcher (1997) reports an interesting work on obsidian provenance. The colour, translucency and texture of obsidian depends on the site from which it is collected. However, it is not possible to identify different obsidian samples solely on the basis of such broad characteristics. Fulcher uses proton induced x-ray emission (PIXE) data as input to the network and tried to classify into six provenance areas.

BURIAL ANALYSIS: C. Davino et al. (1999) have studied the Iron Age Italian cemetery of Sala Consilina. 173 graves have been selected and described using as input variables the following nominal variables: Preservation, Burial length, sex/age, depth, Quantity of grave goods, matter of most frequent

grave goods (clay, metals), type (according to the presence/absence of weapons and other features). The goal was to calculate a classification rule for age/sex, based on grave attributes. Also as an example of burial analysis, Suzanne Bell and Richard Jantz (2002) have published the results of an interesting project to classify skeletal remains. The goal was to distinguish the anatomic features of individuals from sites at both banks of river Missouri at different periods. Results suggest that the river was an effective barrier between sites and peoples since cranial measurements were sufficient to distinguish them.

SPATIAL ARCHAEOLOGY AND SITE CLASSIFICATION: Neural networks are also particularly adaptable for use with spatial data and Geographic Information Systems (GIS). C. Reeler (1999) reports on the application of a neural network / fuzzy logic hybrid system to spatial and other excavation data from Maori pa sites in New Zealand. Zupanek and Mlekuz (2001) have approached site classification through a comparison of roman inscriptions and the number of monuments dedicated to various deities. The idea was to discriminate between two sites (Emona and Poetovio) in modern day Slovenia, based on the religious backgrounds of both towns. Available data are in form of dedicatory inscriptions. Site Classification is also a necessary step for archaeological predictive modelling. It tries to develop and establish methods suitable for predicting potential locations of archaeological sites in geographic space. (see Ducke 2003)

TELEDETECTION AND REMOTE SENSING: The analysis of remotely sensed imagery from earth observation satellite systems such as Landsat-TM is usually achieved by machine-oriented pattern recognition techniques like backpropagation neural networks to detect land types and archaeological features. This image is first processed to extract all the line features or other distinct features (colour, spectral information) which are the candidates of archaeological features, then the supervised neural network is employed to tell whether a candidate archaeological pixel is really an archaeological element or not by its spectral characteristics. Furthermore, some postprocessing steps are taken to refine the result. If we are interested in a particular kind of archaeological element (a site, a megalith monument, an ancient road, etc.) we may construct specialist detectors.

CONCLUSIONS

In this paper we have seen that the aim of learning in social sciences and the humanities can always be viewed as a prediction process. But to produce good “cultural predictions”, learning obviously must make use of “properties of the cultural world that enable good predictions”. The ability to fuel prediction is the property that makes *regularity*, and regularity can also be described in terms of *useful redundancy*, based on the existence of *similarity effects* between cultural observables (size, shape, contents, composition, texture, location).

We have shown some archaeological datasets to discover a way to reconstruct the process or process which generated them. In all those cases, certain properties in the data were exploited, which can be referred as patterns, trends, regularities, harmonies and so on.

As used here, neural networks should be seen as an iterative non linear mapping algorithm, which is able to predict the output from a given set of unknown inputs, provided the network has been *trained* using known data. The network requires input data and a desired response to each input. Neural networks take this input-output data, apply a learning rule and extract information from the data. Essential to this learning process is the repeated presentation of the input-output patterns. Training is accomplished by sequentially applying input vectors, while adjusting network weights according to a predetermined procedure. The more data presented to the network, the better its performance will be. Training modifies selectively the network parameters so that application of a set of inputs produces the desired (or at least consistent) set of outputs, even though there is no pre-requisite for internal similarity between all inputs related to the same output. Just relational similarity is enough for calculating the mapping.

REFERENCES:

- BARCELO, J.A., 1995a, "Back-propagation algorithms to compute similarity relationships among archaeological artifacts". In *Computer Applications in Archaeology*. Edited by J. Wilcock y K. Lockyear. Oxford: British Archaeological Reports.
- BARCELO, J.A., 1995b, "Seriación de Datos Arqueológicos Ambiguos o Incompletos. Una Aplicación de las Redes Neuronales". In *Aplicaciones Informáticas en Arqueología. Teoría y Sistemas*. vol. II. 1995. Bilbao, Denboraren Argia
- BARCELO, J.A. 1996 *Arqueología Automática. El uso de la Inteligencia Artificial en Arqueología* Sabadell: Editorial AUSA (Cuadernos de Arqueología Mediterránea, 2).
- BARCELO, J.A., PIJOAN-LOPEZ, J., 2004, CUTTING or SCRAPPING? Using Neural Networks to distinguish kinematics in use-wear analysis. In *Enter the Past. The e-way into the four dimensions of cultural heritage*. Edited by W. Börner. Wien. (in press).
- BELL, S., CROSON, C., 1998, "Artificial Neural Networks as a tool for Archaeological Data Analysis". *Archeometry* 40 (1), pp. 139-151.
- BELL, S., JANTZ, R., 2002, "Neural Network Classification of Skeletal Remains". In *Archeological Informatics: Pushing the Envelope*. Edited by G. Burenhult. Oxford,
- DAVINO, C., Di MARTINO, R., MOLA, F., VISTOCCO, F., 1999, Riconoscimento automatico di Forma in archeologia: il caso delle necropoli di Sala Consilina. *Workshop "Intelligenza Artificiale per i Beni Culturali"* Bologna, 14 Settembre 1999. <http://studi131.casaccia.enea.it/enea/apps/aiia/consili.pdf>
- DIAZ, D., CASTRO, D., 2001, "Pattern Recognition applied to Rock Art". In *Archaeological Informatics: Pushing the Envelope*. Edited by Göran Burenhult. Oxford: ArchaeoPress (BAR Int. Series S1016), pp. 463-468.
- DRIES, M.H. van den, 1998, *Archeology and the application of Artificial Intelligence. Case studies on use-wear analysis of prehistoric flint tools*. Archaeological Studies Leiden University No. 1., Faculty of Archaeology, University of Leiden.
- DUCKE, B., 2003, "Archaeological Predictive Modelling in Intelligent Network Structures". In *The Digital Heritage of Archaeology. Computer Applications and Quantitative Methods in Archaeology*. Edited by M. Doerr and A. Sarris. Hellenic Ministry of Culture, Archive of Monuments and Publications., Heraklion (Greece).
- FINE, T.L., 1999, *Feedforward Neural Network Methodology*. New York, Springer.
- FULCHER, J., 1997, Neural Networks for archaeological provenancing. *Handbook of Neural Computation*. Edited by E. Fiesler and R. Beale. Institute of Physics/Oxford University Press.
- GIBSON, P., 1996, An archaeofaunal ageing comparative study into the performance of Human analysis versus hybrid neural network analysis. *Analecta prehistorica Leidensia* 28(1), pp. 229-232
- KULKARNI, A., 2001, *Computer Vision and Fuzzy-Neural Systems*. Prentice Hall, Upper Saddle River (NJ).
- MEHROTRA, K., MOHAN, C.K., RANKA, S., 1997, *Elements of Artificial Neural networks*. Cambridge (MA), The MIT Press.
- OPENSHAW, S., 1994, "Neuroclassification of Spatial Data". In *Neural Nets: Applications in Archaeology*. Edited by B.C. Hewitson and R.G. Crane. Kluwer Academic Publ. Dordrecht (Hol.)
- PAO, Y.H., 1989, *Adaptive Pattern Recognition and Neural Networks*. Addison-Wesley. Reading (MA).
- REELER, C., 1999, Neural Networks and Fuzzy Logic Analysis in Archaeology. In *Archaeology in the Age of the Internet*. Edited by L. Dingwall, S. Exon, V. Gaffney, S. Laflin and M. van Leusen. Oxford, ArchaeoPress (BAR Int. Series S750).
- RUMELHART, D.E., 1989, "The Architecture of Mind: A Connectionist Approach". *Foundations of Cognitive Science* Edited by M.I. Posner. Cambridge (MA): The MIT Press.
- THORNTON, C., 2000, *Truth from Trash. How Learning Makes Sense*. The MIT Press. Cambridge (MA).

WASSERMAN,P.D., 1989, *Neural Computing. Theory and Practice*. New York, Van Nostrand Reinhold.

ZEIDENBERG,M., 1990, *Neural Networks in Artificial Intelligence*. Chichester: Ellis Horwood.

ZUPANEK,B., MLEKUZ,D., 2001, "Counting the uncountable: A quantitative approach to the religious differences between the Roman Towns of Emona and Poetovio". In *Computing Archaeology for Understanding the Past*. Edited by Z. Stancic and T. Veljanovski. ArchaeoPress (BAR International Series, 931). Oxford.