Computational Intelligence in Archaeology. State of the Art

Juan A. Barceló[1]

[1]Department of Prehistory. Universitat Autònoma de Barcelona. Spain.

Abstract
Computational (or "artificial") intelligence is not just about robots. It is about understanding the nature of intelligent thought and action using computers as experimental devices. It also deals with the nature of inferential mechanisms and how computer programs allow us to discover how we produce inferences. In this paper I introduce some of the key points in Computational Intelligence in Archaeology, exploring the implications in our discipline, both theoretically and methodologically, of Machine Learning tools and techniques. Theoretical and practical aspects of computer programs able to reproduce the same tasks archaeologists do are reviewed in this paper. The question of whether it is possible to automate the archaeological knowledge production is of both great theoretical interest and increasing practical importance, because knowledge and information are being generated much faster than they can be effectively analyzed. Computable archaeology—if you do not like the expression "automatic archaeology"—is the proper way of exploring new ways of answering the questions we have not yet answered.

***Keywords**: artificial intelligence, expert systems, neural networks, machine learning*

1. INTRODUCTION

Is it possible to build a machine to do archaeology? Will this machine be capable of "interpreting" and "explaining" cultural heritage? So-called "intelligent" machines inspire instinctive fear and anger by resembling an ancestral threat—as a rival for our social position as more or less respected specialists. But robots are here, around us. I have never heard of a claim against washing machines "intelligently" selecting the best way to wash a specific fabric, or a camera with an "intelligent" device to measure luminance deciding by itself the parameters in which to take the picture. So, why be afraid of a machine classifying a prehistoric tool and deciding "intelligently" its origin, function, and/or chronology? Critics seem to think that computer programs are guilty of excessive simplification, of forcing knowledge or distorting it, and of failing to exploit fully the knowledge of the expert, but it seems to me that it is archaeology that is "narrow minded", not computer programs. The saddest thing is that archaeologists do not know how they know archaeological matters.

My personal approach is based on a reality that archaeologists and cultural heritage scholars could not evaluate 15 years ago: computer programs do work in real science, not only in archaeology. Perhaps they are more successful in other "harder" sciences,[1] but we cannot deduce from this fact that archaeology is a *different* kind of science.

In other scientific domains, the performance of humans at a particular task has been used to design a robot that can do the same task in the same manner (and as well) as a human. In many different domains it has been shown how 'robot scientists' can interpret experiments without any human help. Such computer programs generate a set of hypotheses from what is known about a scientific domain, and then design experiments to test them. Don't panic! I am not arguing that an artificial archaeologist will replace human archaeologists because it works better and cheaper than us.[2] We all know that Artificial Intelligence will eventually produce computer programs whose activity may seem dazzling, but it will not produce robotic persons. Computational intelligence in archaeology will *do* a lot, but it won't *be* a lot. Computational mechanisms cannot by themselves carry the weight of a scientific explanation. No machine is ever likely to provide an adequate explanatory analogy for the human brain or mind. Machines will not produce for free the categories we need for explaining past social action. "Real" machines are too simple and limited in their functions. Nevertheless, computer-based models mimic human behavior, and therefore they are good models of what archaeologists do rather than abstract models of brains or minds. The purpose is to understand how intelligent behavior in archaeology is possible. I

[1]S. H. Liao, "Knowledge Management Technologies and Applications—Literature Review from 1995 to 2002," *Expert Systems With Applications* 25 (2003): 155–164; S. H. Liao, "Expert System Methodologies and Applications—A Decade Review from 1995–2004," *Expert Systems With Applications* 28 (2005): 93–103.

[2]J. A. Barceló, "A Science Fiction Tale? A Robot Called Archaeologist," in *The World Is In Your Eyes. Proceedings of the XXXIII Computer Applications and Quantitative Methods in Archeology Conference*, ed. A. Figueiredo and G. Velho (Tomar, Portugal, 2005) 221–230; J. A. Barceló, "Towards a True Automatic Archaeology. Integrating Technique and Theory," in *Layers Of Perception. Advanced Technological Means to Illuminate Our Past. Proceedings of the XXXV Computer Applications and Quantitative Methods in Archaeology Conference*, ed. A. Poluschny et al. (Bonn: Dr. Rudolf Habelt GmbH, 2007) 413–417.

suggest the use of computer programs in such a way that the shortcomings of *natural* archaeology may be avoided.

2. THE MECHANICAL BASIS OF ARCHAEOLOGICAL KNOWLEDGE PRODUCTION

Archaeological artifacts have specific physical properties because they were produced so that they would have those characteristics and not some other. They were produced in that way, at least partially, because those things were intended for some given uses and not some other; they were tools, or consumed waste material, or buildings, or containers, or fuel, etc. If objects appear in some locations and not in any others, it is because social actions were performed in those places and at those moments. Therefore, archaeological items have different shapes, sizes, and compositions. They also have different textures, and appear at different places and in different moments. That is to say, the changes and modifications in the form, size, texture, composition, and location that nature experiences as the result of human action (work) are determined somehow by these actions (production, use, distribution) having provoked its existence.

In that sense, I am considering archaeology as a problem solving task:

- *Why is the present observation the way it is?*
- *What action or process caused what is seen now?*

In other words, why do the observed material entities have specific values of size, shape, texture, and composition; and why do they appear at some specific spatial and temporal location?

The main assumption is that some percept (*archaeological description*) is related to a causal affirmation about the causal event (social event, work activity) having produced the perceived evidence (*archaeological explanation*). In our case, it means to *predict* the cause or formation process of some archaeological entity, given some *perceived* evidence of the effect of this causal process. In its most basic sense, then, the task may be reduced to the problem of detecting localized key perceptual stimuli or features, which are unambiguous cues to appropriate causal events. For instance, a distinctive use/wear texture on the surface of a lithic tool, and not on others, predicts that these tools have been used to process fresh wood. We infer that at some moment a group of people was cutting trees or gathering firewood. Alternatively, we can consider that the shape of some pottery vases predicts their past use as containers for wine, and then we have evidence of wine production and trade; likewise, the composition of some graves predicts the social personality of the individual buried there and hence the existence of social classes. Here the output is not the object (trees or firewood, wine, social elite), but a causal affirmation: cutting trees or gathering firewood, wine production and trade, social power and coercion.

Can we implement this framework on a computer? Ideally, to solve such an archaeological problem we would need to know the solution beforehand. The reader may be surprised at this characterization of archaeological problem solving. Archaeological problems can be defined as "some material effect of social action in the past we wish to explain and we do not know how." Now we see that the past is knowable, only if it is already known. It seems a tricky way to solve problems! There is, however, nothing wrong in this approach. By making use of some previously stored knowledge, an automated archaeologist would infer from empirical data what it is that gave rise to those data. Explanation occurs when a perceptual input matches a perceptual memory containing a description of each causal event the system is expected to *recognize* or *identify* (fig. 1).

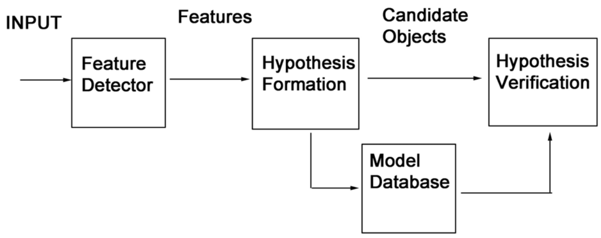

Figure 1. *A model for an archaeological recognition system.*

The model database contains all the models known to the system. The information in the model database depends on the approach used for recognition; it can vary from a qualitative or functional description to precise parametric equations. The feature detector applies operators to the input and identifies locations of features that help in forming causal event hypotheses. Using the detected features in the input, the hypothesizer assigns likelihoods to those events that may have produced the observed evidence. The knowledge base is organized using some type of indexing scheme to facilitate elimination of unlikely causal event candidates from possible consideration. The verifier then uses causal theories to verify the hypotheses and refines the likelihood of explanations. The system then selects the causal event with the highest likelihood, based on all the evidence, as the correct event.

Although there can be many criticisms to this approach, its advantage is that it is a practical and efficient way to solve archaeological problems and to explain archaeological evidence noted at the archaeological site.

3. Archaeological applications in the domain of Expert Systems

Therefore, "knowledge representation" is the key aspect, not laws, which are inviolate but explicit mappings that can be changed, and indeed are always changing, in a reflexive relationship that allows the archaeologist to accommodate new information. Given some empirical data (observations) about a particular archaeological case, and some bit of associative knowledge ('if…then' hypotheses and interpretations considered valid in social, anthropological, or historical theory), the archaeological problem can be explained in terms of the knowledge stored in the knowledge base. In other words, given some visual input and a candidate explanatory causal model, a correspondence can be established between them. This means that a small number of features are identified as matching features in the input and the model. Based on the corresponding features, a decision rule linking visual features with their causal process (social activity) is uniquely determined. The recovered decision rule is then applied to the model. Based on the degree of the match, the candidate causal event is selected or rejected. To be accepted, the match must be sufficiently close, and better than that of competing solutions.
An expert system is:

- *a computer system that is programmed to mimic the procedures and decisions that "experts" make;*
- *a domain specific knowledge base combined with an inference engine that processes knowledge encoded in the knowledge base to respond to a user's request for advice.*

The primary goal of expert systems research is to make expertise available to decision makers and technicians who need answers quickly. Today's expert systems deal with domains of narrow specialization. For expert systems to perform competently over a broad range of tasks, they will have to be given very much more knowledge. That makes these kinds of computer systems nothing more than a discrete plan for expressing scientific research, because they contain descriptions of intended courses of explanation. In that case, a specific explanation is created by searching through a range of possible explanations until the knowledge necessary to generate that explanation is found somewhere in memory. The procedure may be as follows: (1) during sensing, information from various sensors is collected and integrated into a central representation of the environment; (2) a number of possible explanations is generated and one explanation is chosen and finally applied.

This is not just a theoretical assumption. It implies programming some computer systems that act like archaeologists explaining their data. The technology really works, as it has been shown in many practical applications.[1] Automated typologies are the preferred domain of application but there are many other domains where expert systems technology has been applied. There are computer systems to mechanize the process of microscopic sample classification for ancient wood taxonomy determination, or to help archaeologists to interpret the results of archaeometric analyses, within the framework of provenance studies. Such programs produce one (or several) "diagnoses" according to the geographic origin of raw material, from a database of analyzed samples of known origin provided by the user. Other classificatory programs have been proposed in zooarchaeology and osteology. Even an expert system from the field of paleontology for the determination of a dinosaur species has been published. It helps the paleontologist to identify creatures from field data. Other systems help scientists to decode decorative patterns in pottery or rock-art, or to interpret the iconography of Greek ceramics. Given an input formed by the iconographic features of the personages who appear represented in one particular vase, the system answers with a reference to the mythological role present in that scene. Automated classification and diagnosis is also possible in epigraphy. In the domain of conservation analysis of archaeological materials, some prototype expert systems have also been proposed.

Of direct interest to archaeologists and humanists are the important applications of expert systems technology to solve geographical and geoscientific problems. The idea seems to be to build a full Geo-Expert System to answer questions in a seemingly intelligent way based on facts contained in a GIS and on the procedures and data available in a Digital Remote Sensing System. In earth resources applications, computer programs have demonstrated the possibility of incorporating spatial knowledge for land use prediction. Knowledge-based systems for aerial photo interpretation have been developed. For Remote Sensing, expert systems that help to detect relevant features in a landscape have been published. There are also some interesting applications in geomorphology, which can be useful to archaeologists. Applications to social analysis, that is to say the use of expert systems to explain social action, has not yet been fully explored.

Interest in expert systems has vanished in recent years, both in computer science and in archaeology. What seemed to be an interesting tool in the early 80s, never found the place it really merited in archaeology, compared with the situation in other similar domains. The real cause seems to lie in the poorly developed formal aspects of our discipline, even today. The post-modern criticism of the early 90s and its reification of subjectivism was an insurmountable obstacle to any effort that tried to analyze "objectively" the way we

[1] J. A. Barceló, *Computational Intelligence in Archaeology* (Hershey, VA: The IGI-Global Publishing Group, 2008) 38–60.

think. Within the last two decades, the view of problem solving based on pre-fixed plans and searching in restricted knowledge bases using well-defined operators for activating already existing sequences of explanations has come under scrutiny from both philosophers and computer scientists. The reliance on declarative expressions (expert systems rules) seems to be misplaced. The fundamentally unrepeatable nature of everyday life and human existence gives reality a significance that cannot be understood in terms of pre-defined, well-structured declarative expressions. This position argues that a cultural heritage scholar's understanding of archaeological, historical, or social data is rooted in the practical activity of coping with the everyday world. An explanation cannot be properly understood if considered independently of the context in which it occurs.

The relative success of expert systems is due to their working within a world in which the range of meanings for terms is circumscribed within a carefully selected micro-world. When the closed world is violated, the intelligent machine will not be able to function correctly. Explanatory knowledge cannot be defined by necessary and sufficient conditions. Archaeologists, historians, anthropologists, and sociologists do not have exact or complete definitions readily available. Rather, they are creating the boundaries of their concepts when there is a demand for it. These "blurred" concepts cannot easily be made operational. Many concepts seem to have a rather "generic" definition, which shapes up by instantiating the concepts with concrete objects. That is, our concepts do not have sharp boundaries initially, and the boundaries are drawn incrementally during use of the concept and probably also during use of other more or less related concepts. In fact, concepts are not fixed entities; rather, they are constructed on each usage by combining attribute values that are appropriate to the context. That raises the question of what mechanism constructs these unstable concepts. Obviously, it is not an expert system with its pre-fixed rules and facts!

Formally speaking, only one expert system was presented at CAA 2009 (although the paper was not submitted for publication: L. J. Dibble, "An Application of Rule-based Eco-cultural Niche Modeling to Archaeological Modeling"). This "intelligent" technology was used to predict the location of archaeological sites by using climate and fossil data. The papers published in this volume of the proceedings by Zhou et al. ("Towards Indexing and Data Mining All the World's Rock Art"); Keogh et al. ("Automatic Construction of Typologies for Massive Collections of Projectile Points and Other Cultural Artifacts"); and Mom and Drenth ("Continuity and Change: A Study of Shape of Late Neolithic and Early Bronze Age Vessels" [on the website]) fit this subject perfectly well.

4. Archaeological Applications in the Domain of Automated Discovery

Expert systems are useful, very useful indeed, because many archaeological problems can be structured in terms of a single template-matching mechanism. However, a template-matching scheme could only work provided that we had precompiled rules for all events to be explained. To explain social action produced in the past, an expert system would need a universal knowledge base covering the entire domain of interaction. Unfortunately, this is almost impossible to achieve, because it implies the existence of an infinite number of rules that would have the ability to recognize every unique archaeological evidence for what it is, and then to select an appropriate explanation for each possible historical state.

Archaeologists generally do not know why archaeological observables have the shape, size, texture, composition and spatiotemporal location they have. Instead, we have sparse and noisy observations or measurements of perceptual properties, and an incomplete knowledge of relational contexts and possible causal processes. This is a kind of *inverse problem*, where the consequence is known (observed), and the cause must be inferred.

Programming computers to be able to solve an inverse problem is a cross between statistics and computer science. We can formalize this inferential task in terms of a kind of "automated learning:"

Given:

- *an initial description of a theoretical entity;*
- *an instance of this entity;*
- *an explanation of the association between the concept and its instance; and*
- *some operating criteria*

Determine:

- *a generalization of the instance that substitutes initial description and is related to the explanation and operating constraints.*

In other words, the idea is to program a system that is able to look for common features between positive examples of the causal relationship to be predicted and common differences between its negative examples. This task is exactly like an example of a truth-function learning problem:

1	1	0	1	1	→	1
1	0	0	0	0	→	0
0	1	1	1	0	→	1
1	1	0	0	1	→	0
0	0	0	0	0	→	?

Concept learning problems have the same form, except that target outputs are either "yes" or "no" (or "true"=1 and "false"=0). Inputs that map onto "yes" are treated as positive examples of a particular concept. Inputs that map onto "no" are treated as negative examples (i.e. counterexamples). The process of finding a solution to such a problem is naturally viewed as the process of calculating the *communalities* among positive examples. As such, it is a variation of the philosophical theories seeing *induction* as a process involving the exploitation of similarity.

This implies that an automated archaeologist will learn explanatory concepts such as "15th century", "cutting", "killing", "social elite", or any other concept, provided it has enough known instances for the underlying event, and a general background knowledge about how, in this situation, a human action has generated the observed modification of visual appearances that it is using as perceptual information. When subsequently asked to determine whether novel instances belong to the same causal event, those instances that are similar to instances that are characteristic of a single event or of a single class of events will tend to be accepted. For instance, a machine will understand what a house, a castle, a burial, or a tool is when it learns how a prototypical house, a prototypical castle, a prototypical burial, or a prototypical tool was made, and under which social and economic conditions such objects existed.

This approach is a surrogate for experiment design. Experimental analysis is the process whereby the antecedents of a phenomenon are manipulated or controlled and their effects are measured. An obvious archaeological example is modern use wear analysis. By replicating lithic tools and using them for a determined period of time performing some activity—e.g., cutting fresh wood—we will be able to test the relationship between kinematics, worked material, and observed use wear on the surface of the tool. When laboratory replication is not possible, archaeologists are limited to mere observation. Ethnoarchaeological data can be also used to generalize observations and to learn explanatory general principles.

Computer scientists are intensively exploring this subject and there are many new mechanisms and technologies for knowledge expansion through iterative and recursive revision. Artificial Intelligence offers us powerful methods and techniques to bring about this new task. Fuzzy logic, rough sets, genetic algorithms, neural networks, and Bayesian networks are among the avenues we have to explore. Although statistical reasoning is still giving its support to all these methods, it is not classical statistical inference. Artificial Intelligence paradigms differ from usual classification and clustering methods in that they are (in comparison at least) robust in the presence of noise, flexible as to the statistical types that can be combined, and able to work with feature (attribute) spaces of very high dimensionality. In addition, they can be based on non-linear and non-monotonic assumptions, they require less training data, and they make fewer prior assumptions about data distributions and model parameters. The huge number of learning algorithms and data mining tools makes it impossible to review the entire field in a single paper.[1] Free computer programs like *Weka*[2] or *Tanagra*[3] can be explored to discover how to extract meaning and knowledge from archaeological data.

The most basic inductive algorithms are designed to find a conjunctive description for a single concept *C* that covers positive instances of *C* and that fails to cover negative instances. In this way, we can represent the solution to an inverse problem as a logical conjunction of Boolean features, values of nominal attributes, limits on the values of numeric attributes, or some combination of them. It is usual to refer to each component of such conjunction as a *condition* or a *test*. Alternatively, concept hierarchies provide a framework for memory organization, and a considerable amount of machine learning research has taken this approach. Such hierarchies can be represented as a decision trees consisting of nodes and branches. Each node represents a separate concept, typically with its own associated intentional definitions. The links connecting a node to its children specify an "is-a" or subset relation, indicating that the parent's extension is a superset of each child's extension. Typically, a node covers all of the instances covered by the union of its descendents. In fact, such a decision tree can be seen as a collection of rules, with each terminal node corresponding to a specific decision rule.

Inductive decision trees are increasingly applied in archaeology. Modern applications range from sex determination of buried human bodies to the discrimination of geo-archaeological soil data. In any case, it is in archaeometry where these methods have found their greatest popularity in recent years. Decision trees also seem relevant to paleoecological research. In CAA 2009 at Williamsburg some papers presented automated approaches to learning and discovery. Keogh et al. ("Automatic Construction of Typologies for Massive Collections of Projectile Points and Other Cultural Artifacts") use decision trees for discriminating among a series of projectile points. Zhou et al. ("Towards Indexing and Data Mining All the World's Rock Art") give an application of Support Vector machines and neural networks. Maaten et al. ("Visualization and Automatic Typology Construction of Pottery Profiles") apply affinity propagation methods, which can be seen as an interesting alternative to the most usual techniques for automated discovery. Märker et al. ("The Application of a Georelational Database and Data Mining Technologies for Predictive

[1] For a complete overview of recent developments in automated learning algorithm design and archaeological applications, see J. A. Barceló, *Computational Intelligence in Archaeology.* (Hershey, VA: The IGI-Global Publishing Group, 2008) 73–130.

[2] http://www.cs.waikato.ac.nz/ml/weka/.

[3] http://eric.univ-lyon2.fr/~ricco/tanagra/en/tanagra.html.

Site Modeling for the Paleolithic of the Iranian Plateau") consider regression trees for predictive site modeling using topographic and paleoecological data as input. Of related interest is Märker, Hockschield and Kanaeva ("A Multidisciplinary Integrative Geo-relational Database for Spatio-temporal Analysis of Expansion Dynamics of Early Humans").

Alternatively, we can use neural networks as a non-linear fitting mechanism to find regularities in a set of data. An Artificial Neural Network (ANN) is an information processing paradigm that is inspired by the way biological nervous systems, such as the brain, process information. It is composed of a large number of highly interconnected processing elements (neurons) working in unison accepting numeric inputs and sending numeric outputs. Neurons are organized in such a way that incoming vectors (descriptions) are sequentially transformed into output vectors (archaeological explanations) (fig. 2).

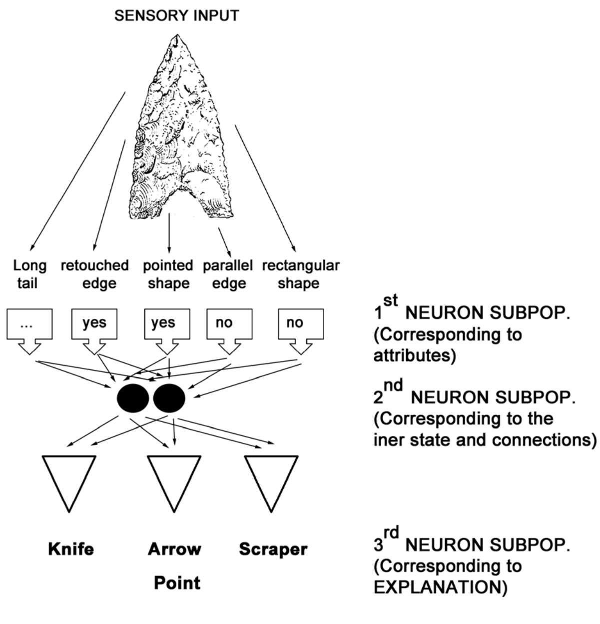

Figure 2. *A three-layer Neural Network topology, with a hidden layer.*

ANNs, like people, learn by example. An ANN is configured for a specific application, such as pattern recognition or data classification, through a learning process. Learning in biological systems involves adjustments to the synaptic connections that exist between the neurons. This is true of ANNs as well. In general, upon repeated presentation of various real examples and under the steady pressure of a learning rule or algorithm that makes small adjustments in the connections among artificial neurons, the network slowly but spontaneously generates a set of internal representations, one for each of the several features it is required to detect. The overall result is that after learning, the network contains a number of processors chained together in such a way as to produce the appropriate outputs, given a set of inputs. During learning, a network will typically develop a way of organizing its representations so that different inputs come to be represented as belonging to partitioned classes or groups (which may themselves be hierarchically ordered into various subgroups).

Given the particular vector representation of input data, images can be easily transferred into a neural network. The procedure is similar to that of expert systems, but here archaeological observables and archaeological explanations are no longer represented in terms of sentences, but as numbers. This fact allows the intelligent processing of archaeological image data.

We can cite the use of this technology in rock-art research, lithic arrow-point shape classification, the functional classification of lithic tools according to use wear descriptors, the interpretation of ancient sites according to their spatial features, the reconstruction of whole pottery vessels, the historical classification of ancient Mesopotamian seals, and the recognition of written characters in ancient documents, coins, and epigraphic inscriptions. Human and animal bone materials found in archaeological sites have also been investigated using neural networks.[1]

5 ARCHAEOLOGICAL APPLICATIONS IN THE DOMAIN OF INTELLIGENT IMAGE PROCESSING

Archaeology is a quintessentially visual discipline. Among all archaeological features, the most important factor in the recognition and/or explanation of an item is visual. Tasks such as identifying a pottery type, identifying decorative patterns or use wear in archaeological materials, recognizing archaeological structures in a satellite or aerial image, identifying layers or buildings at the site, and interpreting burials or settlement patterns can be considered to be within the purview of visual analysis. Visual perception makes us aware of such fundamental properties of objects as their size, orientation, shape, color, texture, spatial position, and distance, all at once. Visual cues often tell us about more than just optical qualities. In particular, the mechanical properties of a thing of any kind are often expressed in its image.

Human beings have the ability to recognize and classify images, identifying interesting patterns and single objects in them. Computers and robots can do this as well. Computer vision has been defined as a process of recognizing elements of interest in an image; it can be described as the automatic logical deduction of structures or properties of the three-dimensional objects from either a single image or multiple images and the

[1]For a complete overview of recent developments in neural network techniques and archaeological applications, see J. A. Barceló, *Computational Intelligence in Archaeology* (Hershey, VA: The IGI-Global Publishing Group, 2008) 73–130.

recognition of objects with the help of these properties.[1] Visual explanation occurs when a perceptual input matches a perceptual memory that contains a description of each causal event the system is expected to *recognize* or *identify*. Here, visual recognition means the reasoning process during which the social action's *observable* effects are used to specify the conceptual identity of the causal action. At this level, we should make this distinction:

- *Event recognition can be defined as the process of finding and "labeling events [in the real world] based on known causal models," that is, event recognition is the process of deciding what category of causal processes an observed effect belongs to.*
- *Event identification can be defined as the process of deciding which individual event it is, rather than deciding what category of causal processes it belongs to.*

This is exactly the inference mechanism we reviewed when dealing with Expert Systems: categorization in the guise of *associationism*. That is, the meaning of an object is accessed when its visual appearance activates a category representation linked to known interpretations via associations in memory. This is the basis of what has been called *pattern matching*. Pattern matching is actually a very broad concept, and it is useful to distinguish among types of matching. Pattern completion has been defined as the mapping of an incomplete pattern onto a completed version of the same pattern. Pattern transformation is the mapping of one pattern onto a different, related pattern. Pattern association is the arbitrary mapping of one pattern onto another, unrelated pattern. Finally, pattern recognition has been defined as the mapping of a specific pattern onto a more general pattern (that is, the identification of an individual as an exemplar of a class). In statistical terms, one first extracts a sufficient set of characteristic features from the primary input patterns, and then applies statistical decision theory for the identification and the classification of the latter.

Comparing an internal model with an external input is then the basis for perception understanding. The recognition of one input constitutes an internal cue, which facilitates explanation together with the external cues available from outside the brain. The outcomes of preliminary classifications should be combined to obtain patterns that are more global. They will in turn serve as input patterns to higher-level recognition devices. Thus, a problem will be solved by explaining something, and with the help of that result, explaining further.
To automatically solve a visual problem, we need a set of mappings that can be classified into three categories: a) the visual competences that map different visual features to each other; b) the problem solving routines that map visual features to explanatory concepts or representations of various kinds residing in memory; and c) the learning programs that are responsible for the development of any map. In other words, an automated archaeologist should determine whether visual data "it currently sees" corresponds to a causal event "it already knows." Recognition requires knowledge about how social action happens, and about the specific changes generated by all related social and natural processes. To design or analyze such a vision system amounts to understanding the mappings involved.

A system like this resembles an associative memory. During the recall stage, a cue pattern is presented to the system by activating visual input units. This causes signals to be sent and to activate the output processors. If the associative mechanism runs properly, then the pattern of activation in the output will be the pattern that was originally associated with the cue pattern. Visual input is acquired in the form of a vector of intensities (feature detectors), and used as a cue pattern to retrieve its associated explanation, which is represented as a vector of activity in the memory's output. The advantages are obvious:

- *When a previously stored (that is, "familiar") pattern is "seen" by the system, it is amplified, and the system responds with a stronger version of the input pattern.*
- *When an unfamiliar pattern is "seen" by the system, it is dampened, and the response of the machine is shut down. This is a kind of unfamiliar response.*
- *When part of a familiar pattern is "seen", the system responds by "filling in" the missing parts. This is a kind of recall paradigm in which the part constitutes the retrieval cue, and the filling in is a kind of memory-reconstruction process.*
- *When a pattern similar to a stored pattern is "seen", the system responds by distorting the input pattern toward the stored pattern. This is a kind of assimilation response, in which similar inputs are assimilated to similar stored events.*
- *Finally, if a number of similar patterns have been stored, the system will respond strongly to the central tendency of the stored patterns, even though the central tendency itself was never stored.*

Such an associative memory, however, is not limited to the association of only those specific individual objects that the robot has seen before. If such were the case, the mechanisms underlying archaeological automatic explanation would be of limited use. As archaeologists, we must identify a range of novel visual data as corresponding to a given type of object. Generalization is part of our ability to identify objects and events; we typically can identify social actions that have been performed in the past even when the visual appearance of their material consequences in the present does not exactly match what we know of previously memorized cause/effect associations. The capability for archaeological recognition implies, then, the existence of some previous form of learning, in which the abstract potentially explanatory categories have been created and defined. The goal of recognition is to perform these

[1] A. D. Kulkarni, *Computer Vision And Fuzzy Neural Systems* (Upper Saddle River, NJ: Prentice Hall, 2001).

identifications correctly, in the sense that identification reflects a meaningful property of the world that is independent of the particular data that is being interpreted.

There are different ways to implement such mappings using computational intelligence technologies. The task is to extract the *statistical central tendency* of a series of visual exemplars (the learning set) in such a way that the computer program encodes information not just about the specific exemplars, but about the stereotypical feature-set displayed in the training data. That is, it will discover which sets of features are most commonly present in the exemplars, or in commonly occurring groupings of features. In this way, semantic features statistically frequent in a set of learning exemplars come to be both highly marked and mutually associated. "Highly marked" means that the connection weights about such common features tend to be quite strong. "Mutually associated" means that co-occurring features are encoded in such a way that the activation of one of them will promote the activation of the other.

The easiest way to create an associative memory for archaeological image data is by assuming that there is a roughly fixed set or vocabulary of "supposed" descriptive regularities shared by a single population of objects, which are also distinctive enough. Partial identification of individualized parts of the input is carried out by specialized shape detectors, processed, and eventually decoded. At the highest level, a decision mechanism selects the concept corresponding to that represented by the cognitive detector activated by the highest quantity of partial identifications.

Traditionally, however, archaeological visual input has been translated into a set of universal picture stereotypes used as subjective bits of information, e.g., "round," "ovoid," or (even worse) by user-defined stereotypes such as "hat-shaped," "cigar-shaped," or "kidney-shaped." This kind of identification-based analysis is a misleading way of solving the archaeological visual problem. In fact, it is not a visual analysis, because the original visual input is being "described" in non-visual terms (words). Low-level recognitions are assumed to be known, but no criteria are given about their reliability. It is the human user who feeds the computer system with an interpreted input, in which each feature contains the result of a previous inference. In this way, the receptive field properties of low-level visual feature detectors do not encode the salient features of the input image, but rather the previous knowledge the user has about the features characterizing the archaeological evidence.

For a long time computer scientists and archaeologists have realized that the only way of reasoning with images is feeding an intelligent computer system with images and not with "words." Images are no longer "described" but acquired and automatically encoded as vector arrays. They contain all the useful information to derive geometry and texture for any explicative purpose. However, the reconstruction of detailed, accurate, and photo-realistic 3D models from external images is a difficult task, in particular for large and complex archaeological evidence sets.

In many cases, a vector that encodes the two-dimensional coordinates of the edge defining the boundaries of the object seems enough for intelligent image analysis. At the CAA Conference in Williamsburg, many papers advocated the use of laser scanners and similar equipment for doing this task. The idea is to rely on geometry and coordinate measuring rather than on linguistic descriptions. More specifically related to the intelligent analysis of archaeological images, Zhou et al. ("Towards Indexing and Data Mining All the World's Rock Art") use generalized Hough transforms for the geometric hashing of original images of rock art. This is a kind of *harmonic analysis*, shape-unrolling methods that convert observed boundaries or edges to a function of the coordinates of points (or pixels) delimiting the contour. Keogh et al. ("Automatic Construction of Typologies for Massive Collections of Projectile Points and Other Cultural Artifacts") transform the original artifact's contour into a one-dimensional "time series" representation. But instead of using the complete one-dimensional vector, they select a *small* subsection. They call such subsections *shapelets*, which invokes the idea of a small "sub-shape". Martínez-Carrillo et al. ("A Proposal of Ceramic Typology Based on the Image Comparison of the Profile") describe pottery profiles using anchor points and Euclidean coordinates of morphometric landmarks. Koutsoudis and Chamzas ("3D Pottery Shape Similarity Matching Based on Digital Signatures") analyze 3D polygonal meshes of complete ceramic vases. Zhou, Geng, Wu, and Shui ("A System of Pottery Recovery and Repair") also consider the necessity of geometrical modeling of image data as input to a pattern matching mechanism. The same subject is considered in Kleber and Sablatnig ("Scientific Puzzle Solving: Current Techniques and Applications"). Because in many cases laser scanning data acquisition and polygon meshes can be too complex for posterior clustering, Maaten et al. ("Visualization and Automatic Typology Construction of Pottery Profiles") prefer an approach based on nonlinear dimensionality reduction using "shape contexts." The key idea behind shape contexts is to sample a set of points from the shape contour and to describe these points with local descriptors—the shape contexts—that measure the relative angle and distance to the other points that were sampled from the shape contour.

Another very important domain of computational intelligence approaches for image understanding is remote sensing and satellite imaging. There are two different domains of application within this field:

- *When remote sensing data are intensity measurements to be reconstructed as an image. This is the case of laser scanning (3D-scanners) or the different modalities of geoelectric/georadar/geomagnetic surveying.*
- *When remote sensing data are images or part of an image (satellite imaging or aerial photos).*

In the first case, computational intelligence techniques like artificial neural networks have been successfully applied to a number of geophysical modeling problems, including parameter prediction and estimation, classification, filtering, and optimization. In archaeological geophysical surveying, neural networks can be used to interpolate the possible nonlinear spatial trend among magnetic differential measurements obtained in an archaeological geophysical survey and derive estimates of feature burial depths, allowing a three-dimensional reconstruction of buried subsurface remains to be made. The neural network approach potentially offers several advantages in terms of efficiency and flexibility over more conventional data interpolation techniques.

In the second category of remote sensing data, the input is not an array of sensor measurements, but an aerial or a satellite image. Remotely sensed images are digital pictures composed of pixels showing grey-level values. In many satellite or remote sensing cases, such values are the intensities of specific spectra of electro-magnetic radiation of either form of reflection or emission. Because different types of objects have different physical natures in terms of reflection, absorption, and emission, these values of two or more layers are used to categorize the pixels into several groups. The idea is then to distinguish between the various categories of spatial features of interest to archaeologists. It can be a difficult task, because archaeological features comprise a complex spatial assemblage of disparate land-cover types—including built and/or linear structures, numerous vegetation types, bare soil, and bodies of water—each of which has different reflectance characteristics. Conventional image classification techniques assume that all the pixels within the image are pure, that is, that they represent an area of homogenous cover of a single land-cover class. This assumption is usually untenable with pixels of mixed land-cover composition.

In employing machine-learning approaches, the idea is to use image data (brightness, greenness, wetness, and ratio indexes) and geographical information (forest, grass, water, archaeological elements, etc.) to train an input-output nonlinear relationship model. The resulting network can be exported and used for new satellite images, where map data have not been interpreted, and these geographical values may be predicted. The input data typically comprises a set of multi-spectral data, although it may also include measures of image texture or ancillary data. Supplemental information, such as soils or elevation attributes, and even non-numerical data, e.g. ground cover classes or soil types that might assist in the classification, can be easily integrated. In the output layer, there is one unit for each class in the classification.

6 Computational Intelligence in the Domain of Anthropological and Historical Modeling

Solving archaeological problems implies answering a *double causality* question:

- *Given the perception of visual inputs, the automated archaeologist should explain what social activity produced in the past the evidence perceived in the present.*
- *Once it knows what social activity was performed, where, and when, the automated archaeologist should explain why such activities were performed there and then, and in what way.*

So far we have only presented examples of the first kind. It is time to present some possible applications of the second kind of archaeological knowledge production.

Let us begin with standard social explanation. It is usual in the social sciences to classify people according to social attributes. Computational intelligence tools can help in such a classification. In the social sciences, a neural network can classify a population into homogenous groups, using factors such as age, sex, and other socio-economic variables to infer social status or position. The obvious archaeological example of this kind of analysis is the explication of ancient burial practices.

Spatiotemporal modeling is another approach to historical explanation. Neural networks have been applied in this domain, in ecology, geography, and historical dynamics.[1] In such examples, a neural network is trained on sets of dependent variables (outputs) measured at known spatial or temporal locations (inputs) to generalize how such ecological or social aspects are spatially or temporally related. Ecological applications show that neurocomputation is a viable technique and has advantages over linear models. Examples are very diverse, from the classification of soil structure based on soil sample data to the prediction of changes in the dominant species of grassland communities based on climatic input variables.

The most promising area of research is that of social simulation using computational intelligence algorithms. The idea is to represent human societies using computational units that simulate the acting of different social agents. An Artificial Prehistoric Society is then a complex set of computational reactive units simulating how a group of people behaved in the past. There is an

[1] Relevant examples are quoted in detail in J. A. Barceló, *Computational Intelligence in Archaeology* (Hershey, VA: The IGI-Global Publishing Group, 2008) 297–323.

increasing number of examples in the specialized literature, including very interesting studies of hunter-gatherer societies and the beginnings of social life. Artificial societies are also being programmed for studying the origins of agriculture and the dynamics of prehistoric and ancient societies. At CAA 2009, two relevant papers were presented at the conference but not submitted for publication (M. Hinz, "Where Do You Want To Go Today? Pathfinding, Algorithms and Agent-based Modeling," and M. D. Harris, "Applying a Neutral Agent-based Model of Lithic Material Procurement to the Middle Atlantic Region"). Somewhat related to this discussion is the paper by Whiteley, Moore and Goel ("Beyond the Marsh: Settlement Choice, Perception, and Spatial Decision-making on the Georgia Coastal Plain") describing the use of a cell-based simulation for modeling ancient spatial decision-making.

7 CONCLUSIONS

Two different views on archaeological knowledge production have been presented here:

1) *Archaeological knowledge is viewed as something that can be stored, coded, matched, and displayed. That means that information is derived from external objects and flows into the system via the senses. It is denotational because it is an encoding. An intelligent computer memory is just a storehouse of denotational encodings.*
2) *Archaeological knowledge is not given but created as transformations of stimuli. Information does not exist in the world waiting to be extracted by a rational agent, but rather, the agent is situated in meaningful contexts, in which information should be defined as a function of the local needs and concerns of the agent. Perceiving a world implies distinguishing "possibilities for action" and not naming or identifying per se. That is to say, it can be understood as recognizing the circumstances to act with or upon. This means that the contents of perception (and hence, the structure of the phenomenal world) is largely determined by the self-organized dynamics of the cognitive system and pre-rational dispositions that are embodied in the cognitive agent. Being a perceiver, the automated archaeologist should literally create a phenomenal world, because the process of perception first defines relevant distinctions in the sensory environment.*

Consequently, two different and indeed opposite approaches to the use of computational intelligence for research efforts appear:

1) *we can build an "automated archaeologist" simply by telling it what it needs to know;*
2) *we can build it as a learning machine.*

Both approaches have their advantages. They are often presented as competing paradigms, but since they attack cognitive problems in different ways, we should see them rather as complementary methodologies.

Bringing artificial intelligence into archaeology introduces new conceptual resources for dealing with the structure and growth of scientific knowledge. The discussion is between what is considered an *artificial* way of reasoning (computer programs) and a *natural* way of reasoning (verbal narrative). Critics of computationalism insist that we should not confound scientific statements with predicate logic operations, since discursive practices or argumentations observed in a scientific text are not "formal." By that reasoning, they are tributary, to a certain extent, from natural language and the narrative (literary) structure from which scientific texts derive. I take the opposite approach: scientific problem solving stems from the acquisition of knowledge from a specific environment, the manipulation of such knowledge, and the intervention in the real world with the manipulated knowledge. The more exhaustive and better structured the knowledge base is, the more it emulates a scientific theory and the easier the solution to the scientific problem will be, and the more adequate the interpretations we will get.

ACKNOWLEDGEMENTS

This paper originated in the presentation text I gave at the opening of the special session on Computational Intelligence in Archaeology at the CAA 2009 in Williamsburg. I thank all the researchers present at the session as well as the authors and speakers. Some of the ideas put forward in the text were developed as a response to the discussions at the end of the session, where many important questions were asked. I am also grateful to Bernard Frischer for allowing me to read the original papers once they were submitted by the authors.

BIBLIOGRAPHY

Barceló, J. A. "A Science Fiction Tale? A Robot Called Archaeologist," in *The World Is In Your Eyes*. *Proceedings of the XXXIII Computer Applications and Quantitative Methods in Archaeology Conference*, edited by A. Figueiredo and G. Velho, 221–230. Tomar, Portugal, 2005.

Barceló, J. A. "Towards a True Automatic Archaeology. Integrating Technique and Theory," in *Layers of Perception. Advanced Technological Means to Illuminate our Past. Proceedings of the XXXV Computer Applications and Quantitative Methods in Archaeology Conference*, edited by A. Poluschny, K. Lambers, and I. Herzog, 413–417. Bonn: Dr. Rudolf Habelt GmbH, 2007.

Barceló, J. A. *Computational Intelligence in Archaeology.* Hershey, PA: The IGI-Global Publishing Group, 2008.

Kulkarni, A. D. *Computer Vision And Fuzzy Neural Systems*. Upper Saddle River, NJ: Prentice Hall, 2001.

Liao, S. H. "Knowledge Management Technologies and Applications—Literature Review from 1995 to 2002." *Expert Systems With Applications* 25 (2003): 155–164.

Liao, S. H. "Expert System Methodologies and Applications—A Decade Review from 1995 to 2004." *Expert Systems With Applications* 28 (2005): 93–103.