Gilbert, Nigel. (1995c). Emergence in social simulation. In N. Gilbert & R. Conte (Eds.), *Artificial societies: The computer simulation of social life* (pp. 144-156). London: UCL Press.

Chapter 8

Emergence in social simulation

Nigel Gilbert

Every discipline is based on a unique foundation of epistemological assumptions and concepts. This means that even when one discipline develops so that it begins to share its concerns with another, there may be little or no contact because the practitioners are, literally, speaking different languages. Even if contact is established, the neighbouring disciplines may still have nothing to say to each other because, while a topic may be common, the questions being asked and what count as interesting answers differ so greatly. To bridge the gap, considerable work in translation has to be done.

This situation seems to be developing in the case of artificial intelligence (AI) and the much older discipline of sociology. Both are interested ostensibly in the behaviour, structure and properties of collections of actors or agents. Yet the literature on distributed artificial intelligence (DAI) makes little or no reference to sociology, or even to its major concepts. Similarly, there are very few sociologists who have considered the relevance of DAI to their concerns.

In this chapter, I shall show first that sociologists have also struggled with one of the basic conceptual problems that those interested in simulating societies have encountered: the problem of understanding "emergence" and, especially, the relationship between the local and global properties of complex systems. Secondly, I shall indicate some ways in which DAI simulations may have oversimplified important characteristics of specifically human societies, because the actors (agents) in these societies are capable of reasoning, and do so routinely, about the emergent properties of their own societies. This adds a degree of reflexivity to action which is not present (for the most part) in societies made up of simpler agents, and in particular is not a feature of current DAI simulations.

The macro and the micro level

The issues I shall address stem from the fact that both societies and the computational systems with which distributed AI are concerned are composed of many interacting agents (also known as people, actors or members). These systems can therefore be described either in terms of the properties and behaviour of the agents, or in terms of the system as a whole. The former mode of description focuses on the "micro" level, that is, the features of individual agents and their local environment (which they can perceive directly), while the latter focuses on the "macro" level, that is, the global patterns or regularities formed by the behaviour of the agents taken as a whole. The problem is how to characterize the relationship between these two types of description, in particular when the macro properties can be said to "emerge" from the micro-level behaviour.

Sociological approaches to emergent properties

Methodological holism

The relationship between local and global properties is one which has exercised sociologists since the foundation of the discipline. For example, Emile Durkheim, one of its founding fathers, emphasized the external nature of social institutions (a macro-property of a society) and argued that they imposed themselves on individuals at the micro-level. It is worth quoting at some length from his methodological writings, because he was addressing issues very similar to those which continue to worry people who construct social simulations. After having defined a category of facts that he called "social facts", he wrote

> Another proposition has been . . . vigorously disputed . . . it is the one that states that social phenomena are external to individuals . . . The states of the collective consciousness are of a different nature from the states of the individual consciousness: they are representations of another kind. The mentality of groups is not the mentality of individuals; it has its own laws . . . To understand the way in which a society conceives of itself and the world that surrounds it, we must consider the nature of society, not the nature of the individuals. (Durkheim 1895: 65–6)

Thus, for Durkheim, there are social representations which can be examined independently of the individuals that make up the society. These social facts are

> a category of facts with very special characteristics: they consist of ways of acting, thinking and feeling that are external to the individual and are endowed with a coercive power by virtue of which they exercise control over him. (Durkheim 1895: 69)

This view of the relationship between an individual and society was later developed by Parsons (1952) into a "normative conception of order" which emphasized the role of internalized norms in ensuring the integration of society through shared values and obligations. Subsequent theorists have extended this tradition, which has become known as "methodological holism" (O'Neill 1973). It asserts that the social behaviour of individuals should be explained in terms of the positions or functions of these individuals within a social system and the laws which govern it.

Methodological individualism

In contrast, "methodological individualists" see macro-phenomena accounted for by the micro-level properties and behaviour of individuals. The individualists' position demands that all the concepts in social theory are analyzable in terms of the interests, activities, etc. of individual members of society.

> If social events such as inflation, political revolution, "the disappearance of the middle classes", etc. are brought about by people, then they must be explained in terms of people; in terms of the situations people confront and the ambitions, fears and ideas which activate them. In short, large scale-social phenomena must be accounted for by the situations, dispositions and beliefs of individuals (Watkins 1955: 58).

A rather sterile debate between these two camps continued for much of the 1970s and 1980s. With the benefit of hindsight, it is now possible to argue that while there was some truth in both, neither was a particularly helpful or revealing way of conceiving the relationship between global and local behaviour. It is the case, however, that most, if not all, current simulations of human societies in essence adopt one or other of these positions, often without making this explicit.

Structuration theory

Perhaps the most influential recent sociological approach to the relationship between macro and micro is the theory of structuration. The theory argues that there is a duality between society and knowledgeable human agents. Agents are seen as reproducing in action the structuring properties of society, thereby allowing social life to be reproduced over time–space (Giddens 1984). This rather dense, and possibly even opaque, statement of the theory needs some considerable explanation.

First, by "structuring properties", Giddens means things such as institutional practices, rules and resources. These properties are the means by which social structure or "society" is produced and reproduced. Human action has these structuring properties "embedded" in it, so, as people act, they contribute to the reproduction of society. Also ("reflexively" according to Giddens), human action is both constrained and enabled by social structure, for this is the medium through which action is performed. Structure is at the same time both the outcome of knowledgeable human conduct and the medium that influences how conduct occurs (Giddens 1984: 25).

Complex adaptive systems

We can now consider whether structuration theory can be linked to or has anything to offer computational views of emergence in simulated societies. The question of emergence has probably been considered most systematically by those interested in complex systems (Stein 1989, Jen 1990, Stein & Nadel 1991). Such systems have the following general characteristics (Forrest 1990):

1. The system consists of a large number of interacting agents operating within an environment. Agents act on and are influenced by their local environment.
2. There is no global control over the system. All agents are only able to influence other agents locally.
3. Each agent is driven by simple mechanisms, typically condition–action rules, where the conditions are sensitive to the local environment. Usually, all agents share the same set of rules, although because they may be in different local environments, the actions they take will differ.

A typical example of such a complex system is an array of cellular automata. Each cell can take on one of a number of internal states. State changes occur according to the results of the firing of simple condition–action rules dependent only on the states of neighbouring cells.

Some complex systems are also adaptive. In such systems, agents either learn or are subjected to a process of mutation and competitive selection, or both. In learning systems, agents are able to modify their rules according to their previous "success" in reacting to the environment. In systems involving competitive selection, agents' rules are altered, either at random or by using some specific learning algorithm, and then those which are more successful are copied while the less successful are deleted from the system. An important characteristic is that agents adapt within an environment in which other similar agents are also adapting, so that changes in one agent may have consequences for the environment and thus the success of other agents. The process of adaptation in which many agents try to adapt simultaneously to one another has been called "co-evolution" (Kauffman 1988).

Emergence in complex systems

Because complex systems, whether adaptive or not, consist of many agents, their behaviour can be described either in terms of the actions of the individual agents or at the level of the system as a whole. In some system states the global description may be very simple (e.g. if the agents are either not interacting or interacting in repetitive cycles, the global description might be that "nothing is happening"), or exceedingly complex (e.g. if the agents are in complete disequilibrium). In some circumstances, however, it may be possible to discover a concise description of the global state of the system. It is in these latter circumstances that it becomes possible to talk about the "emergence" of regularities at the global level.

An example relevant to social simulation can be found in the work of Nowak and Latané (e.g. 1993) which aimed to model the effect of social influence. This draws on Latané's theory of social impact which states that the impact of a group of people on an individual's opinion is a multiplicative function of the persuasiveness of the members of the group, their social distance from the individual and the absolute number of the members. For example, in a simulated world in which there are only two mutually exclusive opinions, the opinion adopted by an individual is determined by the relative impacts on the individual of those propounding the two opinions. At any moment in time, an agent's opinion is determined by the multiplicative rule that implements the theory of social impact, and is dependent only on the opinions of other agents. The simulation fits the definition of a complex system.

The outcome of Nowak and Latané's simulations is that coherent clusters of opinion emerge and remain in dynamic equilibrium over a wide range of assumptions for the parameters of the model. The population does not move wholly to adopt one opinion, and minority views are not completely expunged. Nowak and Latané show that the clusters of individuals with the same opinion can be visualized as patches on a two-dimensional grid in which the nodes represent the members of the population. These clusters can be said to have emerged and to form regularities.

In systems in which there is emergent behaviour, it is convenient to think of the emergent properties of the system as influencing the actions of the agents. Thus, not only do the agents' actions at the local level, when aggregated and observed at the global level, constitute the emergent behaviour, but also the global emergent behaviour can also be said to influence the local actions of the agents, in a form of feedback. For example, in Nowak and Latané's simulation, one can speak of an opinion cluster influencing the opinions of individuals outside the cluster. But it must be remembered that in the standard complex system, there is no explicit mechanism for such feedback to take place: agents are affected only by their local neighbours. The idea of feedback, or "feed-down", from global to local level is merely a convenient shorthand.

Levels of emergence

So far, for simplicity, it has been assumed that there are clearly distinct "levels": the micro and the macro. In fact, the situation is more complex than this. For example, it may be the case that individual identity is best regarded as an emergent phenomenon, where the micro-level "agents" are sub-cognitive, such as neurons. And societies are perhaps best considered as emergent phenomena arising from the interaction of social institutions. Thus it is better to consider a complex hierarchy of levels of emergence, rather than a straightforward division between micro and macro.

A continuing worry of those who are interested in emergent systems, and particularly of those who are simulating social process, is whether the emergent behaviours they observe are in some sense "programmed in" and an inevitable consequence of the way in which agents are constructed. Some criterion is required which will distinguish emergent behaviour from behaviour which is predictable from the individual characteristics of the agents. The description of complex systems suggests that a candidate criterion is that it should not be possible to derive analytically the global emergent behaviour solely from consideration of the properties of agents. In other words, emergent behaviour is that which cannot be predicted from knowledge of the properties of the agents, except as a result of simulation. This kind of definition, of course, provides an obvious explanation for the popularity of simulation as an investigative technique in the field of complex adaptive systems.

This criterion, however, is not sufficient to end the worries of those who are concerned about whether they have *really* achieved emergence. For it is always possible that at some future time analytical methods will be developed that can be used to derive global properties from local ones. In a few cases, this seems to have happened. For example, Forrest and Miller (1990) define a mapping between classifier systems (i.e. systems consisting of nodes that use the genetic algorithm to synthezise patterns of connections that collectively co-evolve to perform some function) and Boolean networks whose properties are, relative to classifier systems, much better understood. This mapping opens the way to predicting the emergent properties of classifier systems where previously predictions were impossible. The example shows that if we define emergence in terms of an inability to find an analytical solution, any particular emergent property stands the risk of being demoted from the status of emergence at some time in the future. This suggests that emergence may be neither a stable nor an especially interesting property of complex systems: what are interesting are the systems' macro properties and the relationship of those macro properties to the micro ones.

Complex adaptive systems and human societies

The above account of complex adaptive systems is generic, in the sense that it is intended to apply equally to many disparate domains. For example, there has been work applying these ideas to the evolution of antibodies, stock market crashes, ecological dynamics, computer networks and so on (Waldrop 1992). However, I am particularly concerned with understanding human societies. Can these accounts be applied unchanged to human societies, or is there something special about these societies which would indicate that such models are completely inappropriate, or must be amended or developed in order to provide insights into their structure and functioning?

Debating whether there is a boundary between humans and other animals and, if so, its nature, is a matter of some controversy, for example with the argument that great apes deserve the same moral status as humans (Vines 1993). The question of whether human societies differ in any fundamental way from non-human societies threatens to become mired in the same moral debate. Nevertheless, it is clear that, while an exact boundary line may be difficult to draw and some characteristics of human societies are also to be found to some degree in some animal societies, human societies are significantly different from any animal society. Does this also mean that present theories of complex adaptive systems, while perhaps useful for understanding non-human societies, are lacking in some significant way as models of human societies? I believe that the answer to this question is that such theories do at present fail to model a crucial feature of human societies, and that the preceding discussion, especially the notion of structuration, can help to identify what this is and how it may be corrected.

A fundamental characteristic of humans, one that is so important that it makes the societies they form radically different from other complex systems, is that *people are routinely capable of detecting, reasoning about and acting on the macro-level properties (the emergent features) of the societies of which they form part.* The overall lesser intelligence of animals prevents this occurring in non-human societies. The consequences of this capability to "orientate" to macro-level properties are far-reaching.

Our discourse is permeated by references to the macro-level properties of our society. For instance, we are not only able to speak a language using a common lexicon in which words have shared[1] meaning, but we are also able to describe and comment on our lexicon; that is, as members of a human society, we can observe and reason about this macro-level property of our society.

Let us now consider another example, one where a current simulation has been examining a theory relating to human society, and which tries to provide explicitly a relatively sophisticated approach to modelling human sociality, and compare its implicit characterization of the relationship between macro and micro with sociological accounts that theorize society in terms of structuration.

Simulating human institutions

Institutions are among the most pervasive emergent features of human society. An institution is an established order comprising rule-bound and standardized behaviour patterns. Examples include the family, tribes and other collectivities, organizations, legal systems and so on. Sociological approaches to institutions can be considered in terms of the three positions described at the beginning of this paper: methodological holism, individualism and structuration theory.

In brief, the holists' position sees individuals being socialised into membership of an institution (for example, the family, or the nation) by parents and teachers. The socialization process consists of the transmittal of the rules and norms of behaviour of the institution to new members. These rules and norms then guide members in institutional action. In contrast, the individualists see institutions as arising from negotiation among the members, each pursuing his/her interests. An approach from structuration theory might focus on the way in which individual action and interaction in everyday life generates institutions through rationalizations of action, in which the institution itself becomes acknowledged by individuals as part of the process of accounting for their actions. As Bittner (1974), although himself not an advocate of structuration theory, put it with reference to a particular kind of institution, the organization: "Organisational designs are schemes of interpretation that users can invoke in as yet unknown ways whenever it suits their purposes" (Bittner 1974: 77). In other words, institutions, as macro-level features, are recognized by members, who use them to warrant their own actions, thus reproducing the same features.

The EOS simulation

We can now examine an example of a simulation based on DAI principles to see whether it fits neatly into any of these theoretical perspectives on the relationship between macro and micro. I have been associated with the EOS (Emergence of Organised Society) project since its inception, although Jim Doran and Mike Palmer are the people who have done all the work (Doran et al. 1994, Doran & Palmer, Chapter 6 in this volume). The objective of the project is to investigate the growth of complexity of social institutions during the Upper Palaeolithic period in south-western France, when there was a transition from a relatively simple egalitarian hunter-gatherer society to a more complex one, with centralized decision-making, specialists and role differentiation, territoriality and ethnicity (Cohen 1985, Mellars 1985, Doran 1992). The project is intended to examine a particular theory about the causes of this transition, one which takes seriously the effect of human cognitive abilities and limitations. Consequently, the simulation that is at the heart of the EOS project is based on a computational model that incorporates a number of agents having AI capabilities that enable them to reason about and act on their (computational) environment. The agents, implemented using a production system architecture, are able to develop and store beliefs about themselves and other agents in a "social model". Doran et al. (1994: 204) describe the social model thus:

> The social model is particularly important. It is designed to record an agent's beliefs about the existence of groups of agents and the identities of their leaders. Associated with beliefs about leaders is other information, for example, the size of its immediate following. Some followers may themselves be leaders. An agent may well appear in its own social model as a leader or follower in particular groups. Also held within the social model are beliefs about territories. There is no expectation that the information held in an agent's social model will be either complete or accurate.

The social model is important because it has an effect on how an agent behaves in relation to other agents. For example, an agent's membership of a semi-permanent group depends on the agent becoming aware of its membership of the group and then treating fellow members differently from non-members.

Agents start without any knowledge of groups or other agents, but with a need to obtain a continuing supply of "resources", some of which can only be secured with the co-operation of other agents. If an agent receives a request to co-operate with the plans of another agent, it will do so, or it will attempt to recruit agents to a plan of its own. Agents that successfully recruit other agents become group leaders and the other agents become followers. Such temporary groups continue to work together, with group membership recorded in the members' social models, unless other processes (such as the agent's movement out of the proximity of other members) intervene.

A number of observations can be made about this simulation. First, in terms of the definition offered previously, it is a complex system, although not an adaptive one. Secondly, group membership is an emergent property of the behaviour of the agents. Thirdly, agents pursue their own interests to obtain resources, forming "groups" as and when necessary through negotiation with other agents. These groups have no existence other than their representation in agents' social models.

In particular, the agents, even though their membership in a group is recorded in their social model, have no awareness of the group as an entity in its own right with which they can reason. Agents not in a group have no way of recognising that the group exists, or who its members are. In short, the simulation permits agents to reason about micro-properties that they perceive in their local environment, and from this reasoning and the consequent action, macro-level properties emerge, but because the agents have no means of perceiving these macro-properties, the model as a whole fails to match the capabilities of human societies.

Structuration theory begins to suggest what additional functionality could be built into the agents. It proposes that actions are sedimented into "structuring properties", such as institutional practices, rules and typifications. Thus it would be necessary for the EOS agents to *recognise* that they are members of groups and to realize what this implies about members' actions; and do so explicitly and in a way that results in a representation they can subsequently reason with. Their own actions, as group members, would reinforce ("reproduce") these explicit representations of the group and its practices.

In the present simulation, a "group" emerges from agents' actions, but it remains an implicit property of the agents' social models. As would be expected from a simulation that adopts the individualistic approach implicitly, the existence of the "group" can be observed by the researcher either by looking at the patterns of interaction at the global level or by "looking inside the heads" of the agents, but it is not visible to the agents themselves. An alternative way of implementing the simulation, rightly rejected by the research team, would have been to build the notion of a group into the system at a global level so that the agents' actions were constrained by their group membership. In this case, which corresponds roughly to the holists' position, the group would not have been an emergent property of the simulation, but a constraining "social fact".

A third possibility is to design the simulation so that agents themselves are able to detect the emergence of groups and react ("orientate") to this structure, so that they become no less privileged in this respect than the observer. This latter course is the approach that might be advocated by structuration theorists.

Conclusion

This chapter originated from a feeling that, despite the sophistication of the present-day work, the models being used to simulate human societies are lacking in crucial aspects. Few computational models of human societies have supported features which could be argued to correspond to "meaning", "structure" or "action", although to sociologists, concepts such as these are central to understanding human society. One of the suggestions of this chapter is that more sophisticated models of human interaction are essential if we are to progress with the simulation of human societies. Human societies are not just ant hills writ large. However, pinning down clearly, and in a computationally precise way, the nature of the difference between non-human and human societies is not straightforward.

The second suggestion of this chapter is that sociological debates about the nature of human society can contribute to our understanding of these issues, even if a certain amount of effort is required to "translate" these debates into forms that are accessible to computational modelling. We have seen that one theory – structuration – provides some useful suggestions about how the relationship between macro- and micro-level properties might be conceptualized and even some broad indications about how a simulation might be implemented. It must be noted, however, that structuration theory is certainly not the only position on these issues to enjoy respect among sociologists, and that other present-day ideas may be equally, or more, fruitful.

The third suggestion is that one capacity of humans which is probably not widely shared by other animals is an ability to perceive, monitor and reason with the macro-properties of the society in which they live. It seems likely that this ability is related to, and is perhaps even a consequence of, the ability of humans to use language. Much of the lexicon of modern adults is, for example, composed of concepts about macro-properties.

The fourth and final suggestion, lurking behind the rest of the argument, is that human societies are dynamic; that is, they are (re-)produced by people as they proceed through their everyday life. Routines, typifications, structures, institutions and so on, are not reflections of some external reality but are human constructions. In particular, the "macro-properties" that are recognized by members of human societies are not fixed by the nature of the agents (as they would be in other complex systems), but are an accomplishment, the outcome of human action.

Acknowledgements

I am grateful to the EOS project team, especially Mike Palmer and Jim Doran, and the participants at SimSoc '92 and '93 for opportunities to discuss these issues.

Notes

1. I am simplifying here; meanings are indexical and negotiated in the context of use in a way which we do not yet properly understand, at least from a computational perspective. See Hutchins and Hazlehurst, Chapter 9 in this volume.