

Environment design for emerging artificial societies

Nigel Gilbert Stephan Schuster
Matthijs den Besten Lu Yang

Centre for Research on Social Simulation
University of Surrey, Guildford GU2 7XH, UK
n.gilbert@soc.surrey.ac.uk

Abstract

The NewTies project is developing a system in which societies of agents are expected to develop autonomously as a result of individual, population and social learning. These societies are expected to be able to solve the environmental challenges that they are set by acting collectively. The challenges are intended to be analogous to those faced by early, simple, small-scale human societies. Some issues in the construction of a virtual environment for the system are described and it is argued that multi-agent social simulation has so far tended to neglect the importance of environment design.

1 Introduction

The goal of social simulation is to develop models that shed some light on the functioning of human societies. The advantages of a simulation approach to understanding human societies include the requirement to express theories in complete and unambiguous terms; the opportunity to derive the implications of proposed social mechanisms; and the possibility of performing experiments on the simulated society (Gilbert, 2005). As a result of these advantages, there has been a rapid growth in the popularity of social simulation over the last decade (Gilbert & Troitzsch, 2005).

There are two main current approaches to the construction of simulation models of society. One approach starts with data observed or collected from a human society and tries to find a model that reproduces the observations. This approach, which generally yields results that are complex but can be compared directly with the observed data, has been labelled KIDS (Keep It Descriptive) (Edmonds & Moss, 2004). The other approach, named KISS (Keep It Simple) (Axelrod, 1997), begins by attempting to simplify the putative social phenomena to its essence and models only an abstract version of the society. The model tends to be easier to explore and understand, but validation against human societies is much harder.

This paper, like the NewTies project of which it is part¹, takes a third approach. We aim to see whether an artificial society can ‘construct itself’ with only the bare minimum of experimenter provided rules or theory. We take our inspiration partly from work on the evolution of language, which has shown that, given a capacity to learn, artificial agents are capable of developing a simple ‘language’ with which to communicate (see Cangelosi & Parisi (2002) for an overview). Initially agents utter only random noise with no information content, but through repeated interactions, some of which are rewarded, the agents gradually develop a shared lexicon (‘a consensus on a set of distinctions’ Hutchins & Hazlehurst, 1995:161).

If agents can develop a lexicon from ‘nothing’, could they also develop a shared culture? This is the hypothesis underlying the NewTies project. Taken strictly, the answer must be ‘yes’, since language and thus a lexicon *is* an important part of human culture. But we wish to see whether agents can develop culture in a wider sense, as a set of shared behaviours and understandings of the society and environment in which they live. This culture, like the shared lexicon, must be developed collaboratively by the agents from ‘nothing’. This means that we give the agents the ability to learn, but do not

¹ New and Emergent World models Through Individual, Evolutionary, and Social learning (NEW TIES), <http://www.newties.org>

direct them about what to learn, and initialise them with a bare minimum of knowledge about their worlds. However, the agents are given a rich and extensive simulated environment in which they have to learn how to survive.

The next section of the paper reviews the types of learning available to agents. The following two sections introduce the environment and the challenges that the agents face. The fifth section describes the proposed interface between the environment and an agent. Agents perceive their surroundings and act in their world through this interface. The paper concludes by emphasising the importance of the design of the environment for social simulation and suggesting that this aspect has too often been neglected.

2 Learning

The agents are constructed to be able to learn in three ways:

a. Individual learning through trial and error.

Agents act according to their genetic predispositions, overlaid with random variations. Some actions are more effective than others. Those actions that succeeded in the past are remembered and the agent is then more likely to repeat those actions than the others.

b. Population learning through reproduction and selection

Agents with predispositions to carry out effective actions more frequently are more capable and are therefore more likely to reproduce, transferring a version of their genetic material to their offspring. Thus the population of agents as a whole will tend to become more successful over the course of many generations.

c. Social learning

Neither individual nor population learning require any communication between agents. However, these types of learning could be the means by which the agents begin to develop a language for communication. If they do so, they can start to use a more direct and effective mode of learning: that of one agent teaching another.

3 Environmental challenges

If agents are to learn, they must have some motivation to do so. In the NewTies project, that motivation is ultimately that of their survival. Agents are placed in an environment which they find individually and collectively challenging. Unless they master survival in this environment they will 'die'. This notion is operationalised by constructing environments in which there is a limited amount of 'food' to provide agents with energy and requiring the

agents to maintain at least a minimum energy level. At first, agents have to act on their own, since they have not yet learnt to act collectively. Those that manage to collect sufficient food from the environment may survive long enough to breed, while those that are less successful are more likely to 'starve'. The environment thus imposes a strong selection pressure on the agents. Eventually, the agents may discover how to communicate and then be able to engage in collective action. This is likely to be more effective than individual acts in obtaining food from the environment.

The fine detail of the environmental challenge is an extremely important factor in the agents' development. If obtaining food is too easy, the agents will not need to learn much, and will probably not do so. If the environment is too unfriendly, all the agents will die of starvation before they have had a chance to learn anything. Secondly, if the environment requires agents to engage in activities which they are not able to carry out, the agents will surely fail, since they are only able to increase their knowledge through learning, but not their repertoire of basic actions. For example, humans have learned to fly, not by growing wings, but by learning how to build aircraft. Thirdly, the long-term objective of the research is to understand human societies better. The environment must set challenges that are analogous to those faced by humans if there is to be even the possibility of reading across from the simulation to human development.

We have designed four environmental challenges, each based on a well studied aspect of human society. In the descriptions below, the human system is first summarized, the challenge stated in terms of the simulated environment, and the observable outcome that might be expected is specified.

3.1 The Kula Ring

A complex system of visits and exchanges among the Trobriand Islanders of the western Pacific was first described by Bronislaw Malinowski (1922 [1978]). Necklaces were exchanged in one direction among the residents of a chain of islands and arm-bands exchanged in the opposite direction (hence the notion of a ring). These exchanges did not primarily serve an economic function but created a network of social obligations among peoples which could be depended upon at various times in an individual's life. In particular, the social network seems to have been the basis for economic relationships such as trading food for pottery.

The challenge parameters:

Food is distributed in spatial patches and the amount of food in a patch varies over time. The overall quantity is more than enough to feed the population, but there may be short-term local shortages. These can be alleviated by trading or by theft. Trade is

less costly in energy, but requires the prior development of mutual trust by the traders.

Expected outcome:

The establishment of a 'gift-exchange' system in which not only food but also tokens are exchanged.

3.2 Herders in a semi-arid area

Nomadic herding is another human solution for dealing with variable and uncertain shortages. Herders and their cattle move to where food is available, leaving exhausted areas until the grass has re-grown. This requires herders to find ways of managing common pool resources (the grass) so that no individual herder overgrazes the grass. The human solution involves well developed status hierarchies and no private property.

The challenge parameters:

Food is randomly distributed with the mean level of food just sufficient to support the population. The rate of food growth varies randomly over time. Food is perishable. Some food must be left uneaten on each patch since subsequent growth is proportional to amount of food left uneaten.

Expected outcome:

Agents leave uneaten food when they move away, even if they leave hungry.

3.3 Central place theory

Walter Christaller developed Central Place theory in 1933 (King, 1985) to explain the size and spacing of cities that specialize in selling goods and services.

The theory consists of two basic concepts:

- threshold -- the minimum market needed to bring a firm or city selling goods and services into existence and to keep it in business
- range -- the average maximum distance people will travel to purchase goods and services

The theory predicts that settlement size will follow the rank size rule. It works well for human settlements.

The challenge parameters:

The distribution of types of food is such that agents need to trade food with other agents. The food types vary in their transportability. Agents can move to find the best location to maximise their income from trade.

Expected outcome:

Agents settle into spatial clusters separated by relatively empty areas. The size of the clusters is power law distributed.

3.4 Branding

When producers produce and consumers consume complex goods (i.e. ones with a large number of distinct attributes), and there are a large number of

producers and consumers, search problems occur. Producers find it hard to locate consumers that desire goods having the precise set of attributes that a producer is selling, and consumers find it hard to identify producers with the desired goods. One 'solution' to the problem each side faces is for producers to brand their range of goods (targeting them at a subset of consumers) and for consumers to use the brand as the major preference criterion. Similar processes may help to account for prejudice and discrimination among human populations.

The challenge parameters:

Agents have characteristic sensible attributes ('tags'). Agents seek to locate other agents with a similar or identical set of tags (through movement and communication), but this search is expensive. Agents are able to create additional tags (the brand) by collecting tokens and carrying them around.

Expected outcome:

Agents either generate one additional tag or specially distinguish an existing tag and this becomes a linguistic category that labels agents and leads to differences in behaviour towards those agents that are labelled and those that are not.

4 The virtual environment

An environment that offers these challenges to agents must be sufficiently rich in features to allow each challenge to be constructed, but also no more complicated than necessary. Any features beyond the minimum required would slow down the simulation and, crucially, make the agents' task of learning how to manage in the environment more difficult, because they would need to learn to disregard irrelevant features.

The environment we have designed consists of a very large simulated flat surface over which the agents are able to move. The surface is divided into small patches or 'locations'; an agent or other object is of a size that it occupies exactly one location. A virtual clock counts 'time steps', used primarily to synchronise the agents' actions. To remain in accord with the real world, agents do not have direct access to their location on the surface, nor to the time. They are, however, able to detect geographical features ('places') and the relative position of the 'sun', an object which slowly traverses the surface, crossing it once per simulated day (there is no night – the sun is always visible). Places are bounded areas of the landscape which differ from the rest of the surface in having a varied, but lesser degree of roughness, making it easier for agents to move within places than in the wilderness outside places.

On the landscape are a number of objects as well as the agents: tokens, plants, and paving stones. Tokens are distinguishable, moveable objects, some

of which can be used as tools to speed up the production of food, but most of which have no intrinsic function, but can be employed by agents as location markers, symbols of value ('money'), or for ritual purposes.

Plants are the source of food. They are annuals, living for one year. At the beginning of the year, eating them gives agents little energy, but as the year progresses, they ripen and become better food. In the 'autumn', their energy value decreases again, and is entirely lost at the end of the year when they die. However, before they die, they produce two seeds, one at the parent plant's location and one in an adjacent location. If a seed is the only one in the location, it grows, but if there are more than one, only one will survive. If a plant is picked by an agent, it starts decomposing and will lose all its goodness if not consumed or replanted within a few days.

Agents lose energy (the rate depending on the roughness of the location) when they move over the landscape. The effort required to move can be reduced by building roads. Roads are constructed from paving stones laid end to end.

With these simple ingredients, we can construct scenarios corresponding to each of the challenges. For example, the Trobriand Islands can be represented as places, with the rest of the surface (having a very high value of roughness) representing the sea. The varied availability of food among the Islands (and the seasonal availability of crops) can be represented by arranging the plants in the places. The agents can learn to use tokens as symbolic gifts. Economic trading between islands could involve exchanges of food and of token tools. The other challenges could be modelled by constructing 'scenarios' in similar ways. For example, the 'branding' challenge would involve agents trading many similar but not identical tokens between themselves, with search being costly (i.e. the roads are rough).

5 Agent interface

To survive in this environment, agents need to be able to perceive the landscape and the objects in it, and also need to be able to act on objects and other agents. Moreover, it is expected that experiments will be carried out using a variety of agent designs, possibly including agents constructed outside the NewTies project, and so a simple and precisely specified interface between the agents and the environment is desirable.

At each time step, every agent is given a slice of computational resource. During this step, it must complete two phases in sequence: a *perceive* phase and an *act* phase. During the perceive phase, an

agent is given the following information about the environment:

- a. a list of the attributes (type, characteristics, colour, heading, and weight) of each object located within a segment defined by the direction in which the agent is facing, plus or minus 45°. The information returned about each object also includes its distance and direction from the agent and, if the object is an agent, its age and sex. These data do *not* include any direct indicator of the objects' identities; the agents have to infer these from the objects' attributes..
- b. A list of the places in which the agent is located (places can overlap, so there may be more than one).
- c. The agent's current energy level.
- d. A list of the attributes of all the objects that the agent is currently carrying.
- e. The roughness at the current location.
- f. The result of the action performed in the Act phase of the previous time step, if any.
- g. A list of messages that other agents have sent during the preceding Act phase.

The agent is able to process this information as it wishes, and can then carry out one action, chosen from the following:

- **Move:** The agent moves from its present location to an adjacent location in its forward direction.
- **Turn left / turn right:** the agent rotates in the indicated direction by 45 degrees.
- **Pick up object:** The agent acquires the object. The object remains with the agent until the agent puts it down or eats it (if the object is food).
- **Put down object:** The agent puts the object down at the current location.
- **Give object:** The agent transfers an object in its possession to another agent. The receiving agent must be in an adjacent location.
- **Take object:** The agent takes an object from another agent. The donating Agent must be in an adjacent location.
- **Build/improve road:** The agent builds (if there is no road already) or improves (i.e. reduces the roughness of) the road at the current location.
- **Talk to agent:** The recipient agent must be 'visible' to the speaker (An agent cannot talk to another agent while facing away from that Agent, but the hearer does not have to be facing the speaker). A character string emitted by the speaker is conveyed to the listener. The effect is that both the

listener and the speaker are given the character string during the next Perceive phase.

- **Shout:** A character string emitted by the shouter is conveyed to all agents within a short distance (including the shouter itself) during the next Perceive phase.
- **Hit:** The agent chooses, first, the amount of energy to expend on the blow, which must be less than the current energy level of the Agent, and, second, which agent will be the victim (the victim must be in an adjacent location). Both the aggressor agent and the victim lose energy proportional to the ratio of the weights of the aggressor and the victim. If the victim's weight decreases to zero or less as a result of the violence, the victim dies.
- **Eat food:** The agent must already be carrying the food (see Pick up object). The energy of the food is added to the agent's energy and the food 'disappears'.

The information given to agents about their environment is intended to reflect the information which would be available to a human. Particular care is taken not to give agents information which would not be accessible to people. For example, the identity of other agents is not provided, only some descriptive characteristics through which agents may be recognised. However, there is no guarantee that all agents will necessarily have a unique set of these characteristics. Also, in a small group, only a subset of the characteristics may in fact be needed to distinguish agents. Utterances are labelled by the system, not with the identity of the speaker, but with its characteristics for the same reason. Speakers hear their own utterances reflected back to them, again because this is the experience of humans, who are able to monitor their own speech.

Initially, agents will have no common lexicon and therefore no understanding of what other agents say to them; we expect, in the light of studies on the evolution of language, that in time the agents will develop a shared vocabulary and ultimately a shared idea of grammar (see Vogt & Divina (2005) for details on language evolution in NewTies). However, because of the design of the agents and the environment, it is not necessary or even likely that this vocabulary will be entirely composed of utterances (i.e. 'words'). Because talking is just one of the actions available to agents, it would be expected that some actions other than talking will come to take on meaning for the agents – in the same way as human gestures, for example, can substitute for or even be preferred to speech for conveying some meanings. This is in contrast to current studies of the evolution of language, which have generally taken a more purely linguistic approach to interaction.

Although the list of possible actions may seem long, it is intended to be the minimum set that would enable the challenges to be met by the agents while yielding social behaviour comparable to that of human societies. For instance, the actions 'give object' and 'take object' are required in order to make trade a possibility. Without these actions, the only way to transfer an object from one agent to another would be for one agent to put the object down and another subsequently to pick it up. However, there would be no way for the first agent to guarantee that the second agent is the recipient, and thus directed personal transfers (required for trade) would be difficult or very risky. The justification for the 'hit' action (aside from the fact that violence is an endemic feature of human societies) is that without violence, private property cannot be preserved. An agent wanting an object in the possession of another could simply remove it and the owner would have no recourse if there were no possibility of violence. To match the human situation, an aggressor will only be effective if it is stronger (i.e. heavier) than the victim, so we can expect weak (light) individuals to be subject to theft which they cannot resist, at least until a protective social system evolves.

In this environment, agents have only one overriding 'motivation': to obtain sufficient food to survive². Human requirements are of course more complex, involving not just a reasonably balanced diet, but also warmth and water, but we are assuming that 'food' is an adequate abstraction for these more complex needs.

It is intrinsic to the implementation of population learning that agents are born, reproduce and so pass on their genotype, and die. New agents result from the coupling of a male and a female agent (hence agents need to have a gender) and are born in an adjacent location to their parents. Parents have no predisposition to attend to their offspring, but because they are nearby, are likely to interact with them more than with other agents. Parental care of offspring is likely to be selected for since neglected children will find survival even more difficult than their parents (since they have had no opportunity for individual learning). To enable adults to identify children, one of the characteristic features of agents, perceptible by other agents, is their age.

6 Conclusions

We have outlined a design for an environment which can be tuned in ways that are expected to promote the emergence of agent social behaviour to

² There is no need for the agents to have this motive 'hard-wired' by the experimenter; agents that are not so motivated, or that are motivated to gather food, but are not effective in doing so, simply die from starvation.

solve environmental challenges analogous to those that human societies have been able to overcome.

If such behaviour does arise, the simulation could serve as an invaluable test bed for examining a wide range of social theories. Its great advantage is that while one cannot experiment on human societies, one can on artificial societies. It will be possible, for example, to determine the conditions under which particular social phenomena emerge and survive in a way undreamt of by social theorists who can observe only a small number of human societies as cases on which to test their ideas. Even these few societies have been subject to an unknown amount of cross-fertilisation (for example, it is believed that the practice of agriculture was only discovered in two or three places in the world's history; all other agriculture was learned by copying these early innovations (Smith, 1995)).

Nevertheless, there must be some caveats about making too close a link between the simulation and human societies. On the one hand, the simulated agents are lacking many of the qualities of humans, and we do not know to what extent the differences between humans and the agents are important for the generation of analogous social phenomena (for example, we noted above that the simulation does not treat 'warmth' as a distinct need for the agents, although in cold climates it is for humans).

On the other hand, what we observe in human societies is one outcome from an unknown number of other possibilities. For example, it has been pointed out that, although most simple societies engage in some form of trade with other communities, the Kula Ring is unique. No other society has ever been discovered in which there is a two-way flow of symbolic goods. It follows that if the agent society does not generate an institution resembling the Kula Ring, this may simply be because an alternative institution has evolved, as it did in the great majority of human societies faced with similar challenges. This is of course a question that can be explored using the simulation: the experiment can be repeated many times to see whether a Kula phenomenon ever appears.

In contrast to most social simulation research, we have been almost exclusively concerned in this paper with the design of the environment; what in the environment is perceived by the agents; and the actions that the agents can take on the environment. The 'internal' design of the agents has been given little attention because it is entirely generic: agents are required to have:

- a means of generating actions as a function of their past history and current perceptions (but the form of this (phenotype) function is not of direct interest other than to the extent that it is affected by the agent's genotype),

- a genotype which, through some reproduction process, is able to generate copies with variation, and
- an algorithm for categorising objects and associating them with actions (including uttered 'words').

The details of *how* these internal processes work is little consequence for the simulations proposed here (which is not to say that these processes are trivial or easy to design). Their only important features is that they should be effective and efficient. Perhaps the fact that the agents can be black boxes, and yet the simulation can be interesting, should not be surprising, for this is the case with human societies also. We have only the flimsiest understanding of how humans 'work', yet both our social scientific and our everyday understanding of how societies work is increasingly sophisticated.

Acknowledgements

The NewTies project is partially supported by the European Commission Framework 6 Future and Emerging Technologies under contract 003752. Opinions are the authors' own and do not represent those of the Commission. The ideas in this paper have greatly benefited from interaction with other partners in the project (Vereniging voor Christelijk Wetenschappelijk Onderwijs; Eötvös Loránd University; Napier University; and Stichting Katholieke Universiteit Brabant).

References

- Axelrod, R. (1997). Advancing the art of simulation in the social sciences. *Complexity*, 3, 16-22.
- Cangelosi, A. & Parisi, D. (Eds.) (2002) *Simulating the evolution of language*. London: Springer-Verlag.
- Edmonds, B., & Moss, S. (2004). From kiss to kids: An 'anti-simplistic' modelling approach. from <http://cfpm.org/%7Ebruce/kiss2kids/kiss2kids.pdf>
- Gilbert, N. (2005). Agent-based social simulation: Dealing with complexity. from <http://www.complexityscience.org/NoE/A/BSS-dealing%20with%20complexity-1-1.pdf>
- Gilbert, N., & Troitzsch, K. G. (2005). *Simulation for the social scientist* (Second ed.). Milton Keynes: Open University Press.
- Hutchins, E., & Hazlehurst, B. (1995). How to invent a lexicon: The development of shared symbols in interaction. In G. N. Gilbert &

- R. Conte (Eds.), *Artificial societies*. London: UCL Press.
- King, L. J. (1985). *Central place theory*: Sage.
- Malinowski, B. (1922 [1978]). *Argonauts of the western pacific*. London: Routledge and Kegan Paul.
- Smith, B. (1995). *The emergence of agriculture*. New York: Scientific American Library.
- Vogt, P. & Divina, F. (2005) Language evolution in large populations of autonomous agents: issues in scaling. Proceedings of AISB 2005