Archeologia e Calcolatori
7, 1996, 1231-1245

ARTIFICIAL SOCIETIES AND COGNITIVE ARCHAEOLOGY

"With the possibility of constructing artificial systems, a new methodology of scientific inquiry becomes possible"

CONTE, GILBERT, 1995, 4.

1. INTRODUCTION

In this paper I shall describe an approach to the study and understanding of social processes which has recently become prominent: systematic experimentation with "artificial societies" created on computers. I shall consider the contribution that this new research tool can make to what is sometimes called "cognitive" archaeology. By way of illustration, I shall focus particularly on how artificial societies techniques may be used to enhance our understanding of the role played by collective belief and misbelief systems in the initial emergence of certain types of complex society, and describe some relevant experimental work.

2. ARTIFICIAL SOCIETIES

The phrase "artificial societies" typically refers to the systematic exploration of the properties of societies of information processing "agents" (or "actors" or "animats"), which exist within simulated spatial environments created on one or more computers (GILBERT, CONTE 1995; EPSTEIN, AXTELL 1996). In essence agents are entities which decide their actions in the light of their perceived surroundings and past experience.

They range from a relatively simple "reflex" type, whose behaviour is straightforwardly determined by a small number of condition-action rules, to those endowed with higher-level cognitive abilities such as planning, albeit in limited form. Typically, agents have manifestations within the environment which implies inter-agent observability. Inter-agent communication of simple kinds, and various types of agent reproduction, can also easily be programmed.

The technical development of artificial societies is largely driven, not by social scientists (although often in cooperation with them), but by computer scientists who are concerned with the understanding of the essentials of social phenomena from their own perspective. Also it derives in part from the current trend to reject the strong emphasis on rationality and logic which has long existed in artificial intelligence research and that, many would say, has hindered AI research from its earliest beginnings.

But such computer-based experiments are not entirely new in archaeology. Computer simulations have been used to address archaeological problems on many occasions in the past two decades and more. What, exactly, is new? Two things are relatively new, apart from a continuing and remarkable increase in available computing power. Firstly, there is now an emphasis on the systematic discovery of what is possible in societies, rather than on what may or may not have actually occurred at some particular location and time in the past (for a discussion of this shift of emphasis see, for example, Conte, Gilbert 1995). Secondly, explicit modelling on a computer of cognitive processes is now feasible given the progress made by artificial intelligence and related techniques.

The particular significance of the latter development is that over the past decade archaeology itself has moved a little (in UK at least) back towards the view that the role and content of cognition in prehistoric societies is something that archaeologists can and should address.

3. Cognitive archaeology

In the UK, "cognitive archaeology" is particularly associated with the name of Colin Renfrew (e.g. 1982, 1994). He defines it to mean "the study of past ways of thought as inferred from material remains" (Renfrew 1994a, 3) and sees it (*contra* the polemics of the anti-processualists) as a necessary extension of processual archaeology to include cognitive phenomena, inspired, of course, by cognitive science (itself often led by artificial intelligence research). More specifically, he is seeking insights into the cognitive representations ("cognitive maps" or "mappae" in Renfrew's terminology) with which human beings reason, and hence into ideologies and the functional role that they can play.

For Renfrew, an important component of cognitive archaeology is the nature and role of the archaeology of religion (Renfrew 1994b). Renfrew discusses the nature of religion, notes that it involves some framework of beliefs, and sees as basic, but difficult to express adequately, what I shall describe as a sense of awareness of and a subordinate relationship towards some supernatural power or powers. He further emphasizes the central role of the "religious experience" in the creation and maintenance of religious belief systems, which brings into consideration feeling and affect as well as belief. His practical emphasis, of course, is on the recognition of religious phenomena in the archaeological record and on their interpretation.

4. Artificial societies and cognitive archaeology

At first sight any connection between human belief systems and computers seems remote. But on closer inspection important and useful linkages may be developed. It is quite possible to identify, in a consistent and defensible way, the beliefs that an agent in an artificial society created on a computer holds about its world. Furthermore, it is quite possible to discuss and to experiment with the formation of sets of collective beliefs common to members of a community of agents, and to relate such beliefs to the "truth" in the created world, so that it becomes possible to study, the creation and impact on the society and its future of collective *mis*beliefs. It should be noted that by "belief" I here typically mean a descriptive belief about the environment.

Remarkably, it is just as likely that an artificial society will settle to a collective system of misbelief as that its beliefs will be uniformly true. *Indeed some degree of collective misbelief is almost inevitable* (DORAN 1994; DORAN 1995). In almost all cases, the limited access of individual agents to the reality of their world (in both space and time), together with actual errors of perception, communication and generalization make sure of that.

But a system of misbeliefs need not damage a society of agents that holds it. Far from it. It is easy to devise both mathematical and informal examples where misbeliefs are beneficial to a community. An informal example is the community of robots in a laboratory all of whom wrongly believe that passing another oncoming robot on the right is impossible and whose movements are thereby greatly facilitated in most circumstances.

There is a controversial issue here. As just implied, it is natural in computer-based artificial societies to distinguish sharply between what an agent believes, and what is actually there in the computer created world. Discrepancies are apparent. But many social scientists find a sharp distinction between belief and misbelief unconvincing and unacceptable as applied to our own social world, and therefore see here a major argument against the validity of computer based models of society.

Perhaps even more controversial is that we can begin to model *affect*. Worldwide, there have been several projects of this type (e.g. the Oz project at Carnegie Mellon University; BATES, LOYALL, REILLY 1992) and more are ongoing. It seems that it is possible to encompass within artificial societies analogues of certain types of emotional dynamics as well as aspects of individual rationality. Hence it may prove possible to address religious experience and its impact. I shall not attempt to explore this possibility here, important though it is. Before passing on, however, I should perhaps make quite clear that there is no suggestion that computers, or the agents created within them, actually *feel* emotions any more than there is *real* water inside a computer running a mathematical model of the Mediterranean Sea.

5. THE EMERGENCE OF SOCIAL COMPLEXITY

I shall now try to illustrate how the new technology of "artificial societies" can usefully link up with a "cognitive archaeology" by considering the initial emergence of a certain type of social complexity.

5.1 *Social non-complexity and social complexity*

By social complexity in human societies, I have in mind a combination of integration, ranking and centralised decision making. Thus I take social complexity to mean:

- coordinated behaviour by a relatively large ranked community (beyond individual family size), led by
- one individual (or a small group) which reflects upon the community and its environment, and takes decisions and issues commands accordingly, and whose commands are normally followed, all this on a semi-permanent basis.

By contrast, non-complexity refers to a society, such that:

- there is only loose cooperation in the community as a whole, and
- no individual or small group holds any permanent leadership role

These are imprecise definitions, but they do serve to indicate the transition from non-complexity to complexity that I wish to consider. Note that these definitions are intended to encompass both societies of relatively mobile hunter-gatherers and societies of sedentary agriculturalists, on the presumption that the underlying integrative processes apply to both types of society.

5.2 *Causes and trajectories*

The question now to be asked is: what causes may be ascribed to the transition from non-complexity to complexity and what particular trajectories may be followed? It would surely be an error to assume that there could be only one cause or trajectory.

The immediate cause for a transition may be a clear change:

- in the environment
- in the typical cognitive abilities of agents
- in the mix of agents in existence

Alternatively, it may lie in:

- an incremental process which has crossed a threshold, or even
- some seemingly very minor chance event

It is easy to hypothesize more specific factors, and a wide range of proposals may be found in the research literature with varying degrees of evidence offered in support of them (see, for example, JOHNSON, EARLE 1987).

Some of the factors more frequently suggested are:

Factors involving the relationship between the population and its environment

- population increase
- technological advance
- environmental stress

- resource affluence
- resource and/or population concentration
- circumstances favoring centralised decision making e.g. necessity of risk management, opportunities for large scale hunting, territorial defence
- environmental circumscription

Cognitive, affective and social factors
- increase in group cognitive ability
- development of enabling belief systems
- extension of kin relationships
- "promotion" of social components
- impact of trade networks
- emotional manipulation by aspiring leaders
- overriding of mechanisms which restrain aspiring leaders
- warfare

Of course, several factors may operate simultaneously. A number of those just listed will be prominent in the discussion that follows, notably population concentration, the emergence of enabling belief systems, and the extension of kin relationships.

5.3 *Causes and trajectories in artificial societies*

In practice, proposals about causal factors responsible for observed transitions in human societies are difficult to take forward. The evidence from the archaeological record is frustratingly slight. However, and this is at the heart of the new opportunity, trajectories from non-complexity to complexity may be generated and observed *within an artificial society*, and hence their properties, preconditions and relationships may be established and systematically studied at least within that context.

However, for experiments to be informative trajectories must emerge from a "substrate" of processes laid down by the experimenter and not merely be directly enforced by the experimenter. If the substrate layer is sufficiently rich then it will be possible to discover which trajectories are viable and will emerge and which, although they may seem *a priori* plausible, do not. Of course, if the substrate omits the foundations of a whole class of phenomena, then obviously phenomena of that type cannot be observed. If anything is to be learned about social processes in the real world then the substrate processes specified by the experimenter must themselves be broadly realistic. If this requirement is met, then whatever is discovered will be a contribution to the general body of social theory.

It should perhaps be stressed that there are internal structural constraints in any computerbased artificial society, just as there are in mathematical models of a more traditional variety. It is a serious misconception to believe that any behaviour can be obtained merely by "tweaking" the parameters any computer-based society. It is precisely the discoverable implications of the constraints and structural assumptions of the society, which may be highly non-obvious and non-trivial, which constitute potential new knowledge.

6. Illustrative examples

I shall now briefly describe particular examples of computer-based "artificial society" experiments (performed at the University of Essex) which address different but related trajectories from social non-complexity to social complexity as I have defined it.

6.1 *Mellars*

Mellars (1985) suggested that a key factor in the seeming emergence of social complexity in the Upper Palaeolithic period of SouthWestern France might be stable resource concentration, leading in turn to population concentration. He identified specific ecological reasons why resource concentration might occur, and suggested that as a direct consequence a form of social complexity would arise, notably involving centralised decision making. He did not, however, discuss just *how* that social complexity would arise.

6.1.1 Testing Mellars

The EOS project used an artificial society created in a software "testbed" (written in the computer language Prolog) to explore a detailed computation interpretation of the social processes implicit in Mellars' proposal. The interpretation explored may be summarized as follows:

PRIOR COGNITIVE CONTEXT

> *a society of agents able to represent and plan in terms of their immediate physical and social context, and hence to make rational choices, and able to use a generalized notion of a dominance relationship (as might be developed from relationships in the core family — see later)*

IMMEDIATE CAUSE

> *resource/population concentration in the environment, and limited total available resource, requiring effective cooperation if the society is to survive*

==> OUTCOME

> *a semi-stable and effective multi-level decision hierarchy*

Notice how at the heart of this interpretation is a notion of rationality and its effectiveness.

The computer experiments performed and reported in detail elsewhere showed that multilevel hierarchies of agents, with centralised decision making of a type, did indeed arise in conditions of population concentration (DORAN *et al.* 1994; DORAN, PALMER 1995).

6.2 *Rappaport's hypothesis*

RAPPAPORT (1971), following BERGSON (1935), formulated a very different hypothesis, originally couched in the now unfashionable terminology of cybernetics. This may be summarized as:

> *The level of intelligence reached by human beings* impedes *social cooperation based on rational choice, because it is too easy for individuals to see and follow alternative social options, with consequent discordancy.*
>
> *Hence a* special belief system, *providing strong enforcement of social conventions, is needed to establish control and enable cooperation. Rappaport particularly sees* religious or sanctified *belief systems as fulfilling this need.*

Clearly, what is different from Mellars' view, and therefore from the EOS project work, is that Rappaport is suggesting that self-interested rationality *cannot* bring about the necessary cooperation (for example, at subsistence tasks), and that rather a type of religious belief system is necessary. This analysis applies whether or not there is population concentration, although it is reasonable to suppose that the particular nature of the belief system and its consequences might differ depending upon the degree of population concentration.

6.3 *Evidence for Rappaport's hypothesis*

Some support is to be found for Rappaport's hypothesis in the EOS experiments themselves. Although, as stated earlier, multi-level hierarchies with a form of centralized decision-making do emerge in these experiments, and in circumstances of population concentration, they are not usually very effective (depending on the precise properties of the environment). The presence of rational decision making at all levels in the hierarchy, rather than on mere following of commands, tends to lead to slow and therefore inefficient responses.

Also relevant is recent work by CONTE and CASTELFRANCHI (1995) who have used an artificial society to study the impact of social norms in groups of agents. They have shown the ways in which norms, by restricting the choices available to the individual, clearly benefit the group as a whole, notably in the control of aggression.

Mathematical studies by MOSES and TENNENHOLTZ (1991) support Rappaport's hypothesis by drawing attention to and achieving insights into what these authors call the «Golden Mean Problem». This is the problem of finding a compromise between overly restrictive social rules and overly liberal ones – where the issue is not one of ethics in any abstract sense but of the effectiveness of the society from the point of view of its designer.

6.4 *The origins of collective belief systems*

Rappaport's hypothesis immediately prompts such questions as: «What is a "religious" or "sanctified" belief system? Where does such a thing come from? What is its typical structure, and how exactly does it function to enable cooperation?».

It is clear that for Rappaport 'sanctity' relates to a combination of authority and unfalsifiability. *A sanctified belief system controls just because it asserts the existence of great power and cannot be logically refuted.* But Rappaport does not fully discuss the origins of such systems of belief. Before I attempt to do so there is a difficult to be faced. It is apparent that one personally committed to a particular religious belief system will tend to answer the question of its origins very differently for that system (and quite possibly for all such systems) from those who are not so committed. Here I shall proceed on the assumption that the belief systems at issue are *not actually true*. If a particular belief system is true, then it is intuitively much easier to explain how it comes into existence, for it is presumably a matter of observation rather than, in some sense, error.

An answer to the question: «How do misbelief systems originate?» may now be offered by reference to the individual cognitive level. At the heart of adaptive processes of cognition (which human beings certainly deploy) are processes which generate and vary concepts.

By "concepts" I here mean structured representations, internal to the agent, which define categories and relationships between categories in the world external to the agent. These categories may or may not correspond to actual external entities. Concepts *are* sets of linked beliefs. Adaptive processes vary the structure of these representations in systematic or partially random ways, often in response to incoming sensory evidence.

Because they are inevitably working with partial information, it is in the nature of such processes that they often generate new representations not entirely in accord with external reality. They are, in a sense, fallible hypothesis generators. Much is proposed, both accurate and inaccurate, and much of what is proposed cannot immediately or easily be refuted by the agent, even when it is in fact inaccurate.

For example, a special case of concept variation which is particularly relevant here is the generalization of a "parent-of" relationship to that of a "parent-figure", that is, another agent recognized on some particular grounds as akin to a parent. This extended concept, although a misbelief in a certain sense (assuming there is no such category as "parent-figure" designed into the environment) may well be useful. However, further extensions for example the proposition "all parent-figures are friendly" can clearly be damaging to the individual.

The effect of generative processes operating upon internal representations, coupled with processes of inter-agent communication and inheritance from one generation to the next, is that there is a continual modification of and extension to the pool of beliefs available to the community and that many of these beliefs are inaccurate. Where, however, a particular set of (mis)beliefs offers a competitive advantage to a (sub-)population, then it will tend to persist and stabilize (compare DAWKINS 1989).

This extension to Rappaport's original hypothesis locates the origin of collective (mis)belief systems in error prone processes of internal concept manipulation. It is, of course, strongly influenced by what has been shown possible at the computational level in artificial intelligence systems and hence in artificial societies. Agents in artificial societies can certainly be designed to modify and communicate sets of beliefs in this way.

6.5 *Testing Rappaport*

I now briefly describe computer work at Essex which illustrates, and in certain respects amplifies, aspects of the foregoing discussion.

6.5.1 Projecting family concepts

The work first to be described concerns agents' internal processes of concept manipulation. MAYERS (1995) has extended the EOS research to address the role played by family relationships and their cognitive extension in the formation of different types of complex society (compare TODD 1985). Specifically, he has defined and implemented in the EOS testbed generalized family relationships (whereby, for example, one agent may see another unrelated agent as a "parent figure" and act accordingly) as generic concepts within agents. Necessarily, this includes specification of the circumstances under which an instance of the relationship is to come into being (i.e. the relevant concept instance is created within an agent) and of the particular behavioral effects it will have. He has also discussed specific mechanisms which would create generic concepts of this type.

Mayers has performed initial experiments which explore the relationships between small scale societies relying upon basic family relationships and the circumstances under which they will re-structure in terms of extended family relationships. His work compares interestingly with ongoing simulation studies of the !Kung san kinship system by READ (1995) whose emphasis is on the contrast between genealogical and "conceptual" kin relationships and the practical impact on group dynamics and survival of variations in the latter.

6.5.2 Pseudo-agents and their significance

The final piece of experimental work to be described here concerns what may be called "pseudo-agents": agents which are believed to exist by some or all agents in the artificial society, but in fact do *not* exist in the environment of the society.

At Essex there has been developed a software testbed (SCENARIO 3, written in the computer programming language C) which supports an artificial society scenario intended to address the formation by agents of representations of such "pseudo-agents" and their impact. The testbed features:

- a two dimensional spatial environment
- mobile agents and immobile resources
- agent perception of their surroundings
- agent internal representations of other agents and of resources (inc. memory limitation and forgetting)
- agents moving towards and harvesting resources for energy (in mutual competition)
- death by starvation or by ageing
- (asexual) reproduction of agents
- friendship and information passing
- killing

In the foregoing list words such as "harvesting" and "killing" must, of course, be understood to denote relatively simple events within the testbed which may nevertheless be so denoted without gross misrepresentation. For example, "harvesting" is said to occur when an agent located at a resource reduces the energy level of the resource to zero, and increments its own internal energy store by a corresponding amount. Energy is used up by an agent (decrementing its energy store) as an agent moves around in the testbed. "Killing" involves two agents meeting and one possibly becoming "dead" (and deleted from the world) and the killer acquiring the killee's energy store.

One agent maintains a representation of another when it holds information (not necessarily accurate) about the other and about some of the other's characteristics.

One feature of the testbed requires special explanation. This is the *friendship* relation which is intended to be representative of social relationships in general. An agent may decide that another agent, which it encounters, is its *friend* (in practice the decision is simulated by a random probability). If agent X "thinks of" agent Y as a friend, then:

- X passes information about resources to Y whenever it can
- X never attempts to kill Y

Further, agent X will not attempt to kill Y if X and Y have a believed friend, say agent Z, in common.

Observe that friendship is not necessarily symmetric. X may treat Y as a friend whilst Y does not so treat X.

6.5.3 Resource agents and cults

Recurring features of primitive religions are the localization in the environment of supernatural entities, and the attribution of humanlike or spiritual characteristics to the inanimate. Of the innumerable examples that might be cited, two that are prominent in the research literature are the Tsembaga of New Guinea who are reported to have associated different types of spirits with different localities in their territory (RAPPAPORT 1984, 38-41), and the ancient Zapotec of Oaxaca, Mexico who saw the sky and the earth as animate (MARCUS, FLANNERY 1994, 57).

Now the testbed described in the preceding section may be set so that from time to time an agent "agentifies" a resource, that is, *wrongly* comes to "think of" a resource as if it were an agent. We may call a pseudo-agent like this a *resource agent*. Once an agent forms a representation of a resource agent, that representation may be passed to other agents by inter-agent communication. It may also be passed from one agent to its offspring. If the circumstances are right, therefore, the representation may spread.

In its "thinking" an agent does not distinguish between resource agents and "real" agents, so an agent may even regard a resource agent as a "friend" and act towards it accordingly (of course, messages sent to a resource agent go nowhere). When a set of agents all come to believe that they have the same resource agent as a friend, we may call that a *cult*. The resource agent in question may be called the cult head. The advantage to the members of a cult is that *they will not kill one another*. This follows from the properties of the friendship relation given earlier.

6.5.4 Experimental results

What is found experimentally (DORAN 1995) is that even a very low-frequency possibility of agents coming to believe in resource agents regularly leads to the formation of large and enduring cults and that, all other things being equal, killing in the society is then greatly reduced and the average population of the society over time increased.

To give the reader a feel for these experimental trials and what happens within them, there follows a summary account of key events in one typical trial. The agent and resource identifiers are exactly as they appear in the testbed and its output:

Resource number 11 was initially 'conceived' as a friendly agent, 110000000rrr, by agent number 248 on time cycle 251.

On cycle 268 this resource agent had just one 'host' agent in a population of 3 – not agent 248, which died before cycle 268, but agent 8000254. A cult around 110000000rrr built up thereafter (comprising descendants of 8000254) typi-

cally with about 30 member agents, and lasted for hundreds of time cycles. Initially the cult was also around a 'dead' real agent 20000267 which itself had lived for only one time cycle, but was also regarded as a friend by agent 8000254. Memory of 20000267 was lost on time cycle 287.

Note the appearance of a "dead" agent in the account. The potential of dead agents as cult heads was not anticipated, though obvious enough in hindsight. In fact, a resource agent is better than a dead agent as a cult head because the former, unlike the latter, can be "seen" and awareness of it thereby refreshed.

It should perhaps be stressed that these experimentally observed phenomena are not entirely straightforward to obtain. The agent society embodies many structural parameters, and different combinations of settings for these parameters lead to very different outcomes. Current experiments are focussing on the dynamics of multiple competing cults each exploiting several resources, and with each agent possibly a member of several cults.

6.6 *Discussion*

The experiments just described address certain key aspects of Mellars' and Rappaport's approaches to the emergence of social complexity. They illustrate and explore (a) a connection between rationality and the emergence of multi-level, centralised decision-making, (b) how agents may manipulate internal cognitive representations of, for example, genealogical relationships and (c) how, arising from this, collective misbelief with certain religious characteristics may come to exist and persist in an agent population with a certain type of longer-term benefit for that population.

But it is clear that there is much of relevance that is *not* touched upon by these experiments. For example, Rappaport's notion of sanctity supporting a central authority is not addressed. In the experiments last described, there is no sense in which cult heads have power, nor is power attributed to them by the cult members. And no linkage is offered between collective misbelief and multi-layered centralised decision making.

These experiments comprise, in fact, just a small number of pieces from a computational jigsaw. At a more abstract level, much needs to be done to characterize what are the specific structural types of collective belief system that are possible and their particular relevance to trajectories to social complexity.

A further major deficiency is that the "religious experience", which both Rappaport and Renfrew emphasize is at the heart of sanctity systems, is ignored. I have earlier indicated that computational treatments of affect are beginning to appear in the research literature, so that artificial societies research can possibly contribute even in this direction.

Perhaps the most important insight to come from the experiments, apart from the simple insight that such work is possible at all, is to discover just how rich and intricate is the connection between the cognition of the individual and the macro-behaviour of the society. Indeed, it seems most unlikely that this connection can be fully understood without computer based experimentation. For example, as noted earlier TODD (1985) has proposed that different basic family types give rise to different types of emergent socio-political organization. This suggestion can only be assessed in a very uncertain way by using empirical data. By contrast, the creation of an artificial society incorporating the appropriate low level cognitive and social processes as a substrate (not, it must be admitted, an easy task) would enable Todd's predictions to be challenged and assessed in a much sharper way. Do the anticipated correlations actually occur in the computer created world and if so under precisely what conditions?

7. The relevance to practical archaeology

Cognitive archaeology, as does all archaeology, rests on the practical recovery and interpretation of the material record. How can the insights provided by experimentation with artificial societies aid this fundamental activity?

If we can pin down the relationships between environmental and social preconditions and emergent macro-behaviour (even if only within an acceptably plausible artificial society) then at least the practical archaeologist is provided with useful guidelines as to what is and is not likely.

In fact, experimental work of this type always reveals a multitude of variable parameters for the artificial society in question and reveals just how complex is the dependence upon them of the patterns of emergent phenomena observed. But if we are (as we surely should be) systematic in our experimentation, then we can learn that in the parameter space of the artificial society:

- some particular patterns of emergent behaviour are impossible
- some particular patterns of emergent behaviour appear only in certain regions of the parameter space.

For example, the formation of cults (section 6.5.3) is closely dependent on the "demographic" and belief passing characteristics of the agent population.

From the discovered properties of the parameter space may be derived propositions of the form:

conditions P always give rise to emergent phenomena of type E (if you have P you must have E)

emergent phenomena of type E only arise in conditions P (if you have E you must have P) which may be then used as steps in a chain of inference or as testable predictions. More general forms of such propositions are couched in probabilistic terms.

8. Conclusions

It is clear that creating and experimenting with artificial societies can cast light on social processes, including those involving cognition, and can therefore contribute to the development of cognitive archaeology in potentially very important ways.

But there is no free lunch. Indeed, this lunch is decidedly expensive – the computer-based experimentation required is substantial and technically demanding. As yet I see little willingness to undertake it, even amongst AI and artificial societies specialists.

Argumentation is always easier (and often more fun!) than systematic experimentation. Unfortunately, *meta-analysis is no substitute for new evidence*. I anticipate that the computational social theory needed for the development of cognitive archaeology will in the first instance come from well-resourced disciplines of immediate practical importance. For example, closely relevant issues of organizational ideology arise in new and industrially backed research in computational organization theory. But targeted experimentation led by "cognitive" archaeologists can, and I hope will, contribute to the overall picture at key points.

Jim Doran
Department of Computer Science
University of Essex

BIBLIOGRAPHY

Bates J., Loyall A.B., Reilly W.S. 1992, *An Architecture for Action, Emotion and Social Behavior*, in A. Cesta, R. Conte, M. Miceli (eds.), *Pre-Proceedings of MAAMAW'92 Conference, S. Martino al Cimino (Viterbo), Italy, July 1992.*

Bergson H. 1935, *The Two Sources of Morality and Religion*, (tr. R. Ashley Audra and Cloudesley Brereton) New York, Henry Holt.

Conte R., Castelfranchi C. 1995, *Understanding the Functions of Norms in Social Groups through Simulation*, in N. Gilbert, R. Conte (eds.), *Artificial Societies: the Computer Simulation of Social Life*, London, UCL Press, 252-267.

Conte R., Gilbert N. 1995, *Introduction: Computer Simulation for Social Theory*, in N. Gilbert, R. Conte (eds.), *Artificial Societies: the Computer Simulation of Social Life*, London, UCL Press, 1-15.

Dawkins R. 1989, *The Selfish Gene*, Oxford University Press (new edition).

Doran J., Palmer M., Gilbert N., Mellars P. 1994, *The EOS Project: modelling Upper Palaeolithic social change*, in N. Gilbert, J. Doran (eds.), *Simulating Societies: the Computer Simulation of Social Phenomena*, London, UCL Press, 195-221.

Doran J. 1994, *Modelling Collective Belief and Misbelief*, in M. Keane *et al.* (eds.), *AI and Cognitive Science '94*, Dublin, Dublin University Press, 89-102.

DORAN J. 1995, *Simulating Collective Misbelief*, in *Pre-Proceedings of Simulating Societies '95 Symposium, Boca Raton, Florida, USA, 15-17th, September 1995.*

DORAN J., PALMER M. 1995, *The EOS project: integrating two models of Palaeolithic social change*, in N. GILBERT, R. CONTE (eds.), *Artificial Societies*, 103-125, London, UCL Press, 103-125.

EPSTEIN J.M., AXTELL R.L. 1996, *Growing Artificial Societies: Social Science from the Bottom Up*, Brookings Institution.

GILBERT N., CONTE R. (eds.) 1995, *Artificial Societies: the Computer Simulation of Social Life*, London, UCL Press.

JOHNSON A.W., EARLE T. 1987, *The Evolution of Human Societies*, Stanford, California, Stanford University Press.

MAYERS S.D. 1995, *Modelling the Emergence of Social Complexity using Distributed AI*, Masters dissertation, Department of Computer Science, University of Essex, UK.

MARCUS J., FLANNERY K.V. 1994, *Ancient Zapotec Ritual and Religion: an Application of the Direct Historical Approach*, in C. RENFREW, E. ZUBROW (eds.), *The Ancient Mind: Elements of Cognitive Archaeology*, Cambridge, Cambridge University Press. 55-74.

MELLARS P.A. 1985, *The Ecological Basis of Social Complexity in the Upper Palaeolithic of Southwestern France*, in T.J. PRICE, J.A. BROWN (eds.), *Prehistoric Hunter-Gatherers: The Emergence of Cultural Complexity*, New York, Academic, 271-297.

MOSES Y., TENNEHOLTZ M. 1991, *On Formal Aspects of Artificial Social Systems*, Technical Report CS91-01, January 1991, Weizmann Institute of Science, Israel.

RAPPAPORT R.A. 1971, *The Sacred in Human Evolution*, «Annual Review of Ecology and Systematics», 2, 23-44.

RAPPAPORT R.A. 1984, *Pigs for the Ancestors*, (2nd edition), New Haven and London, Yale University Press.

READ D.W. 1995, *Kinship Based Demographic Simulation of Societal Processes*, in *Pre-Proceedings of Simulating Societies '95 Symposium, Boca Raton, Florida, USA, 15-17th September 1995.*

RENFREW C. 1982, *Towards an Archaeology of Mind*, Cambridge, Cambridge University Press (Inaugural Lecture).

RENFREW C. 1994a, *Towards a Cognitive Archaeology*, in C. RENFREW, E. ZUBROW (eds.), *The Ancient Mind: Elements of Cognitive Archaeology*, Cambridge, Cambridge University Press, 3-12.

RENFREW C. 1994b, *The Archaeology of Religion*, in C. RENFREW, E. ZUBROW (eds.), *The Ancient Mind: Elements of Cognitive Archaeology*, Cambridge, Cambridge University Press, 47-54.

TODD E. 1985, *The Explanation of Ideology: Family Structures and Social Systems* (translated by David Garrioch), Oxford, Blackwell.

ABSTRACT

This paper describes an approach to the study and understanding of social processes which has recently become prominent: systematic experimentation with "artificial societies" created on computers. The contribution that this new research tool can make to a "cognitive" archaeology is considered. It is particularly asked how artificial societies techniques may be used to enhance our understanding of the role played by rationality and by collective belief and misbelief systems, including religious belief systems, in the initial emergence of certain types of social complexity. Experimental work discussed aims to explore the relevant insights of Paul Mellars and of Roy Rappaport. One particular set of computer based experiments demonstrates how, in certain circumstances, social groups with some of the characteristics of "cults" may arise, with long term benefit to their individuals involved.