Computational Social Science: Agent-based social simulation

Nigel Gilbert

Centre for Research on Social Simulation
University of Surrey
Guildford
UK

mailto:n.gilbert@surrey.ac.uk

6 November 2005; revised and updated 20 May 2007

While the idea of computer simulation has had enormous influence on most areas of science, and even on the public imagination through its use in computer games such as SimCity, it took until the 1990s for it to have a significant impact in the social sciences. The breakthrough came when it was realised that computer programs offer the possibility of creating 'artificial' societies in which individuals and collective actors such as organisations could be directly represented and the effect of their interactions observed. This provided for the first time the possibility of using experimental methods with social phenomena, or at least with their computer representations; of directly studying the emergence of social institutions from individual interaction; and of using computer code as a way of formalising dynamic social theories. In this chapter, these advances in the application of computer simulation to the social sciences will be illustrated with a number of examples of recent work, showing how this new methodology is appropriate for analysing social phenomena that are inherently complex, and how it encourages experimentation and the study of emergence.

Social simulation

The construction of computer programs that simulate aspects of social behaviour can contribute to the understanding of social processes. Most social science research either develops or uses some kind of theory or model, for instance, a theory of cognition or a model of the class system. Generally, such theories are stated in textual form, although sometimes the theory is represented as an equation (for example, in structural equation modelling). A third way to express theories is as computer programs. Social processes can then be simulated in the computer. In some circumstances, it is even possible to carry out experiments on artificial social systems that would be quite impossible or unethical to perform on human populations.

An advantage of using computer simulation is that it is necessary to think through one's basic assumptions very clearly in order to create a useful simulation model. Every relationship to be modelled has to be specified exactly. Every parameter has to be given a value, for otherwise it will be impossible to run the simulation. This discipline also means that the model is potentially open to inspection by other researchers, in all its detail. These benefits of clarity and precision also have disadvantages, however. Simulations of complex social processes involve the estimation of many parameters, and adequate data for making the estimates can be difficult to come by.

Another benefit of simulation is that, in some circumstances, it can give insights into the 'emergence' of macro level phenomena from micro level actions. For example, a simulation of interacting individuals may reveal clear patterns of influence when examined on a societal scale. A simulation by Nowak & Latané (1994), for example, shows how simple rules about the way in which one individual influences another's attitudes can yield results about attitude change at the level of a society. A simulation by Axelrod (1995) demonstrates how patterns of political domination can arise from a few rules followed by simulated nation states. Schelling (1971) used a simulation to show that high degrees of residential segregation could occur even when individuals were prepared to have a majority of people of different ethnicity living in their neighbourhood.

Figure 1: The pattern of clusters that emerge from Schelling's model
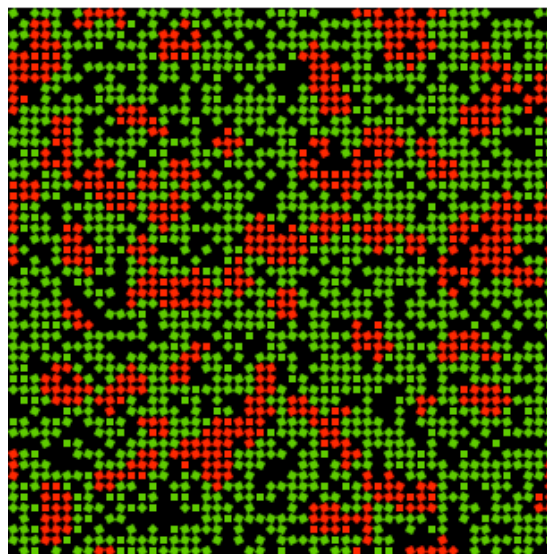
Schelling's study is a good illustration of the kind of work involved in simulation. He modelled a neighbourhood in which homes were represented by squares on a grid. Each grid square was occupied by one simulated household (in Figure 1, either a green or a red household), or was unoccupied (black). When the simulation is run, each simulated household in turn looks at its eight neighbouring grid squares to see how many neighbours are of its own colour and how many of the other colour. If the number of neighbours of the same colour is not sufficiently high (for example, if there are fewer than three neighbours of its own colour), the household 'moves' to a randomly chosen unoccupied square elsewhere on the grid. Then the next household considers its neighbours and so on, until every household comes to rest at a spot where it is content with the balance of colours of its neighbours.

Schelling noted that when the simulation reaches a stopping point, where households no longer wish to move, there is always a pattern of clusters of adjacent households of the same colour. He proposed that this simulation mimicked the behaviour of whites fleeing from predominantly black neighbourhoods, and observed from his experiments with the simulation that even when whites were content to live in locations where black neighbours were the majority, the clustering still developed:

residential segregation could occur even when households were prepared to live among those of the other colour.

Sociology and complexity

The physical world is full of systems that are linear or approximately linear. This means that the properties of the whole are a fairly simple aggregation of the parts. For example, the properties of even such a massive system as a galaxy, with hundreds of millions of component stars, can be predicted precisely using the basic equations of motion. The same applies to aggregations on an atomic and molecular level.

Societies, in particular human societies, are different. They seem to have rather unpredictable features, meaning that it is perilous to make exact predictions of their future development, and their characteristics at any one time seem to be affected by their past histories. For example, the adoption of one of a pair of alternative technologies within a society can be greatly influenced by minor contingencies about who chooses which technology at an early stage in their introduction (see Arthur 1989). This is known as 'path dependence'. It is a sign that human societies, institutions and organisations are complex systems, using 'complex' in the technical sense to mean that the behaviour of the system as a whole cannot be determined by partitioning it and understanding the behaviour of each of the parts separately, which is the classic strategy of the reductionist physical sciences.

One reason why human societies are complex is that there are many, non-linear interactions between their units, that is between people. The interactions involve the transmission of knowledge and materials that often affect the behaviour of the recipients. The result is that it becomes impossible to analyse a society as a whole by studying the individuals within it, one at a time. The behaviour of the society is said to 'emerge' from the actions of its units. There are many examples of emergence in social systems; indeed, it may be that almost all significant attributes of social systems are emergent. For example, markets emerge from the individual actions of traders; religious institutions emerge from the actions of their adherents; and business organisations emerge from the activities of their employees, in addition to the actions of groups such as legislators, lawyers, advertisers and suppliers. We can say that a phenomenon is emergent when it can only be described and characterised using terms and measurements that are inappropriate or impossible to apply to the component units. For example, we can identify the creed of a church or the mission of an organisation, but these terms fit uncomfortably, if at all, when applied to individual people.

While emergent phenomena can also be found in physical systems, a feature of human societies that makes then unique is that people can recognise (and therefore respond to) the emergent features (Gilbert 1995). For example, households not only often cluster in segregated neighbourhoods, but these neighbourhoods are named and can acquire reputations that further affect the behaviour of those living there and others such as employers who may stereotype the inhabitants.

Another important characteristic of societies is that they are result of dynamical processes. The individuals within a society are constantly 'in motion': talking, listening, doing. Society emerges from this constant change. Like a waterfall that exists only so long as the water of which it is formed is moving, a society only exists while its members are living, acting and reacting. Moreover, the units from which societies are formed, that is, people, vary greatly in their capabilities, desires, needs and knowledge, in contrast to most physical systems that are composed of similar or

identical units. For these reasons, while theories of complexity developed for the understanding of natural systems can be illuminating, caution needs to be exercised in applying them directly to social phenomena.

Data

In order to understand complex, dynamical societies, we need appropriate data to base analyses on. Unfortunately, acquiring such data is very hard. The traditional methods of analysis in sociology have been to gather qualitative data from interviews, observation or from documents and records, and to carry out surveys of samples of people. While qualitative data can illustrate very effectively the emergence of institutions from individual action, because of the nature of the data most analyses inevitably remain somewhat impressionistic.

More precision is apparently provided by studies based on quantitative data, but typical survey data has severe limitations if we take seriously the idea that societies are complex and their feature are emergent. Survey data, with just a few exceptions, treats individuals as isolated 'atoms' and pays little attention to the impact of people's interactions with others. The exceptions are data intended for studies of social networks, where respondents are asked about whom they communicate with, are friends with, and so on. However, it is difficult to make such sociometric surveys representative. The result is that much quantitative sociology is based on data that are inappropriate for understanding social interactions. This difficulty applies equally to the study of other complex systems. For instance, if you want to understand the biology of a mouse, taking a random sample of a small proportion of its cells and studying these is unlikely to improve greatly one's knowledge of the mouse's structure and function.

Another problem with much sociological quantitative data, including most surveys, is that they come from measurements made at a one moment in time. But this makes the way in which individuals change, and the effect of these changes, almost invisible to the analyst. Asking 'retrospective' questions about the respondents' past can help, but the answers will inevitably be coloured by their present situation. What are needed are data that track individuals through their life course. Such data are starting to become available with large-scale panel studies, but they are very expensive to collect and still limited in scope.

These limitations of conventional sociological data are fairly well known. The problem is overcoming them. A completely different approach is to build simulation models corresponding to one's theories about society and then to test these against data. In contrast to the inductive methodology of collecting data and then building models that describe and summarise those data, this approach starts from a more deductive perspective. A model is created, calibrated from whatever data is available and then used to derive testable propositions and relationships. The advantage of this approach is that it places much lower demands on the data, while the models can truly reflect the complex nature of societies.

Multi-agent models

A multi-agent model consists of a number of software objects, the 'agents', interacting within a virtual environment. The agents are programmed to have a degree of autonomy, to react to and act on their environment and on other agents, and to have goals that they aim to satisfy. In such models, the agents can have a one-to-one correspondence with the individuals (or organisations, or other actors) that exist

in the real social world that is being modelled, while the interactions between the agents can likewise correspond to the interactions between the real world actors. With such a model, it is possible to initialise the virtual world to a preset arrangement and then let the model run and observe its behaviour. Specifically, emergent patterns of action (e.g. 'institutions') may become apparent from observing the simulation.

Agents are generally programmed using either an object-oriented programming language or a special-purpose simulation library or modelling environment, and are constructed using collections of condition-action rules to be able to 'perceive' and 'react' to their situation, to pursue the goals they are given, and to interact with other agents, for example by sending them messages. The Schelling model described above is an early and simple example of a multi-agent model. Agent-based models have been used to investigate the bases of leadership, the functions of norms, the implications of environmental change on organizations, the effects of land-use planning constraints on populations, the evolution of language, and many other topics.

While most agent-based simulations have been created to model real social phenomena, it is also possible to model situations that could not exist in our world, in order to understand whether there are universal constraints on the possibility of social life. (For example, can societies function if their members are entirely self-interested and rational?) These are at one end of a spectrum of simulations ranging from those of entirely imaginary societies to those that aim to reproduce specific settings in detail.

When an agent-based model has been constructed, it can be run in order to generate output that can be validated against readily observable data. For example, Schelling proposed that local processes of choice about residential domicile, influenced by individuals' and households' perception of the ethnicity of other households within the immediate neighbourhood, would result in residential segregation and yield emergent patterns of clustering, or in the extreme, ghettos with a preponderance of residents of the same ethnicity. It is extremely difficult to gather useful data about individual residential choices, which are made only occasionally and at different times by different residents. A survey of households would only pick up people's retrospective justifications for the decisions that they made about housing choice, which may rationalise decisions based on criteria that they no longer remember clearly. Attempts to tap current attitudes to, for example, living near neighbours of a different ethnicity may also be subject to many kinds of bias. However, in contrast, it is easy to measure people's actual household location and their ethnicity. It is then possible to compare this with the clustering observed after running a Schelling model (e.g. Clark 1991; Sander, Schreiber and Doherty 2000; Bruch and Mare 2006).

Although using a simulation to generate patterns that one would expect to find (if the model is correct) and then comparing these with targeted observations of the social world is considerably easier than trying to obtain detailed data about social processes directly, there are two complications which must be considered. The first is that most models and the theories on which they are based are stochastic. That is, they are based in part on random chance. For example, in a segregation model, the simulation will normally initialise the landscape by distributing agents randomly. As the simulation runs, the agents migrate according to their preferences to locations where they feel more comfortable with their neighbours. The spot where they finish will depend in a complicated way on where they and all the other agents started. The precise pattern of clusters will depend on the chance arrangement of agents at initialisation; re-running the simulation with a new random starting configuration will yield a different pattern of clusters. The important point about this and similar models

is not that they generate a particular pattern of clusters, but that in every case, for a specific set of parameters, some clustering always emerges. The characteristics of these clusters can be assessed using measures such as the mean cluster density (averaged over many runs of the model, each with a different starting configuration), and the variance of the cluster size. It is these 'statistical signatures' that need to be compared with the observed residential segregation, which itself can be considered to be one possible outcome of a stochastic process. Unfortunately, the statistical methods needed to make sound comparisons, given that the distributions of the metrics are unknown but often far from 'normal', are not yet well developed, at least in the social sciences.

The second caveat is that many different models may yield the same emergent patterns. Hence, a correspondence between what one sees emerging from the model and what ones sees in the social world is only a necessary, but not a sufficient condition for concluding that the model is correct. There are many different kinds of processes which can yield clustering and so the fact that households are often ethnically segregated and the Schelling model generates clusters does not prove that the Schelling process is in fact the process followed by households in making migration decisions (Gilbert 2002). All one can do is to gradually increase one's confidence in a model by testing it against observation in more and more ways. In this respect, the methodology of simulation is no different from other approaches in social science.

Examples of agent-based models

Many hundreds of multi-agent social simulation models have now been designed and built, to examine a very wide range of social phenomena. It is not practicable to review all of these, and even describing a representative sample would be a difficult exercise. However, there are dimensions along which models can be arranged (see e.g. Hare and Deadman 2004; Berger and Manson, 2001; David et al 2004). In this section, these will be illustrated by reference to some typical social simulations.

Abstract versus Descriptive. Models can vary in the degree to which they attempt to incorporate the detail of particular targets. An example of a model which aims to be a detailed representation of a specific location and the developments there is the work by Dean et al. (1999) on the Long House Valley, in northern Arizona near Monument Valley. The model covers a time from about A.D. 400 to 1400 and consists of agent households that inhabit a digitized version of the Long House Valley landscape. Agents have rules for determining their agricultural practices and residential locations, as well as for reproduction and mortality. Each run of the model generates a unique history of population, agricultural output, and settlement patterns which can be compared with archaeological evidence from the Valley.

In contrast, a series of papers (Conte and Castelfranchi 1995, Castelfranchi, Conte and Paolucci 1998, Saam and Harrer 1999, Staller and Petta 2001, Flenthe, Polani and Uthman 2001, Hales 2002, Younger 2004) has explored the relationship between norms and social inequality using a very simple 'game' in which agents controlled by 'norms' (i.e. behavioural rules) search a regular grid for 'food', which they consume to maintain their energy level or 'strength'. These authors carried out experiments to study the mean and variance of the distribution of strength under various normative arrangements. Two examples are 'blind aggression', when agents attack other agents to grab their food, regardless of whether the attacked agent is stronger than they are, and 'finders-keeper', when agents respect 'property rights' and do not attack other

agents for their food. The experiments have shown that under some conditions (e.g. when the agents start with more or less equal levels of strength), the finders-keeper norm reduces inequality, but if the agents start with an unequal distribution, holding to the same norm can increase the degree of inequality. These findings are not directly descriptive of or applicable to any real human society or group, although they do raise some interesting questions for the conceptualisation of power and for understanding the origins of social inequality.

Artificial versus Realistic. Although the work on norms mentioned above is highly abstract, it is intended to aid in the understanding of actual human societies. However, some agent-based models are not intended as simulations of human societies at all. A good example is the research of Doran (1997), who investigated the implications if agents were able to see what will befall them in the future (that is, have perfect foresight). Other work on artificial societies has been driven by a desire to engineer groups of cooperative agents to achieve results that single agents could not do on their own. While some of this engineering-oriented modelling takes its inspiration from human societies, much of it assumes a command and control regime that is not a plausible description of real world societies (for examples of such an engineering approach to social simulation, see, for example, Wooldridge 2002).

In contrast, some models are firmly focussed on modelling real social problems. An excellent example of this is Eidelson and Lustick's (2004) research on the effectiveness of alternative defensive strategies against a possible smallpox attack or other major epidemic. Obviously there is neither much experience nor the possibility of experimentation to compare options, such as inoculating a whole population as a precaution versus vaccinating cases after the infection has begun to spread. Their model allows a number of possibilities to be investigated and the most important parameters for confining the epidemic to be identified.

Positive versus normative. Models with clear application to policy domains may tend towards being normative, that is, designed to make recommendations about what policies should be pursued. For example, Moss (1998) developed a model to represent the decision-making of middle managers in crises and was then able to make some tentative recommendations about the appropriate organisational structures to deal with critical incidents. This article is also interesting for the methodology of model building that it recommends. The majority of social agent-based simulations, however, are intended to be positive, that is descriptive and analytical about the social phenomena studied, aiding understanding rather than providing advice.

Spatial versus Network. The agents in some models operate in a spatial environment, often a two dimensional grid of rectangular cells, but sometimes a map of some specific landscape, over which the agents are able to move. In the latter case, the map is often provided by a geographical information system (GIS) (Dibble and Feldman 2004). An example is the model of the recreational use of Broken Arrow Canyon in Arizona (Gimblett, Itami and Richards 2002), which was developed to study policies for protecting the environment and providing a good recreational experience for visitors. Options include building new trails, limiting the number of visitors, or relocating existing trails. The model includes a detailed representation of the environment, including the physical topography of the canyon.

For other models, the physical geography is irrelevant. What is important are the relationships between agents, often represented as a network of links between nodes. For example, Gilbert, Pyka and Ahrweiler (2001) describe a model of an 'innovation network' in which the nodes are high tech firms that each have a knowledge base

which they use to develop artefacts to launch on a simulated market. Some artefacts are successful and the firms thrive; others fail. The firms are able to improve their innovations through research or by exchanging knowledge with other firms. The form of the emergent network and its dynamics observed from the simulation are compared with data from the biotechnology and mobile personal communication sectors and shown to be qualitatively similar.

Complex versus Simple agents. The simplest agents are ones that use a production system architecture (Gilbert and Troitzsch 2005), meaning that the agent has a set of condition-action rules. An example of such a rule could be 'IF the energy level is low, THEN move one step towards the nearest food source'. The agent matches the condition part of the rule against its present situation and carries out the corresponding action. These rules might be explicitly coded as declarative statements, as in this example, or they may be implicit in a procedural algorithm. However, it is difficult to model cognitively realistic agents using such a simple mechanism and so model-builders have sometimes adopted highly sophisticated cognitive model systems to drive their agents. The best known of these are SOAR (Laird, Newell and Rosenbloom 1987) and ACT-R (Anderson and Lebiere 1998). Carley, Prietula and Lin (1998) describe a number of experiments comparing models of organisation in which the agents have cognitive architectures of increasing complexity, from a basic production system to the use of a version of SOAR. They conclude that simpler models of agents are all that is needed if the objective is to predict the behaviour of the organisation as a whole, but more cognitively accurate models are needed to generate the same predictive accuracy at the individual or small group level.

Developing multi-agent models

Developing good multi-agent models is still something of an art, rather than a science. However, there is now some understanding of the steps that usually need to be carried out (Gilbert and Terna 2000; Gilbert 2007). The first is to be sure about the objective of the work. The research question and the model that is to be designed are sometimes clear from the start. More often, one has an idea of the topic, but not anything more precise. It is essential that a general interest in a topic is refined down to a specific question before the design begins. If this is not done, either the design task can seem impossibly difficult or the model can become too encompassing to be helpful.

It is useful to think about narrowing down a research question in terms of moving through a set of layers (see Punch 2000 for a helpful treatment). An area of research contains many topics. More specific is a general research question, usually phrased in terms of theoretical concepts and the relationship between these. The general research question will generate a small number of specific research questions. The specific research questions should be at a level of detail such that their concepts can be used as the main elements of the model.

The social world is very complicated, a fact that modellers are well aware of, especially when they begin to define the scope of a model. The art of modelling is to simplify as much as possible, but not to oversimplify to the point where the interesting characteristics of the phenomenon are lost. Often, an effective strategy is to start from a very simple model, which is easy to specify and implement. When one understands this simple model and its dynamics, it can be extended to encompass more features and more complexity.

The baseline model can be designed to be the equivalent of a null hypothesis in statistical analysis: a model that is not expected to show the phenomenon in question. Then, if an addition to the baseline model is made and the model behaves differently, one can be sure that it is the addition that has the effect. This strategy also has the advantage that it helps to focus attention on the research question or questions that are to be answered. A modeller should always have at the forefront of their attention why they are building the model and what they are seeking to obtain from it.

If the baseline model is simple enough, the first prototype implementation can sometimes be a 'pencil and paper' model, in which the designer (or the designer and a few colleagues) plays out the simulation 'by hand' through a few rounds. This simulation of a simulation can quickly reveal gaps and ambiguities in the design, without the need to do any coding.

Designing a model is easier if there is already a body of theory to draw on. At an early stage, therefore, one should look around for existing theory, in just the same way as with more traditional social science methodologies. Theories that are about processes of change and that consider the dynamics of social phenomena are of course likely to be more helpful than theories about equilibria or static relationships, but any theory is better than none. What the theory provides is an entry to the existing research literature, hints about what factors are likely to be important in the model, and some indications about comparable phenomena. Another function of theory can be to identify the assumptions on which the model is built. These assumptions need to be as clearly articulated as possible if the model is to be capable of generating useful information.

Once the research questions, the theoretical approach and the assumptions have been specified, it is time to begin to design the simulation. A set of issues needs to be considered for almost all simulations, and it is helpful to deal with these systematically and in order. Nevertheless, there is no 'right' or 'wrong' design so long as the model is useful in addressing the research question.

The first step is the definition of the types of objects to be included in the simulation. Most of these objects will be agents, representing individuals or organizations, but there may also be objects representing inanimate features that the agents use, such as food or obstacles. The various types of object should be arranged in a class hierarchy, with a generic object at the top, then agents and other objects as subsidiary classes, and if necessary, the agent class divided into further sub-classes. Each actual object in the simulation will be an example of one of these types (an 'instance' of the class). All instances of a class are identical in terms of the code that creates and runs them, but each instance can be in a different state, or have different attributes.

Once the objects have been decided, one can consider the attributes of each object. An attribute is a characteristic or feature of the object, and is either something that helps to distinguish the object from others in the model, or is something that varies during the execution of the simulation. Attributes function like variables in a mathematical model. Consider each object in turn, and what features it has that differ from other objects. Properties such as size, colour or speed might be relevant attributes in some models. State variables such as wealth, energy and number of friends might also be attributes.

An attribute might consist of a single one of a set of values (for example the colour attribute might be one of red, green, blue or white); a number, such as the energy level of the agent; or a list of values, such as the list of the names of all the other agents that an agent has previously encountered. Sub-classes inherit the attributes of their

superior class, so that, for instance, if all objects have a location, so do all its sub-classes. When the attributes for each class of object have been decided, they can be shown on a class diagram. This way of representing classes and attributes is taken from a design language called the Unified Modelling Language (UML) (Booch et al. 2000) and is commonly used in object-oriented software design.

The next stage is to specify the environment in which the objects are located. If the environment is a spatial one, each object has a location within it (in that case, the objects need to have attributes that indicate where they are at the current time). But there are other possibilities, such as having the agents in a network linked by relations of friendship or trade with other agents. Sometimes it may be convenient to represent the environment as another object, albeit a special one, and specify its attributes. One of the attributes will be the current simulated time. Another may be a message buffer that temporarily holds messages sent by agents to other agents via the environment before they are delivered. Defining the classes, attributes and environment is an iterative process, involving refining the model until the whole set seems consistent.

When this is done, at least to a first approximation, one has a static design for the model. The next step is to add some dynamics, that is, to work out what happens when the model is executed. It is usually easiest to start by considering the interactions of each class of agent with the environment. An agent acts on the environment in one or more ways and the environment acts on the agent. Once lists of the actions of the agents and the environment have been created, one can consider when the actions happen. Against the list of agent actions on the environment, indicate the conditions under which these actions should occur. This table of conditions and actions will lead naturally to defining a set of condition-action rules. Each rule should be associated with a unique state of the agent (a unique set of attribute values and inputs from the environment). After the interactions with the environment have been decided, the same job can be done for interactions between agents.

It is likely that, in working through these lists, it will be realized that additional attributes are needed for the agents or the environment or both, so the design process will need to return to the initial stages, perhaps several times. When a consistent set of classes, attributes and rules has been created, it can be helpful to summarize the dynamics in a sequence diagram, another type of UML diagram. A sequence diagram has a vertical line for each type or class of agent, and horizontal arrows representing messages or actions that go from the sender object to the receiver object. The sequence of messages is shown by the vertical order of the arrows, with the top arrow representing the first message and later messages shown below.

It can also be useful to employ state chart and activity diagrams to summarize the behaviour of agents (Fowler and Scott 1999). A state chart diagram shows each distinct state of an agent and what is involved in moving from one state to another. An activity diagram shows how decisions are made by an agent.

At this stage in the design process, most of the internal aspects of the model will have been defined, although normally there will still be a great deal of refinement needed. The final step is to design the user interface. The components of this interface will be graphical representations of sliders, switches, buttons and dials for the input of parameters, and various graphs and displays for the output, to show the progress of the simulation. Initially, for simplicity it is best to use a minimum of input controls. As understanding of the model improves, and additional control parameters are identified, further controls can be added. Similarly, with the output displays, it is best

to start simple and gradually add more as the need for them becomes evident. Of course, every model needs a control to start it, and a display to show that the simulation is proceeding as expected (for example, a counter to show the number of steps completed). At the early stages, there may also be a need for output displays that are primarily there for debugging and for building confidence that the model is executing as expected. Later, if these displays are not required to answer the research question, they can be removed again.

Even before the coding of a model is started, it is worth considering how the simulation will be tested. A technique that is gaining in popularity is 'unit testing'. The idea is that small pieces of code that exercise the program are written in parallel with the implementation of the model. Every time the program is modified, all the unit tests are re-run to show that the change has not introduced bugs into existing code. As the model is extended, more unit tests are written, the aim being to have a test of everything. The idea of unit tests comes from an approach to programming called XP (for eXtreme programming, Beck 1999), a software engineering methodology that is particularly effective for the kind of iterative, developmental prototyping approach that is common in most simulation research.

When there are many unit tests to carry out, it becomes tedious to start them all individually and a test harness that will automate the process is needed. This will also have to be designed, possibly as part of the design of the model itself, although there are also software packages that make the job easier (see, for example, the open source Eclipse toolset, http://www.eclipse.org/). When the model is working as expected, it will probably be necessary to carry out sensitivity analyses involving multiple runs of the simulation while varying the input parameters and recording the outputs. Doing such runs manually is also tedious and prone to error, so a second reason for having a test harness is to automate analysis. The starting and ending points of an input range can be set and then one can automatically sweep through the interval, re-running the model and recording the results for each different value. To enable this to be done, the model may need to have two interfaces: a graphical one so that the researcher can see what is happening and an alternative test- or file-based interface that interacts with the testing framework.

It is likely that all the output from the first run of a model will be due, not to the intended behaviour of the agents, but to the effect of bugs in the code. Experience shows that it is almost impossible to create simulations that are initially free of bugs and, while there are ways of reducing bugs (for example, the unit test approach mentioned above), one should allow at least as much time for chasing bugs as for building the model. The most important strategy for finding bugs is to create test cases for which the output is known or predictable, and to run these after every change until all the test cases yield the expected results. Even this will not necessarily remove all bugs and modellers should always be aware of the possibility that their results are merely artefacts generated by their programs.

Conclusion

Agent-based models can be of great value to the social sciences and their potential is beginning to be realised. In this chapter, I have shown that such models are especially relevant to simulating social phenomena that are inherently complex and dynamic. These models are also effective at demonstrating the emergence of social institutions from the actions of individual agents, an area where previous methods of analysis can be very weak. A number of examples of work in which multi-agent models have

been used were reviewed to show the range of possibilities now available to the researcher: models may be abstract or descriptive, positive or normative; based on a geographical landscape or represent a social network; while the agents themselves may be simple or very complex. Now that agent-based simulation is a thriving area of research, there is a growing body of experience on how to build models and in the latter part of this chapter, I outlined a typical process of model development. With this as guidance (see Gilbert and Troitzsch 2005 for a more detailed discussion of methods) and an idea of the opportunities for social simulation, I hope that many readers might be inspired to try this approach to social science for themselves on their own research topics.

References

Anderson, J. R. and Lebiere, C. (1998) *The Atomic Components of Thought*. Erlbaum, Mahwah, NJ.

Arthur, W.B. (1989) 'Competing Technologies, Increasing Returns, and Lock-In by Historical Events', *Economics J.* vol.116, p.99

Axelrod, R. (1995) 'A model of the emergence of new political actors'. In N. Gilbert & R. Conte (Eds.), *Artificial Societies*. London: UCL Press.

Beck, K. (1999) *Extreme Programming Explained*. Addison-Wesley, Boston, MA.

Booch, G., Rumbaugh, J. and Jacobson, I. (2000) *The Unified Modeling Language User Guide*. 6th print edn. Addison-Wesley, Reading, MA.

Breiger, R., Kathleen Carley and Philippa Pattison (eds.) (2003) *Dynamic Social Network Modelling and Analysis: Workshop Summary and Papers* Washington: The National Academies Press.

Bruch, Elizabeth E., & Mare, Robert D. (2006). Neighborhood Choice and Neighborhood Change. *American Journal of Sociology, 112*(4), 667-709..

Carley, K. M., Michael J. Prietula and Zhiang Lin (1998) 'Design Versus Cognition: The interaction of agent cognition and organizational design on organizational performance'. *Journal of Artificial Societies and Social Simulation* vol. 1, no. 3, http://www.soc.surrey.ac.uk/JASSS/1/3/4.html

Castelfranchi, C., Conte R. and Paolucci M. (1998) 'Normative reputation and the costs of compliance'. *Journal of Artificial Societies and Social Simulation* 1(3) <http://www.soc.surrey.ac.uk/JASSS/1/3/3.html>

Clark, W.A.V., 1991, 'Residential Preferences and Neighborhood Racial Segregation: A Test of the Schelling Segregation Model'. *Demography* 28:1-19.

Conte R and Castelfranchi C (1995) 'Understanding the functions of norms in social groups through simulation'. In Gilbert N and Conte R (Eds.) *Artificial Societies*. London: UCL Press. pp. 252-267.

David, N. Maria Bruno Marietto, Jaime Simão Sichman and Helder Coelho (2004) 'The Structure and Logic of Interdisciplinary Research in Agent-Based Social Simulation'. *Journal of Artificial Societies and Social Simulation* vol. 7, no. 3 http://jasss.soc.surrey.ac.uk/7/3/4.html

Dean, J.S., Gumerman, G.J., Epstein Joshua, M., Axtell, R.L., Swedland, A.C., Parker, M.T. and McCarrol, S. (1999) 'Understanding anasazi culture change through agent based modeling' in Kohler, T.A. and Gumerman, G.J. eds. *Dynamics in human and primate societies: Agent based modeling of social and spatial processes*, Oxford University Press.

Dibble, Catherine and Philip G. Feldman (2004) 'The GeoGraph 3D Computational Laboratory: Network and Terrain Landscapes for RePast'. *Journal of Artificial Societies and Social Simulation* vol. 7, no. 1 http://jasss.soc.surrey.ac.uk/7/1/7.html

Doran J. E. (1977) 'Foreknowledge in Artificial Societies' In *Simulating Social Phenomena* (eds. R Conte, R Hegselmann, and P Tierna), Lecture Notes in Economics and Mathematical Systems 456, Springer: Berlin. pages 457-469.

Doran J. E. (1998) 'Simulating Collective Misbelief' *Journal of Artificial Societies and Social Simulation* vol. 1, no. 1, http://www.soc.surrey.ac.uk/JASSS/1/1/3.html

Eidelson, Benjamin M. and Ian Lustick (2004) 'VIR-POX: An Agent-Based Analysis of Smallpox Preparedness and Response Policy' *Journal of Artificial Societies and Social Simulation* vol. 7, no. <http://jasss.soc.surrey.ac.uk/7/3/6.html>

Flentge, Felix, Daniel Polani and Thomas Uthmann (2001) 'Modelling the Emergence of Possession Norms using Memes' *Journal of Artificial Societies and Social Simulation* vol. 4, no. 4 <http://www.soc.surrey.ac.uk/JASSS/4/4/3.html>

Fowler, M. and Scott, K. (1999) *UML Distilled*. 2nd edn. Addison Wesley, Reading, MA.

Gilbert, N. (1995). 'Emergence in social simulation' Pp. 144-156 in *Artificial Societies: the computer simulation of social life*, edited by N. Gilbert and R. Conte. London: UCL Press.

Gilbert, N. (2002). 'Varieties of emergence'. Paper presented at the Agent 2002 Conference: *Social agents: ecology, exchange, and evolution*, Chicago.

Gilbert, N. (2007). *Agent-based models*: Sage Publications Inc.

Gilbert, N., and Troitzsch, K. G. (2005). *Simulation for the social scientist*. Second edition. Milton Keynes: Open University Press.

Gilbert, Nigel, Andreas Pyka and Petra Ahrweiler (2001) 'Innovation Networks - A Simulation Approach' *Journal of Artificial Societies and Social Simulation* vol. 4, no. 3, http://www.soc.surrey.ac.uk/JASSS/4/3/8.html

Gilbert, N., and Terna, P. (2000). 'How to build and use agent-based models in social science'. *Mind and Society*, 1(1), 57 - 72.

Gimblett, H.R., R.M. Itami & M. Richards (2002) 'Simulating Wildland Recreation Use and Conflicting Spatial Interactions using Rule-Driven Intelligent Agents'. In H. R Gimblett, editor. *Integrating GIS and Agent based modeling techniques for Understanding Social and Ecological Processes*. Oxford University Press.

Hales, David (2002) 'Group Reputation Supports Beneficent Norms' *Journal of Artificial Societies and Social Simulation* vol. 5, no. 4 http://jasss.soc.surrey.ac.uk/5/4/4.html

Hare, M. and Peter Deadman (2004) 'Further towards a taxonomy of agent-based simulation models in environmental management'. *Math. Comput. Simul*. Vol. 64(1), pp. 25-40.

Laird, J. E., Newell, A. and Rosenbloom, P. S. (1987) 'Soar: An architecture for general intelligence'. *Artificial Intelligence*, 33: 1–64.

Moss, Scott (1998) 'Critical Incident Management: An Empirically Derived Computational Model'. *Journal of Artificial Societies and Social Simulation* vol. 1, no. 4, http://www.soc.surrey.ac.uk/JASSS/1/4/1.html

Nowak, A., and Latané, B. (1994). 'Simulating the emergence of social order from individual behaviour'. In N. Gilbert & J. Doran (Eds.), *Simulating Societies: the computer simulation of social phenomena* (pp. 63-84). London: UCL Press.

Parker, D, Berger, T and Manson, S (eds) (2001) *Agent Based Models of Land-Use and Land-Cover Change*: Report and Review of an international Workshop, Indiana, Indiana University.

Punch, K. F. (2000) *Developing Effective Research Proposals*. Sage, London.

Saam, Nicole J. and Andreas Harrer (1999) 'Simulating Norms, Social Inequality, and Functional Change in Artificial Societies' *Journal of Artificial Societies and Social Simulation* vol. 2, no. 1, http://www.soc.surrey.ac.uk/JASSS/2/1/2.html

Sander, R., D. Schreiber, and J. Doherty (2000) 'Empirically Testing a Computational Model: The Example of Housing Segregation'. pp. 108-115 in Proceedings of the Workshop on Simulation of Social Agents: Architectures and Institutions, edited by D. Sallach and T. Wolsko, Chicago, IL: University of Chicago; ANL/DIS/TM-60, Argonne National Laboratory, Argonne, IL.

Schelling, T. C. (1971). 'Dynamic models of segregation'. *Journal of Mathematical Sociology*, 1, 143-186.

Staller, Alexander and Paolo Petta (2001) 'Introducing Emotions into the Computational Study of Social Norms: A First Evaluation'. *Journal of Artificial Societies and Social Simulation* vol. 4, no. 1,<http://www.soc.surrey.ac.uk/JASSS/4/1/2.html>

Wooldridge, Michael (2002) *An Introduction to Multi-Agent Systems*. John Wiley and Sons Limited: Chichester

Younger, Stephen (2004) 'Reciprocity, Normative Reputation, and the Development of Mutual Obligation in Gift-Giving Societies'. *Journal of Artificial Societies and Social Simulation* vol. 7, no. 1 http://jasss.soc.surrey.ac.uk/7/1/5.html