

27 Contrasting models of Upper Palaeolithic social dynamics: a distributed artificial intelligence approach

Mike Palmer & Jim Doran

27.1 INTRODUCTION

The EOS project (eos, Gk. "dawn"), is a collaboration between the University of Essex, Department of Computer Science (Professor Jim Doran and Mike Palmer), the University of Surrey, Department of Sociology (Professor C. Nigel Gilbert) and the University of Cambridge, Department of Archaeology (Dr Paul Mellars). It has as its primary objective to formulate and experiment with a computational interpretation of an informal model proposed by Paul Mellars (1985) for the emergence of social complexity in the Upper Palaeolithic period in South-western France. A feature of the project is the use of the concepts and techniques of distributed artificial intelligence, or DAI (Bond & Gasser 1988), which enable cognitive as well as social factors to be taken into account.

In this paper (some sections of which are revised from Doran and Palmer, in press), we shall first describe the Mellars model, and then the computer software testbed which has been created, the way(s) in which the Mellars model has so far been formulated within it, and initial experiments which have been performed.

It will become apparent that the EOS project is not a straightforward computer modelling exercise (if there is such a thing). The utilisation of only partly understood DAI techniques, whilst we believe entirely appropriate, means that interpreting the Mellars model becomes a DAI research investigation in its own right. In the short term the EOS project is therefore perhaps best seen as a piece of exploratory DAI research motivated and guided by an archaeological problem, with the contribution of the DAI research to the solution of the archaeological problem taking

place in the longer term. This is not, of course, an unusual relationship between an archaeological problem and a piece of mathematical or computer science research.

Also in this paper we shall explore, using the same DAI approach, a contrasting model of Upper Palaeolithic society proposed by Gamble (1991). It will appear that when couched in DAI computational terms the Mellars and Gamble models are considerably less different than their conventional archaeological formulations suggest. Further, it turns out that the analysis can be used to guide further interpretation of the Mellars model.

27.2 SOCIOCULTURAL TRAJECTORIES TO COMPLEXITY

Much research effort has been put into the study of the long-term dynamics of human society. Although there seem to be the foundations of general theory (e.g. Johnson & Earle 1987) much remains to be done. We are concerned here "merely" with the transition from a relatively "simple" egalitarian hunter-gatherer society to one with a more "complex" structure. We note that even this restricted problem has attracted more controversy than proven insight.

The problem must first be stated a little more precisely. We follow Cohen (1985:99) in meaning by an "egalitarian" society one characterised by: fluid group organisation, freedom of movement and relatively immediate and easy access to resources, immediate consumption, simple division of labour, and relatively direct personal leverage on other individuals. By contrast, a "complex" society features centralised decision making,

ranking, technological specialists and other role differentiation, territoriality and ethnicity (Cohen 1985:105; Mellars 1985:285–286). These are, of course, anthropological characterisations. It is far from trivial to translate them into formal mathematical or computational terms as will become apparent.

But how can an egalitarian society thus characterised be transformed, or transform itself, into a more complex one? There has been much debate and certain recurring suggestions concerning, for example, the impact of population increase and/or concentration, of sudden resource scarcity, and of specific climatic changes. In the time-spans concerned, the possibility of significant development in basic cognitive function seems remote, but this does not exclude rapid fulfilment of pre-existing cognitive potential.

27.2.1 The case of South-western France

There is archaeological evidence of just such a transition from simple to complex society during the early stages of the Upper Palaeolithic Period of South-western France and North-western Spain, that is, from around 30,000 to 20,000 years ago. This was about the time of the last glacial maximum when the evidence suggests a congregation of many species into the Iberian peninsula.

The evidence for the development of some form of social complexity (only to be sketched here — for details the reader may refer to Mellars 1985) comprises:

- larger and more abundant archaeological sites;
- exceptional wealth, density and stratigraphic complexity of archaeological material in sites;
- abundant and sophisticated cave art (e.g. at Lascaux);
- elaboration of bone and antler technology;
- high frequency of ceremonial burials;
- abundance of “trade” objects (e.g. marine shells).

The issue is: what explanation can be given of the emergence of complexity in this particular case? We shall give Mellars’ answer to this question shortly.

In passing we should note, (a) that the significance of the archaeological evidence cited is itself a matter of controversy. Not all would agree that it implies the emergence of complexity. And, (b) that the circumstances may well have been much more complex than the simple foregoing statements suggest. There is evidence of short-term temperature rises which Mellars suggests may have had a marked effect on the economic, resi-

dential and social structures of the local human groups (ibid:289), an implication being that the emergence of complexity may have been an irregular or recurrent process.

27.2.2 The Mellars model

Mellars has suggested that, in the case of South-western France, a particular combination of ecological conditions led to population concentration and a degree of stable and relatively large co-residential units involving sedentism over a substantial part of the annual cycle, and that this was the crucial step in the emergence of more complex social structures. The ecological conditions were specifically:

- an exceptional wealth and diversity of food resources;
- a strong pattern of concentration of these resources, both at particular locations and at particular periods in the annual cycle;
- a relatively high degree of stability and predictability in the spatial distribution of these resources from year to year (Mellars 1985:284–5).

The second part of the Mellars proposal, less elaborated, is that population concentration (however caused) leads to social complexity as defined earlier. He suggests as

“appropriate” and perhaps even “necessary” the emergence of certain individuals with increasing status or authority to organise and co-ordinate the activities of other members of the group... [for example] communal hunting activities» (Mellars 1985:285).

It is highly relevant to what follows, that Cohen (1985) has discussed the impact of cognitive stresses and physical congestion in population concentrations. Such problems, he suggests, were solved by aspects of complex organisation such as social stereotyping (a form of cognitive economy) and the co-ordination of groups through differentiation. Cohen’s addition to the Mellars model (perhaps it is really no more than an added emphasis) suggests that the core problem faced by a population concentration is the interaction between specific problems of logistics and cognitive overloading. A society that fails to solve these problems will fail comprehensively.

27.2.3 The importance of cognitive factors

Note that both Mellars and Cohen introduce cognitive factors into their explanations of emergence. The importance of cognition in such explanations has been much more stressed in recent

years, notably by Renfrew (e.g. 1987 and in press). Mithen has recently commented that

«part of the reason that archaeologists continue to have difficulty in explaining change is that they have neglected one locus for that change — people making decisions about what to do» (Mithen 1990:2)

This trend is partly, perhaps, a matter of frustration with the limitations of non-cognitive explanations, and partly recognition that cognitive science and artificial intelligence studies (see later) now provide useful means to handle these cognitive factors, at least in embryo.

Two particular insights coming from cognitive science and artificial intelligence work are these:

- The crucial emergence of the cognitive ability to “distance” reasoning from the immediate situation, notably the ability to predict and to plan, (compare Alexander 1989:459) should be distinguished from natural language usage, and recognised as a prior and equally important development to language use (Bloch 1991).
- The fundamental significance of “cognitive limitation”. The meaning of “cognitive limitation” is that any information processing device (including the human brain) has limitations on its capacity and speed, and must therefore engage in various heuristic devices to get round its limitations, for example, heuristic generalisation, information discard, sub-optimal reasoning, re-use of past problem solutions in similar problem situations. These “heuristic devices” are fundamental components of human behaviour, including social behaviour (Doran *et al.* 1991; Cohen 1985).

It follows from what has been said, that to formulate more precise models along the lines suggested by Mellars and Cohen requires the handling of cognitive factors, notably cognitive limitation, without necessarily becoming involved in a consideration of natural language. This has suggested to us active recourse to the concepts and techniques of cognitive science and artificial intelligence.

27.3 COMPUTER-BASED MODELLING

A traditional computer modelling procedure involves four basic steps:

- 1) ABSTRACT — isolate the essentials of the target system;

- 2) TRANSLATE — cast the essentials into a model within the chosen formal conceptual repertoire;
- 3) EXPERIMENT — establish the properties of the model;
- 4) COMPARE — compare the behaviour of the formalised model with the target system.

Within the EOS project we take as our starting point the informally expressed existing model (Mellars 1985) which partially deals with step 1. Our emphasis is on steps 2 and 3. This, as we shall see, involves the deployment of DAI concepts and techniques at the research frontier. Step 4, the link back to the archaeology, is our long term objective.

For a review of computer-based simulation and formal modelling in archaeology see Doran 1990.

27.3.1 Computer modelling using distributed AI

The research field of Artificial Intelligence is somewhat misleadingly named and is often misunderstood. It is largely concerned with non-numeric computation and, in particular, with achieving operational computational interpretations of aspects of cognition: of, for example, planning, memory, learning, and induction. An important concept of artificial intelligence studies is that of an “agent”: a process, however simple, which collects information about its environment, makes decisions about its actions, and acts. This use of the term “agent” is a little different from its use in some other scientific literature, in which, for example, it may be used to denote anything endowed with causal powers (Bhaskar 1978:49).

“Distributed AI” is a relatively recent development of artificial intelligence studies. It concerns the properties of sets of inter-communicating agents (“multiple agent systems”) coexisting in a common environment. The aim may be to study the properties of such systems in an abstract way, or to design systems of immediate practical use, or — the more relevant case here — to use such a programmed multi-agent system as a model of a human or other real world system. The potential impact upon the social sciences is obvious.

An important facet of DAI research has been the development of various “software testbeds” (eg Bond & Gasser 1988, chapters 1 and 6.3; Doran *et al.* 1991). In this context a testbed is a computer program which provides a platform upon which one can build and experiment with multi-agent systems. Typically it will provide some means of defining and creating a number of agents, an environment in which these agents are to exist, a protocol which they can use to interact

with one another, and means of monitoring the agents collective behaviour when the system is "run". Testbeds differ primarily in their degree of generality or specificity. General testbeds give the experimenter greater flexibility in the range of experiments that can be set up. Specialised testbeds are designed to support experiments in a particular domain or with particular agent architecture's or communication protocols.

It seems clear that DAI (especially of the "experimental testbed" variety) should be able to contribute to the study of social change generally, and to that suggested by Mellars in particular, since it enables the integrated study of both external (environmental) and internal (cognitive) factors.

27.4 THE EOS MODEL AND THE CONCEPT OF EMERGENCE

Our experimental model exists at two levels: as (i) a target set of core abstract processes, based upon the Mellars, Cohen and other formulations and informed by DAI concepts, which we propose as sufficient (maybe not necessary) to capture the transition from a simple to a complex society, and (ii) detailed behavioural specifications for agents (implemented in an ad hoc testbed) from which the proposed core processes are intended to emerge.

The foregoing use of the word "emerge" needs some clarification especially as we have earlier referred to the «emergence of social complexity». The concept of emergence has recently received a great deal of attention, both in the fields of DAI (Steels 1990, 1991; Wavish 1991) and in the philosophy of science (Bhaskar 1978, 1986, 1989; Sayer 1992). Following Bhaskar (1978:98) two senses of the concept may be distinguished. First (our earlier usage) it is used to denote an historical process in that certain phenomena, such as a hierarchically structures society, emerge in time. Second, it is used to suggest that some properties of a complex entity may be, at least in part, a consequence of the relations between its constituent components. It is the latter sense which we employ at this point. The core processes referred to above are supposed in part "emergent" from the properties of the agents and their shared context.

27.4.1 Core processes

We follow Mellars(1985) and focus on what can happen when agents with the elements of human-like cognition, especially planning, share a common environment, are strongly aware of one another, and have to collectively perform one or more complex resource acquisition tasks. Guided

by the DAI repertoire of concepts and techniques we identify three core processes which follow naturally upon population concentration:

- 1) Temporary planned co-operation between groups of agents, involving a temporary leader, to achieve effective resource acquisition and to avoid negative interactions, for example too many agents seeking to acquire the same resource.
- 2) Conversion of temporary groups into semi-permanent groups, with a leader, with agents forming internal representations of the groups, so that once formed a group continues to work together unless something causes it to fall apart.
- 3) "Recursive" development of hierarchical structuring, i.e. groups of groups of groups...

A central point is that an agent's representations of other agents and of groups of agents, in our terms its "social model", will have a determining effect on how it behaves in relation to other agents and ultimately on how effective it is in acting and surviving in its environment (Figure 27.1). For example, group permanency (point 2 above) is a matter of agents becoming aware of their membership of a group and then treating fellow group members differently and having relatively extensive knowledge about them. And agents within a group might have expectations about how the others might behave, enabling quick responses to resource opportunities. An implication here is that an agent must have some "basic" cognitive abilities which dynamically modify its social model in an appropriate way.

We now turn to the EOS testbed in which experimentation with these ideas is being conducted.

27.4.2 The EOS testbed

The EOS testbed has been implemented in "Object Oriented" style in Quintus Prolog and runs on a Sun SparcStation. It allows the creation of several types of object: a world, resources and agents. Messages may pass between the "objects". A scheduler simulates the concurrent activity of objects by activating each in turn. One time unit of the simulation correspond to the activation of each object known to the scheduler. Objects are created by specifying the necessary parameters in a configuration file.

27.4.2.1 The world and its resources

The world keeps a record of the location of the agents and resources. The user can specify the world's dimensions (a square).

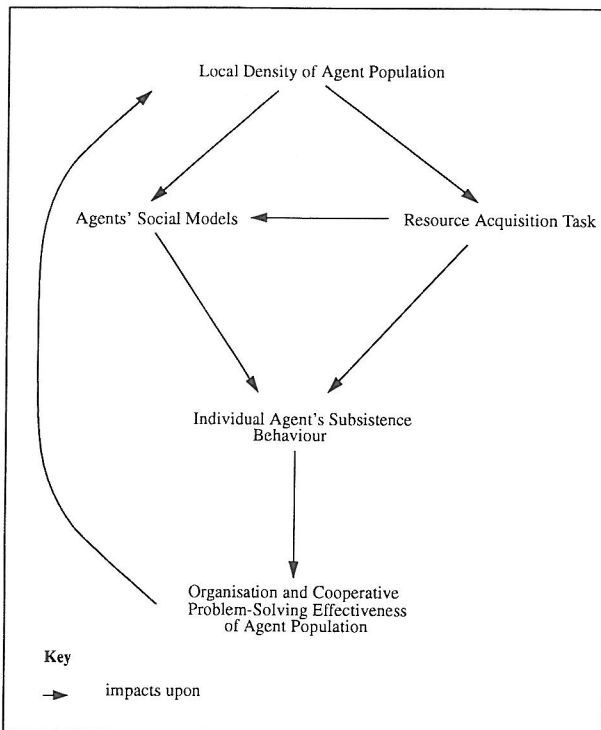


Figure 27.1: The heart of the EOS Model. The local density of the agent population impacts upon the task of acquiring resources and the agents' perception of their social context. These two factors have an effect on the resource acquisition strategies adopted by individual agents (e.g. co-operative). The effectiveness of these strategies will in turn impact upon the local density of the agent population

Resources can be acquired by agents. Resources can vary in a number of ways and the testbed is parameterised accordingly. The following may be specified:

- The number of resource patches to create of each type
- The number of instances each resource location comprises. The reader may think of a resource location as an apple tree for example: when an agent acquires an instance of the resource, the number of instances of the resource at the patch is reduced.
- The type of energy an instance of a resource type can supply.
- The quantity of energy an instance of the resource can supply.
- The number of agents which must simultaneously "attack" the resource patch before any of them can acquire a resource instance. This is an attempt to model resource acquisition tasks that require a group effort. When a resource requires the action of more than one agent to be acquired, we shall sometimes refer to it as "complex".

- The distance within which an agent must be before it can be effective in acquiring a resource.
- The time intervals over which a resource periodically renews itself. This is intended to capture the seasonality of resources.

Using the periodicity of the resources and their regional nature, it is possible to set up any pattern of resources and to have their distributions change through time as desired.

27.4.2.2 Agents

In each (simulated) time unit an agent's energy levels are decremented. Agents seek to stay "alive" by acquiring and consuming one or more types of resource which can supply the appropriate energy. If an agent fails to acquire enough new resource energy its energy will eventually fall to zero and it will "die". This goal of surviving is the only pre-determined goal an agent has, all others are derivative.

The agents are designed with a production system architecture (Figure 27.2). Thus each agent has both a production memory and a working memory. The production memory consists of a number of rules of the form "IF condition THEN action". The working memory consists of, possibly many, facts which can change dynamically. If the condition part of a rule is true, i.e. the condition exists as a fact in working memory, then the rule can be considered for execution, meaning that its action part may be executed.

Cognition and action are carried out by production rules. The difference between cognitive and action rules is that cognitive rules only have an effect on an agent's working memory whereas action rules have an effect on the world. Many cognitive rules can be executed per (simulated) time unit but only one action rule.

The working memory of an agent is divided into four main areas:

- *Message buffer* — where incoming messages are stored. It is cleared at the end of each cycle
- *Resource model* — where an agent keeps a record of resources it knows about.
- *Social model* — where an agent stores its beliefs about itself and other agents. Agents start with no knowledge of any particular groups or other agents, but with *generic* representations (concepts) of agents, groups, leaders and followers available within their social models. (We note in passing that this is *not* equivalent to "pre-programming" a decision hierarchy. To deny agents these generic concepts would

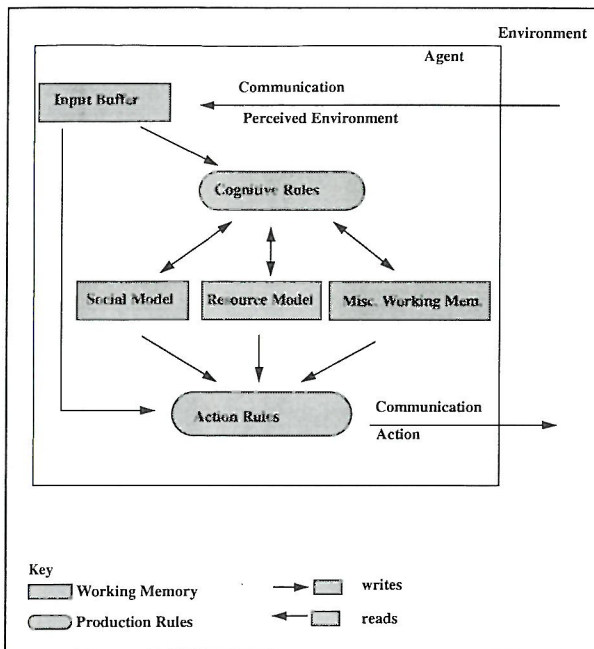


Figure 27.2: Agent Architecture. An agent consists of a set of production rules (cognitive and action) each of which react to a particular set of conditions in its working memory. The working memory contains the agents beliefs about its world including its beliefs about itself and other agents. An agent is aware of the outside world through its input buffer — a part of its working memory which receives incoming messages from the environment, including communication from other agents.

unnecessarily deflect the EOS project into the study of automatic concept learning).

- *Miscellaneous* — where other beliefs are stored, for example information regarding the agent's present state and what it is doing.

The agent parameters which an experimenter must specify are the number of agents of each type, and for each type:

- Locations;
- The speed with which an agent can move, that is, the distance it can move in a time unit;
- Sensory range — that is, how far an agent can detect other agents and resources;
- A list of skills. A resource whose energy type is "a" can only be acquired by an agent with the skill "a";
- A number of energy stores with initial levels;
- An energy level below which the agent is motivated to consume resources;
- A set of rules which comprise the agents' behavioural repertoire;
- The last of these specifications is the most important.

It is important to note that there are no agents in the testbed which directly correspond to groups. A group exists exactly in so far as its members (and also non-member agents) have representations of it within their social models, and behave accordingly.

27.4.3 The testbed interface

A graphical interface provides the user with a plan view of the world with its agents and resources which is updated every scheduler cycle. The interface also provides the ability to interact with the system during experimentation. We plan to add instrumentation facilities, so that the user may quantify or put into statistical form summary information about the dynamics of the multiple agent system.

27.4.4 The agent's behaviour

As stated earlier, it is the production rules which provide the agents with their behavioural repertoire out of which emerges the complex behaviour in which we are interested.

There are about seventy rules, both cognitive and action. For programming convenience and conceptual clarity the rules have been written and organised in such a way that an agent will normally be in one of a number of modes. The rules are divided into those which can only be executed in each of the different modes, those which involve changing mode (meta rules) and those which can work in any mode. The modes include autonomous, bidding, recruiting and executing, and have been designed around a negotiation protocol (Figure 27.3) similar to the Contract Net protocol (Smith, 1980). We briefly describe each mode in turn.

- *Autonomous Mode.* This is the default mode. In this mode an agent acts in an independent non-co-operating way. For example, it can move around at random looking for resources.
- *Bidding Mode.* An agent may go into bidding mode if it receives a message from another agent inviting it to take part in some planned resource acquisition task.
- *Recruiting Mode.* An agent may go into recruiting mode if it believes it can organise a team of agents to perform some resource acquisition task which it has planned.
- *Executing Mode.* An agent goes into executing mode either when it has recruited a full team or if its bid has been accepted. In executing mode, agents go ahead and execute a planned task.

The "planning" referred to above (in "bidding" and "recruiting" modes) involves the instantia-

tion as appropriate of simple plan schemas, implicit in certain production rules, which specify that a number of agents (possibly just one) should jointly act to acquire a particular resource at some time in the future. We anticipate that a more substantial form of planning will be required within agents as our work progresses.

There are rules outside this mode structure some of which, importantly, deal with updating an agent's social model. For example, if an agent observes (by intercepting the appropriate messages) one agent recruiting another, then it will add a group to its social model with both agents as members but with the recruiter as the leader.

The agent behaviour which typically results from these rules is as follows. If an agent is able to collect resources individually then it will do so. If the resources which look the most promising require the co-operation of several agents and an agent has received a request to co-operate then it will do so. If it has not already received a request to co-operate then it will attempt to recruit others to a plan of its own. The priorities associated with each of these types of behaviour, (i.e. "independence" before "bidding" before "recruiting") are plausible and arguably in line with the principle of cognitive economy, but could easily be varied.

Agents which successfully recruit a group to their plan become group leaders, with all members of the group adjusting their social models accordingly. Agents which have become group leaders expect their followers to do as they suggest. Agents which have become followers will normally acquiesce in the plans of their leaders, but will not always do so. For example, a follower agent may in certain circumstances be recruited by a new leader, in which case it will, temporarily at least, reject plans from its old leader.

A leader can receive a request to co-operate involving its whole group rather than just itself and can respond on its group's behalf. It is this step, of course, which leads to the formation of multilevel groups (i.e. hierarchies). Followers have rules which allow them to pass on resource information to their leaders. This has the consequence, that leaders effectively extend their range of perception.

Faulty social (and resource) models, where agents have incorrect beliefs about the groups to which they belong or about other groups, can arise when an agent wrongly infers the existence of a group in the first place or, more likely, when an agent's beliefs are not updated as the group's membership and other properties change. Naturally, false beliefs lead to plans which fail on execution.

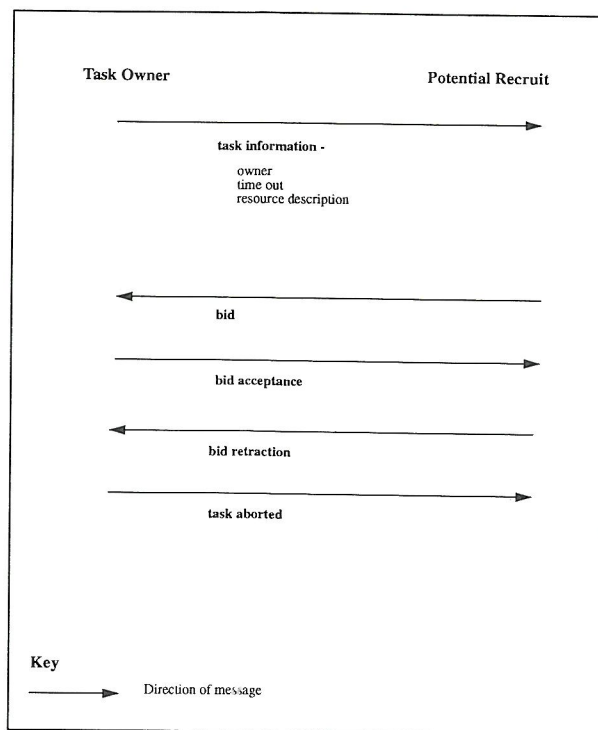


Figure 27.3: Negotiation Protocol. Agents achieve co-operation through a process of negotiation. An agent which is able to plan for the co-operative acquisition of a resource sends out information regarding the task to other agents. These other agents evaluate the information and may decide to bid for a part in the task. If the task owner has received enough bids by some specified time, it will accept those which it favours. An agent whose bid is accepted by a task owner then retracts any other bids it has made. If an agent retracts a bid once the task has been set in motion, the task owner aborts the plan by informing the remaining agents in the task.

27.5 INITIAL EXPERIMENTS

In initial experiments we have demonstrated:

- 1) Agents acting autonomously;
- 2) Crowding — many agents being forced into a small area to acquire patchy resources;
- 3) "Congestion" — agents accidentally interfering with one another's attempts to acquire resources;
- 4) Agents forming groups to acquire complex resources;
- 5) Groups attaining some degree of permanence. Once several agents have co-operated they tend to do so in the future;
- 6) Formation of hierarchical group structures (to several levels of nesting).

Figure 27.4 illustrates some of this collective behaviour. In particular it shows a situation where


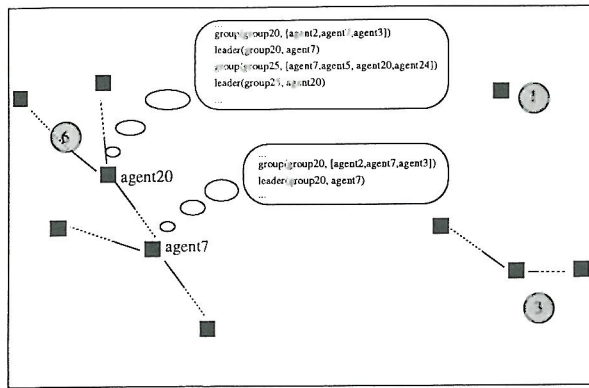


Figure 27.4. Plan view of agents in groups. In particular note how agent20 has formed a 2-level hierarchical group containing 6 agents in order to acquire a complex resource requiring the co-operation of six agents. This has been achieved by agent20 recruiting agent7, the leader of a group of 3, into its existing group of three. Their social models have been modified accordingly (as will have been those of other agents in the vicinity).

the leader of a group of three agents (agent20), in order to acquire a resource requiring co-ordinated action by six agents, recruits another agent (agent7) which is itself the leader of a group of 3 thus forming a two level hierarchy. The social models of both agent20 and agent7 contain references to the fact that agent 20 is the leader of a group of 4 agents one which is agent7, and that agent7 is itself the leader of a group of 3.

In passing we note that inherent in these “emergent” structures are simple forms of ranking and specialisation. For a hierarchical group necessarily implies a natural ranking of the agents within it. Further, once an agent joins a group it will tend repeatedly to be allocated tasks requiring just a subset of its particular skills. This a consequence of leaders only having partial knowledge of their followers.

27.5.1 Territoriality

Having made progress towards understanding and giving computational expression to the processes by which a decision making hierarchy may emerge, it is natural, indeed unavoidable, to address the notion of territoriality. This is an important issue in both Mellars (1985) and Cohen (1985)

and is also implicit in ideas of Gamble (1991; see below).

By territoriality we understand the attempt of particular individuals or groups of individuals to delineate and assert control over particular geographic areas (Sack 1986:19). We follow Dyson-Hudson & Smith (1978:23–24) in the view that territoriality in humans is not innate but is rather a strategy which it is advantageous to adopt in particular ecological situations, specifically when critical resources are abundant and predictable. An important effect of territoriality is cognitive economy — classification by area avoids the need for enumeration and classification by kind (Sack 1986:32).

For us a major issue is whether or not the agents have explicit representations of territories such as can be reasoned about and amended with appropriate cognitive actions. The alternative is for agents to have rules which implement territoriality by reacting directly to some stimulus in the environment. In line with our emphasis on agents’ cognitive and social representations, we assume that the former is the case.

Our task therefore, is to give a computational account of how, and under what conditions, an agent may come to have actual representations of the territories of particular agents and groups, and how these representations affect its behaviour. We assume that an agent has at the outset the generic concept of a territory.

We have implemented the following:

- 1) when an agent perceives some agent (perhaps itself) acquiring some resource it creates an internal (rectangular) representation of the area containing the resource. As new resource acquisitions are observed this representation is enlarged.
- 2) agents create representations of group territories. These are the minimal (rectangular) representations which contain the group members’ territories.
- 3) agents use territorial representations within planning to bound the range of options considered.

We have conducted a small number of experiments involving territoriality. Figure 27.5 shows, in graphical form, the results of an experiment in which the survival of agents with and without the ability to form and use explicit concepts of territories were compared.

Initially there are 10 agents and a number of resource patches containing only one resource instance each. There are just sufficient resources for all agents. The population is run for 1000 cycles.

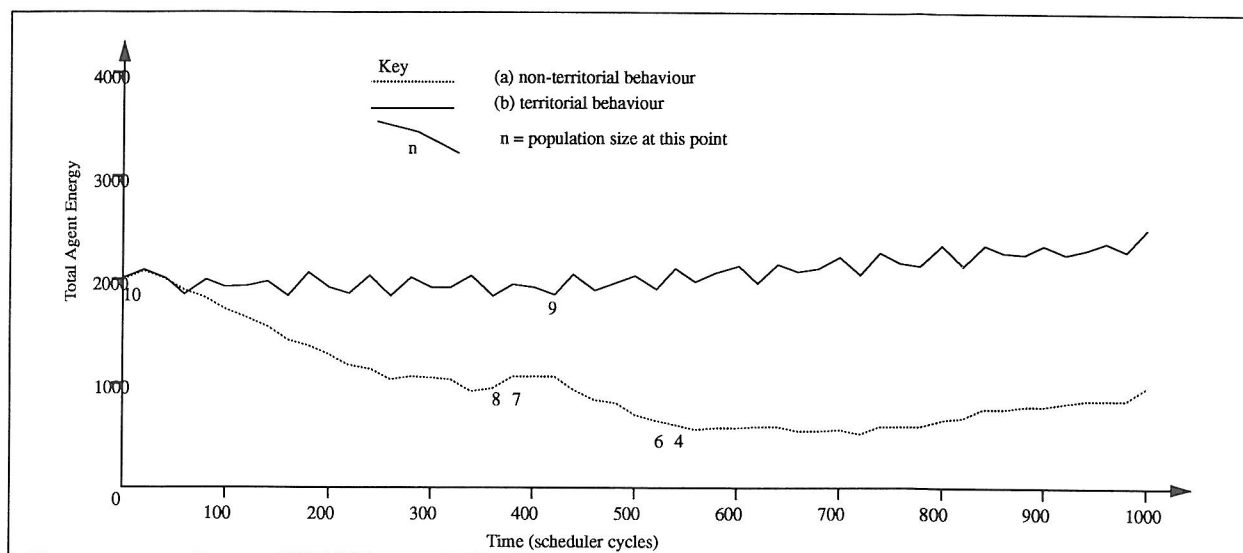


Figure 27.5: Graph showing first 1000 cycles of two experiments contrasting (a) non-territorial and (b) territorial behaviour. In experiment (b) agents form explicit representations of each other's territories — rectangles encompassing the resource locations which agents are observed to collect from. Agents give priority to collecting resources from within their own territories and only collect from the territories of others when there are no resources in their own territory. The graph shows how, without territorial concepts and rules for using them, the agent population as

a whole does less well in terms of the number of agents who survive and the total energy in the system than when such territorial behaviour is introduced. This is because without this behaviour agents often attempt to collect from the same resource location at the same time. If the resource locations are defined in such a way (as they are here) that there is only enough for one agent at any time, then each time there is a conflict only one agent will succeed. In other words disorganisation causes congestion which has an adverse impact on the whole population.

In the case where agents did not have the ability to use the concept of territoriality the agent population quickly falls to less than half its original value with the total energy in the system falling likewise. This happens because several agents often attempt to collect from the same resource patch at the same time. Because each patch contains only one resource instance only one agent can succeed. Some agents fail to acquire sufficient resources and eventually die.

When agents are given the ability to use the concept of territoriality as a factor in deciding where to collect resources only one agent dies and the total energy in system remains at around its initial value. The reason is that territoriality has the effect of reducing conflict, acting as an implicit form of organisation and co-ordination and ensuring that it is rare for several agents to attempt to collect from the same patch at the same time.

27.6 THE GAMBLE MODEL AND ITS RELATIONSHIP TO THE MELLARS MODEL

We now turn to a consideration, from a DAI standpoint, of the relationship between the Mellars model and an important alternative

model, that of Clive Gamble. We hope to show that by identifying the essentials of both models in computational and DAI terms, we can clarify the relationship between them.

Gamble (1986, 1991) has presented an important informal model of social change during the European Palaeolithic generally, and that of S.W. France in particular. He argues that as the climate worsened, people throughout Europe extended their networks of social relations in the form of regional alliances which served as "insurance policies" to provide access to neighbouring resources in times of need (1991:8). He sees the increase in the use of art as reflecting the intensification of the negotiation of alliances, not an increase in the complexity of the social organisation. In regard to South Western Europe in particular, he sees no need to resort to explanations involving population pressure based on affluence or circumscription.

At first sight, the Mellars and Gamble models have little in common. This is partly because their authors are starting from different methodological and theoretical perspectives. Mellars takes as his starting point a detailed consideration of the local archaeological and palaeo-environmental record whereas Gamble prefers to approach the

period from its most distinctive feature, abundant and sophisticated art.

On closer consideration, however, the real points of disagreement between the authors are in fact very limited, turning essentially on:

- 1) the extent to which the resources of the south west were predictable;
- 2) the extent to which the sites of the south west supported large semi-permanent residential groups

We note that both authors take the view that environmental factors probably had a major effect on the social organisation of the local populations. These authors are, in fact, collectively suggesting that two different resource acquisition problems yield the emergence of two correspondingly different types of social structure. Gamble is suggesting alliances in the context of scarce and unpredictable resources. Mellars is suggesting social hierarchies in the context of restricted zones of plentiful and predictable resources.

Our problem, then, is to encompass the two models within a unified DAI framework and so enable direct comparisons to be made between them including experimental comparisons. To do so, we propose two extensions to our existing framework. In overview, we:

- 1) Add to an agent's repertoire of generic concepts that of a "co-operator" agent;
- 2) Elaborate the mechanism whereby instances of generic concepts are formed so that "co-operators" may be identified in reasonable ways.

These two extensions then enable us to relate the types of "social" relationship (and patterns of relationship) that may come into existence between agents to their "objective" circumstances. In greater detail:

Extension I

As previously stated, each agent has, in effect, fixed generic concepts provided within its social model for each of agent, group, leader and follower. It is a small step to add to this set of generic concepts one for "co-operator", with the intention of capturing the notion of habitual symmetric collaboration rather than asymmetric dominance. An agent may then employ instances of this concept to represent instances of habitual co-operation (not necessarily involving itself) and to modify its behaviour accordingly, analogous to the use of the "leader" and "follower" concepts.

Typical behaviour triggered by an agent's regarding itself as within a co-operation relationship would be, for example, transfer of information about resources, and utilisation of an appropriate resource acquisition plan negotiation procedure.

Much hangs, however, on the circumstances in which the generic "co-operator" concept is instantiated, and this issue is addressed in extension II.

Extension II

In the current implementation, instances of the "follower" and "leader" concepts are generated immediately whenever one agent is successfully recruited to take part in a plan proposed by another. Arguably this is much too crude a mechanism. The extension we envisage is that where one agent REPEATEDLY recruits another, then the concepts of dominance will be established, but that where repeated collaboration exists but does NOT show one of the agents as the habitual successful initiator, then both parties will tend to instantiate a co-operation concept.

27.6.1 Objective relations of dependence

The question now is in what circumstances will the situation «one agent repeatedly recruits another» (see above) arise? The answer must involve what Castelfranchi *et al.* (1991) have called the agents' «objective dependence relationships». Indeed, these authors argue (*ibid*:2) that «dependence is undoubtedly the ground relation upon which the whole construction of sociality is based».

Assume there are two agents which can communicate one with another, and that co-operation between them (i.e. a co-ordinated pattern of action) can be mutually beneficial as regards resource acquisition. We see two alternative circumstances in which one agent will repeatedly recruit the other to its own plans, leading to dominance as indicated:

- 1) when the first agent has much greater knowledge than the second of resources, their locations, and how they may be acquired, and
- 2) when the first agent has available to it one or more "potential dominance" actions whose objective effect is significantly to control (positively or negatively) the extent to which the other can acquire resources, without substantial detriment to itself.

We believe that each of these conditions and their consequences can be given a precise computational interpretation. Note that it will surely be

rare for either of them to hold — so that dominance will be unusual within this framework.

The foregoing analysis, sketchy though it is, does seem to illuminate the Mellars and Gamble models, and the contrast between them. The key observation is that neither of the stated conditions are likely to hold in the context of a spatially dispersed population (DAI agents or human groups). But spatial dispersion of relatively small groups is precisely what Gamble envisages, so that co-operation (which between groups becomes alliances) is the pattern of relationship which emerges. Mellars, on the other hand, envisages quite large scale population concentration so that dominance relations become much more feasible (though by no means guaranteed).

This analysis has immediate implications for our computational interpretation of the Mellars model. If dominance relationships are inherently unusual, then we need to capture within the interpretation just how it is that they emerge rather than the usual “co-operation” relationships. This focuses attention of the exact means by which and contexts in which different types of relationship concepts become instantiated in agents’ social models.

27.7 CONCLUSIONS AND FUTURE WORK

We have described work in progress. Although more, and more systematic, experimentation is needed to establish properties of our existing interpretation of the Mellars model reliably, we believe that we have already shown the ability of a DAI-based computational approach to give precise interpretations to such models, and in addition to shed light on the relationship between alternative models.

Further developments of the existing testbed and the processes we have formulated within it are clearly desirable. We mention in particular:

- more powerful cognitive facilities within agents, especially as regards planning and handling of the social model
- implementation of concepts and processes related to co-operation rather than dominance
- implementation of processes by which some kind of simulated archaeological record is generated (c.f. Mithen, 1990), enabling a bridge back to the motivating archaeology to be built.

Acknowledgements

We are grateful to our EOS project colleagues Nigel Gilbert and Paul Mellars for many valuable discussions relevant to the contents of this paper. However, they should in no way be held responsible for errors, omissions or other infelicities within it.

The EOS project is funded by the UK Joint Council Initiative in Cognitive Science/HCI by grant no SPG8930879.

References

- Alexander, R.D.
1989 Evolution of the Human Psyche. in Mellars, P. & Stringer, C. (eds.) *The Human Revolution*. Edinburgh University Press. pp. 455–513.
- Bhaskar, R.
1978 *A Realist Theory of Science*. (2nd edition) Harvester Wheatsheaf.
1986 *Scientific Realism and Human Emancipation*. Verso.
1989 *The Possibility of Naturalism*. (2nd edition) Harvester Wheatsheaf.
- Bloch, M.
1991 Language, Anthropology and Cognitive Science. *Man* (N.S.), 26:183–198.
- Bond, A.H. & L. Gasser (eds.)
1988 *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann.
- Castelfranchi, C., M. Miceli & A.Cesta
1991 Dependence Relations among Autonomous Agents. in Donald D. Steiner & Jurgen Muller (eds.) *Pre-Proceedings of the 3rd European Workshop on Modelling Autonomous Agents and Multi-Agent Worlds*.
- Cohen, M.N.
1985 Prehistoric Hunter–Gatherers: the Meaning of Social Complexity. in T. Douglas Price & James A. Brown (eds.) *Prehistoric Hunter–Gatherers: the Emergence of Cultural Complexity*. Academic Press. pp. 99–119.
- Doran, J.E.
1990 Computer-based Simulation and Formal Modelling in Archaeology: a Review. in A. Voorrips (ed.) *Mathematics and Information Science in Archaeology: a Flexible Framework*. Studies in Modern Archaeology, Vol 3. HOLOS-Verlag, Bonn. pp. 93–114.
- Doran, J.E., H. Carvajal, Y.J. Choo & Y. Li
1991 The MCS Multi-Agent Testbed: Developments and Experiments. in S.M. Deen (ed.) *Cooperating Knowledge Based Systems 1990. Proceedings of the International Working Conference on Cooperating Knowledge Based Systems*, University of Keele, October 1990. Springer-Verlag, pp. 240–251.

- Doran, J. & M. Palmer
in press The EOS Project: Modelling Prehistoric Socio-cultural Trajectories. due to appear in the *proceedings of the First European Symposium on Computing in Archaeology, held at Saint-Germain-en-Laye, France, November 1991.*
- Dyson-Hudson, R. & E A. Smith
1978 Human Territoriality: An Ecological Reassessment. *American Anthropologist*, 80:21-41.
- Gamble, C.S.
1986 *The Palaeolithic Settlement of Europe*. Cambridge University Press, Cambridge.
- 1991 The Social Context of European Palaeolithic Art. *Proceedings of the Prehistoric Society*, 57, part I:3-15.
- Johnson, A. & T.K. Earle
1987 *The Evolution of Human Societies: from Forager Group to Agrarian State*. Stanford University Press.
- Mellars, P.A.
1985 The Ecological Basis of Social Complexity in the Upper Paleolithic of Southwestern France. in T. Douglas Price & James A. Brown (eds.) *Prehistoric Hunter-Gatherers: the Emergence of Cultural Complexity*. Academic Press. pp. 271-297.
- Mithen, S.
1990 *Thoughtful Foragers. A Study in Prehistoric Decision Making*. Cambridge University Press.
- Renfrew, A.C.
1987 Problems in the Modelling of Socio-cultural Systems. *European Journal of Operational Research*, 30:179-192.
- in press Dynamic Modelling in Archaeology: What, When and Where?. in S. van der Leeuw (ed.) *Proceedings of the Conference on Dynamical Modelling and Human Systems, Cambridge, December 1990*. To be published by Edinburgh University Press.
- Sack, R.D.
1986 *Human Territoriality*. Cambridge University Press.
- Smith, R.G.
1980 The Contract Net Protocol: High Level Communication and Control in a Distributed Problem Solver. *IEEE Transactions on Computers*, C-29(12):1104-1113.
- Sayer, A.
1992 *Method in Social Science: A Realist Approach*. (2nd edition) Routledge.
- Steels, L.
1990 Cooperation Between Distributed Agents Through Self-Organisation. in Yves Demazeau & Jean-Pierre Muller (eds.) *Decentralised AI: Proceedings of the 2nd European Workshop on Modelling Autonomous Agents and Multi-Agent Worlds*. Elsevier Science Publishers B.B. (North-Holland). pp. 175-196.
- 1991 Towards a Theory of Emergent Functionality. in *From Animals to Animats*. MIT Press. pp. 451-461.
- Wavish, P.
1991 Exploiting Emergent Behaviour in Multi-Agent Systems. in Donald D. Steiner & Jurgen Muller (eds.) *Pre-Proceedings of the 3rd European Workshop on Modelling Autonomous Agents and Multi-Agent Worlds*.

Authors' addresses

Jim Doran
Department of Computer Science
University of Essex
Colchester
GB-CO4 3SQ
E-mail: doraj@essex.ac.uk

Mike Palmer
Department of Computer Science
University of Essex
Colchester
GB-CO4 3SQ
E-mail: palmm@essex.ac.uk