MODELLING SOCIALLY INTELLIGENT AGENTS

BRUCE EDMONDS
Centre for Policy Modelling,
Manchester Metropolitan University,
Aytoun Building, Aytoun Street,
Manchester, M1 3GH, UK.
http://www.cpm.mmu.ac.uk/~bruce

ABSTRACT

The perspective of modelling agents rather than using them for a specificed purpose entails a difference in approach. In particular an emphasis on veracity as opposed to efficiency. An approach using evolving populations of mental models is described that goes some way to meet these concerns. It is then argued that social intelligence is not merely intelligence plus interaction but should allow for individual relationships to develop between agents. This means that, at least, agents must be able to distinguish, identify, model and address other agents, either individually or in groups. In other words that purely homogeneous interaction is insufficient. Two example models are described that illustrate these concerns, the second in detail where agents act and communicate socially, where this is determined by the evolution of their mental models. Finally some problems that arise in the interpretation of such simulations is discussed.

MODELLING AGENTS

At the Centre for Policy Modelling (CPM) we are interested in modelling real agents, which can be people or other institutional units (such as firms or departments). We do this by modelling these as intelligent software agents. This perspective means we have different concerns than those concerned with *designing* agents or robots to meet particular goals – what might be called an engineering perspective. In particular we seek veracity over efficiency. We do not claim that the agent architectures and techniques described here result in agents, or groups of agents, that are particularly good at any specific task. We do claim that these techniques result in communities of agents that exhibit behaviours that characterise boundedly rational socially intelligent agents, i.e. are a step towards modelling some key aspects of humans in social and organisational settings.

One corollary of this is that we do not model using reactive agents, since a principal concern of ours is the nature and development of the agents' internal models as they interact with other agents and its environment (see also the reasons in the section entitled Social Intelligence). The purpose of modelling these agents is to discover the emergent behaviour. If we closely specified the agents behaviour (for example by compiling it down into a reactive architecture) we would be needlessly delimiting the behaviour that might result, which would result in us being less informed about the possibilities inherent in a multi-agent situation.

This leaves open the question of the validation and verification of these models - if you constrain them as little as possible, how do you know if (and how) they correspond to reality in any practical way. The answer we have developed is twofold. *Firstly*, to validate the mechanics of your model by separating out clearly the implementation details from any implicit theory of cognition. [5]. We do this by specifying the implementation it in a language with clearly known semantics (in our case a declarative language), and by basing the agents cognition in a known process or cognitive theory. Secondly by verifying the output of the model against real world qualities or data. This issue is discussed in greater detail in [16].

We take the strategy of explicitly representing the agents internal models in a specified language - usually of a quasi-logical or functional variety. This explicit representation makes it possible to limit, examine and analyse the agents' models, as they develop. We are not claiming that humans use such a representation (this is a hotly debated issue) but merely that by having such inspectable and comprehensible models we can easily find out the state of the agent at any time. We find that by providing agents with a suitably expressive internal modelling language and allowing them to develop their own models, we do not introduce any obviously inappropriate behaviour into our agents. The final test of the appropriateness of an agent’ s behaviour is domain dependent and only verifiable with respect to known properties of what is being modelled.

The agents we model have distinct limitations of resources - they are boundedly rational in several respects. They have limited memory, a limit on searches for improved models and a limit of their ability to make inferences from their models. Following what is known about real agents we ensure that their search for new models is incremental, rather than global in nature. The limitations on the current memory cache, especially the agent’ s stock of candidate models, encodes a sharp path-dependency. The nature and framework of this is described in [7].

One particular technique we use is an adaption of the genetic programming (GP) paradigm [10]. Here the internal models belonging to the agents are held as a set of tree-based expressions. The selection among these is based upon their past predictive success or some endorsement-based mechanism. However we do not always use the cross-over operator as this implements a fairly global search process with is unrealistic for our purposes. Instead we prefer other mixes of operators with a bias away from a lot of crossover, and including other operators that are not used in GP at all such as generalisation and specialisation. For a more detailed discussion of this see [7].

MODELLING ORGANISATIONS

One motivation of working towards capturing social behaviour, is that this is important in real (as opposed to planned or theoretical) organisational behaviour. To take two examples: it is known that emergent norms and other emergent social structures can greatly effect collective behaviour [11] and that social factors can make the prediction of organisational behaviour from the micro-level behaviour very difficult for managers as well as academics [6]. This concern with the veracity of modelling behaviour in organisations can be traced back to Herbert Simon’ s studies [20]. Here in addition to the dimension of procedurality we include sociality as a significant factor.

We model such organisations as populations of interacting agents in a given structure. We do not necessarily do this down to the level of individual persons but sometimes stay at the level of departments or even whole firms (if they are themselves interacting). Here a theme we are investigating is the contrast between the official (usually hierarchical) structure of the firm and the unofficial structures that emerge from the social interactions of the individuals.

In order to study such models, we have developed a modelling language called, SDML – a Strictly Declarative Modelling Language*. This allows the flexible and theory-free modelling of intelligent agents in a declarative framework with object-orientated features, [8]. In particular this is particularly well suited for modelling organisations built up in many levels in a composite manner - it allows for better structured large organisational models involving more complex cognitive agents, compared to some other systems [17].

MODES OF COMMUNICATION IN SOCIAL INTERACTION

Communication is a special case of action and perception. For this reason in many organisational models the communication is rudimentary - where this is implemented on a similar basis to other modelled actions. For example, in many economic models only the price is communicated, and then in a global manner to all the agents equally. However communication is so important in social situations and potentially so computationally onerous that effectively it becomes a separate consideration.

One concern, in this regard is the passivity of the communication in many models of interaction. One gets the impression that in many models agents tend to make requests for information and sometimes issue orders but will not, for instance, volunteer unrequested information. Thus many models of communication correspond to a unidirectional mutual (or merely one-way) transmission of information. For example, in many agent models concerned with negotiation between agents there is an assumption that communication will act using the *pull* of information in the form of requests and replies (in this way it extends the query based dialogue used in interrogating databases to one of interrogating other agents). This is in contrast to the mix of *push* and *pull* modes, found in social interaction - where unrequested information is frequently volunteered.

SOCIAL INTELLIGENCE

Social intelligence implies more than mere interaction with other agents plus intelligence. For example, agents might apply their intelligence to trading with other agents without any intelligence being applied to the *process* of relating to the others. Such an intelligence might be without the means of recognising and referring to other agents as individuals (or groups). In such a case its internal models of its social environment would be entirely generic and thus it could not form any individually social relationship with another agent (different from its relation to any other agent it interacts with). Such a lack of social intelligence has advantages, such as the ability to analyse and predict their behaviour in computational communities. However if we are to model many key behaviours in organisations, we need a greater social sophistication.

Thus in denoting the presence of a social intelligence, It would need to be grounded in at least some of the following:

- a relative sophistication of communicative mechanisms (both in the generation and interpretation of messages);
- the ability to represent aspects of other agents (individually or grouped), in order to anticipate their actions (though this need not involve the *explicit* representation of the other's beliefs, goals etc.);
- the ability to distinguish between and refer to different agents, such that different aspects may be captured for each one (or each group), e.g. in modelling their individual reliability as an information source;
- the ability to direct messages locally to specific individuals (or groups of individuals);
- the presence of purely communicative (social) sub-goals (or even top goals).

SOCIAL INTELLIGENCE AND COMPLEXITY

In addition to the aspects of social intelligence listed above, I think there is another aspect to social intelligence - that of coping with the overwhelming complexity that social systems can (and do) produce. This is a complexity which seems to grow exponentially with the size of the society, [4]. In fact it seems to be a hallmark of social systems that such complexity arises due to the variety of individual specialisations and hence relationships that can develop. A society consisting only of homogeneous agents only equipped with global communication mechanisms will not have the same characteristics.

Luhman has argued that one of our social institutions' primary functions is to *filter out* the complexity of the external world*. This perspective highlights some other important aspects of social intelligence, including:

- the intelligent but restrictive selection of information sources;
- the development of rules to structure social interaction - either formally or informally (e.g. emergent social norms, or formal procedure);
- the development of binding long-term relationships (contracts, friendships, etc.).

A socially intelligent agent may thus seek to use institutions which help it deal with the complexity of social reality. The institution may do this by preforming considerable selection and modelling for the individual. If an appropriate institution does not exist (or is inaccessible) the agent may seek to construct one with other agents. The institutions may also regulate the social structures within themselves by various means such as rules, procedures and sanctions. In this way institutions can have a role not only in effectively simplifying the external reality but also in structuring and hence simplifying the social relationships that are internal to it. Such an institution may itself embed itself within a further institution for similar reasons, resulting in a hierarchical structure. Membership of different institutions covering different aspects of life may result in a parallel, matrix structure. Such an emergence of social structure is thus evidence of social intelligence.

EXAMPLE 1 - A MODEL OF EMERGING MARKETS

A three-sector model of emerging market economies has been developed by Scott Moss where the component firms learn how to behave in a newly emerging market economy (and in particular how to respond to the debt crisis). Firms communicate and trade with other firms and build rule-based models of their environment and the other firms. Thus they distinguish between the other firms around them and model them individually. The firms trade (and hence communicate with) only a set of locally accessible firms. Known qualitative behaviour characteristic of such economies (e.g. the characteristic inflation curve of Belarus) was only reproduced when the model was adapted so that firms would copy the observable behaviour of only the other firms they *interacted* with that they judged as successful. This was not the case when they had global access to information. Here the ability to distinguish, identify, and select other firms was critical to the overall behaviour. More about this model can be found in [18] and [14].

Of course, this model only starts to address the concerns I listed above about social intelligence, but it does show how the inclusion of the ability to individuate other agents and interact in a local and specific manner can critically effect the emergent properties of the whole system.

EXAMPLE 2 - EXTENDING THE 'EL FAROL BAR' MODEL

Description of the Original Model

I have extended Brian Arthur's El Farol Bar model [2] to include model-based learning and communication. In this problem a fixed population of agents has to decide whether to go to El Farol's each thursday night or stay at home. It is generally desirable to go (apparently they play Irish music on Thursday's at El Farol's), but not if it is too crowded. The bar is too crowded if more than 60% of the agents decide to go. Each week each agent has to predict if it will be crowded or not. If it predicts that it will be crowded then it does not go and if it predicts it will not be crowded then it does go. The problem is set up so that if most of the agents share the same model then this model will be self defeating, for if most predict that it *will* be crowded, they do not go and so it will *not* be crowded and *vice versa*.

Brian Arthur modelled this by dealing each agent a fixed number of models randomly allocated from a limited number of types, for example ‘the same as last week’, ‘60 minus the number who went two weeks ago’, ‘the average over the last 4 weeks’, or ‘the number predicted by the trend over the last three weeks’. Then each week each agent evaluates all its models against the past record of how many went and find the best predictive model. It then uses this model to predict the number who will go this week and bases its decision on this. All the agents do this simultaneously and in parallel.

The resulting number who go each week seems to oscillate stochastically about the critical 60% mark, similar to the shape in that in figure 1, despite the fact that this model once initialised is strictly deterministic. Each agent’s repetoire of models is fixed and only the current assessment of the models changes from week to week, so that at different times different models will be most successful. The only communication which takes place is the implicit message of the number who go each week.

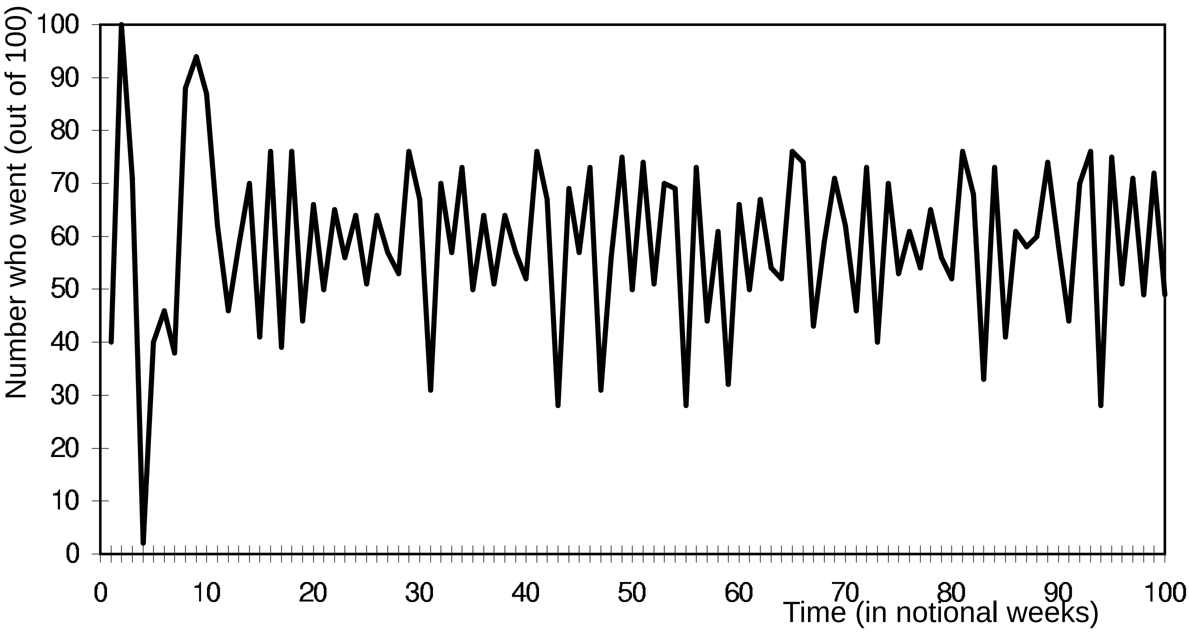

Figure 1. Number of people going to El Farol’s each week in a typical run (with 10 agents)

One of the interesting things in this model is that although each agent is dealt a different menu of models, and all decide to go or not in different combinations and at different times, at the level of whole model they are pretty indistinguishable in terms of their behaviour.

Extending the Model

I have extended this model in several ways. *Firstly,* I have added a rudimentary social structure. A randomised “acquaintance” structure is imposed upon the agents, limiting who they can talk to. *Secondly,* agents have a chance to communicate with acquaintances before making their decision. *Thirdly,* the agents are equipped with a very flexible learning mechanism.

The agent modelling approach broadly follows [7]. Thus each agent has a population of mental models, which broadly correspond to alternative models of its world. These models are each composed of a pair of expressions: one to determine the action (whether to go or not) and a second to determine their communication with other agents. Either of action or communication can be dependent upon communications received, which includes the identity of the agent the communications were received from. These internal models are initially randomly generated to a given depth. Subsequently they are developed in a slow evolutionary manner based either on the past accuracy of the models predictions or some measure of what its past success at gaining utility might be. The agent structure is shown in figure 2. Although the beliefs and goals of other named agents is not explicitly represented, they emerge implicitly in the agents' models.

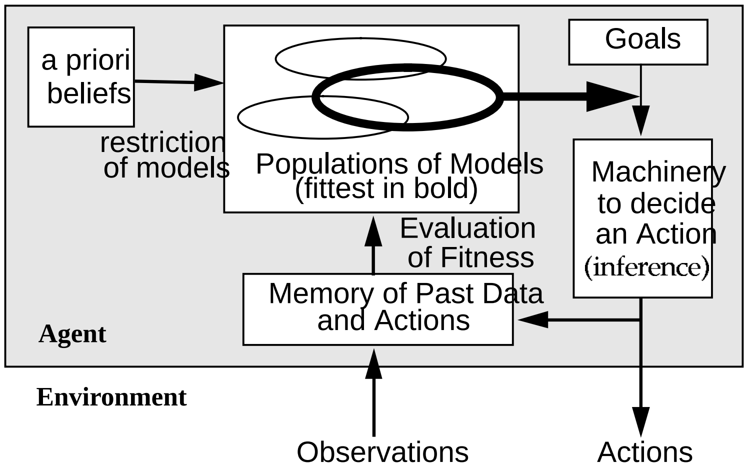

Figure 2. Basic structure of an agent

Each notional week, a new population of models is produced as in a genetic programming manner [10]. In this model we are using some tree crossover but with a high degree of propagation and a few random genes introduced each week. Again I stress that this is not supposed to implement an efficient learning mechanism, but one which deliberately exhibits such qualities such as 'lock-in' and path dependency. The models are evaluated according to a fitness function. The fitness function can be a measure of the past predictive success the models would have had or the utility the agents would have gained by using each model in the past. The best model is then selected and used to determine first its communicative action and subsequently whether to go to El Farol's or not.

The evolution of mental models based upon a GP mechanism is only a rough representation of learning. The cross-over operation is not very realistic for this purpose but does as a first approximation – for a critique of cross-over for these purposes, see [7].

Each model is composed of two parts: one determines what it says and the other whether it will go to El Farol' s. These parts are expressions in a strongly typed language that is specified by the programmer at the start*. A simple but real example model is shown in figure 3 below. Translated this example means: that it will tell its 'friends' that it will go to El Farol' s if the trend predicted over observed number going over two weeks is greater than 5/3 (the total population was 5 in this example); but it will actually go if it said it would go *or* if agent-3 said it will go.

```
talk: [greaterThan [trendOverLast [2]] [divide [5] [3]]]

action: [OR [saidBy ['barGoer-3']] [ISaid]]
```

Figure 3. A simple example model

The possible nodes and terminals of the action and talk models are indicated in figure 4. This language allows a great variety of possible models, including arithmetic, logic, stochastic elements, models based on what other agents have said, on what the agent itself did or said previously and mixtures of all of these. Explicit names for all agents are included in the language as well as some useful arithmetic, and logical constants.

possible nodes for talk gene:
greaterThan lessThan previous times plus minus divide averageOverLast boundedByPopulation wentLag trendOverLast randomIntegerUpTo

possible terminals for talk gene: wentLastTime maxPopulation IPredictedLastWeek 1 2 3 4 5

possible nodes for action gene:
AND OR NOT saidBy

possible terminals for action gene: randomDecision ISaidYesterday IWentLastWeek T F 'barGoer-1' 'barGoer-2' 'barGoer-3' 'barGoer-4' 'barGoer-5' 1 2 3 4 5

Figure 4. Possible nodes and terminals of the tree-structured genes

The amount of utility each agent gains is determined by its action and what the other do: greatest utility is gained by going to the El Farol Bar when it is not too crowded; this is followed by staying at home; going to El Farol's when it is too crowded gains it least utility. Thus each agent is competitively developing its models of what the other agents are going to do. A successful agent learns some way of deciding to go when other do not and avoiding going when many others do.

This model can be seen as an extension of the modelling work in [1], which examines three player games with a similar structure to this model. Their model is, however simpler, for it involves no explicit communication or potential for modelling of other agents and so only touches on the social elements. Revealingly they also notice that being able to distinguish the other players (left and right) is important for complex behaviour to arise.

General Results

Graphs of the results do not show anything surprising. The average fitnesses of the agents populations of models fluctuate wildly at the beginning with a big dif ference between agents but as the simulation progresses they all settle down to around the same value (figure 5).

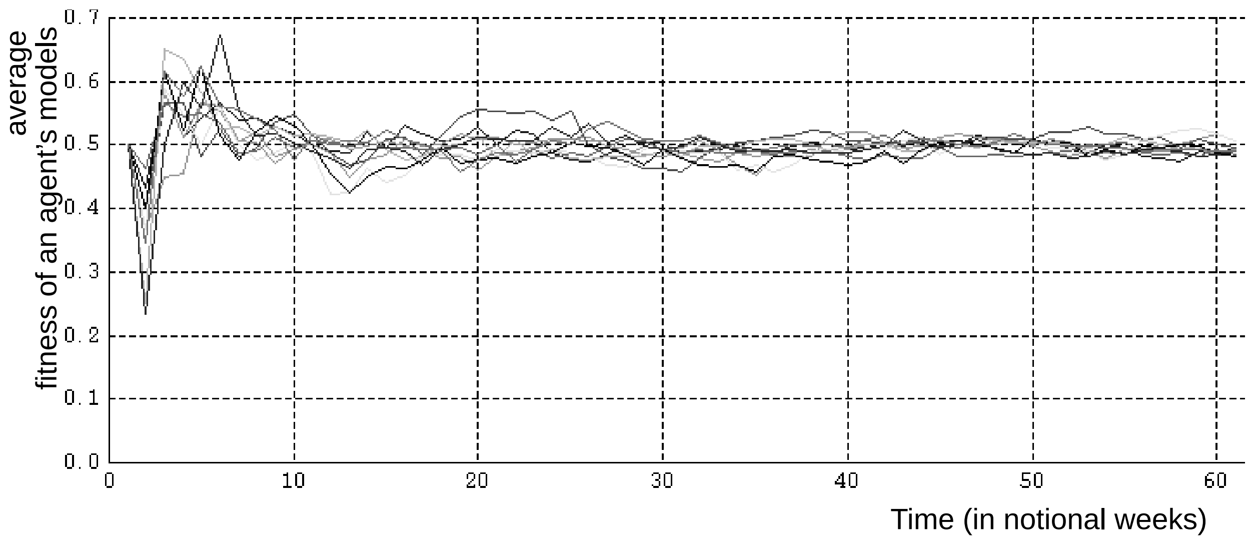

Figure 5. Average fitnesses of the agents' populations of models

The deviance between the fitnesses of the agents' population of models also reduces somewhat but continues to fluctuate at significant levels (figure 6).

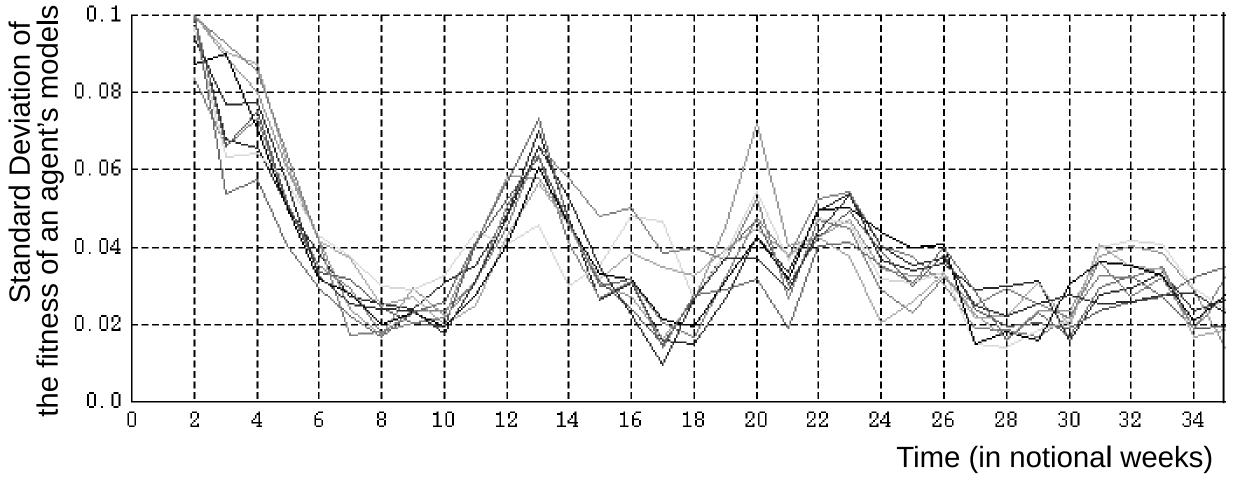

Figure 6. Change in variance (in standard deviations) of the Agents' population of models

The graph of the utility of the agents gained over time shows that different agents are doing best at different periods of time, but none for very long – also that the fortunes of all go up and down somewhat together (figure 7).

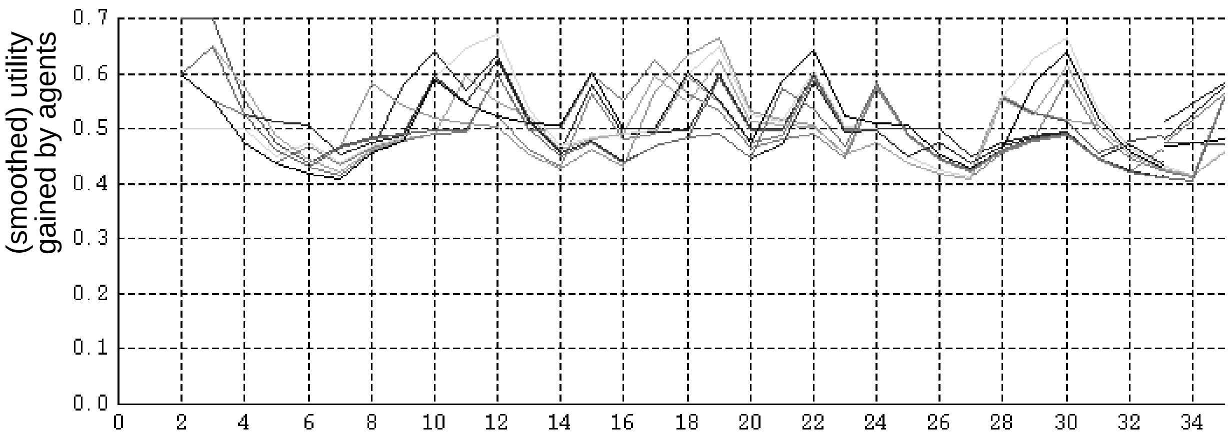

Figure 7. Utility gained by agents over time (smoothed)

When you look at the pattern of who goes and who does not, some of the agents settle for a fixed strategy but some are more dynamic and constantly swap between dif ferent strategies and elaborate old ones (figure 8).

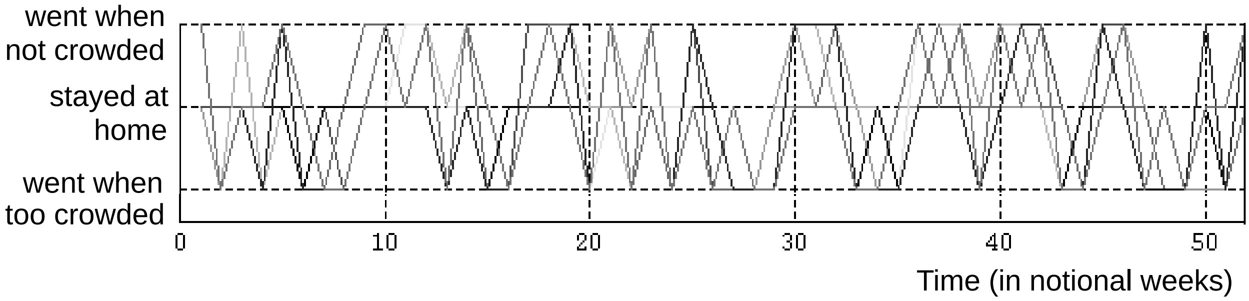

Figure 8. Who did what each week

A Case study from the results

What is perhaps more revealing is the detail of what is going on, so I will exhibit here a case study of the agents at the end of a simulation.

Here I have chosen a five agent simulation at date 100. In this simulation the agents judge their internal models the utility they would have resulted in over the past five time periods. This utility function that agents get is 0.4 if they go when it is two crowded, 0.5 if they stay at home and 0.6 if they go when it is not too crowded (where too crowded means greater than 60% of the total population). This is supplemented with an extra 0.1 of utility for every one of their friends that go if they do.

The friendship structure is chosen at random (within set limits on the minimum and maximum number of friends each can have) at the beginning, and in this case is as show in figure 9.

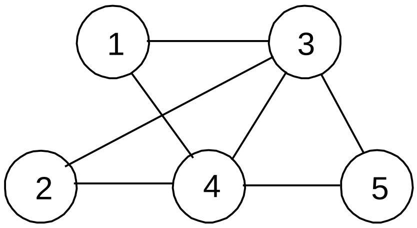

Figure 9. Imposed friendship structure

The best (and hence active) genes of each agent are summarised above in figure 10. I have simplified each so as to indicate is *logical* effect only. The actual genes contain much logically redundant material which may put in an appearance in later populations due to the activity of cross-over in producing later variations.

```
talk-1: averageOverLast(numWentLast) > previous(trendOverLast(numWentLast))
action-1: wentLastTime

talk-2: trendOverLast(numWentLast) - 2 * numWentLag(2)
        > numWentLag(numWentLast)
action-2: NOT Isaid

talk-3: randomNumberUpTo(8) < 8/3
action-3: True

talk-4: averageOverLast(4)/averageOverLast(5) < numWentLag(15)
action-4: (Isaid AND randomDecision) OR (saidBy 2)

talk-5: trandOverLast(20) < numWentLag(2) - averageOverLast(numWentLast)
action-5: randomDecision OR (saidBy 4)
```

Figure 10. Best talk and action genes at date 100

The effect of the genes is tricky to analyse even in its simplified form. For example agent-1 will tell its friends it will go to El Farol's if the average attendance over a previous number of time periods equal to the number who went last time is greater than the predicted number indicated by the trend estimated over the same number of time periods but evaluated as from the previous week! However its rule for whether it goes is simpler - it goes if it went last week.

You can see that for only one agent does what it says indicated what it does in a positive way (agent-4) and one which will do the exactly the opposite of what it says (agent-2). It may seem that agent-1 and agent-3 are both static but this is not so because figure 10 only shows the fittest genes for each agent at the moment in terms of the utility they would have gained in previous weeks. During the next week another gene may be selected as the best.

The interactions are summarised in figure 11, which shows the five agents as numbered circles. It has simple arrows to indicate a positive influence (i.e. if agent-2 says she is going this makes it more likely that agent-4 would go) and crossed arrows for negative influences (e.g. if agent-2 says she will go this makes it less likely she will go). The circles with an “R” represent a random input.

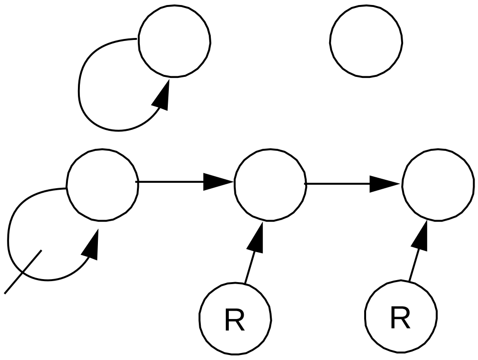

Figure 11. Talk to action causation at week 100

In the original formulation of the El Farol problem one can show that there is a game-theoretic ‘mixed strategy’ equilibrium, where all agents should decide to go according to a random distribution weighted towards going 60% of the time – a behaviour that the deterministic model of Brian Arthur does approximate in practice, [22]. However, *this* model shows the clear development of dif ferent roles*.

By the end of the run described above agent-3 and agent-1 had developed a stand-alone repetoire of strategies which lar gely ignored what other agent said. Agent-3 had settled on what is called a mixed strategy in game theory , namely that it would go about two-thirds of the time in a randomly determined way, while agent-1 relied on largely deterministic forecasting strategies.

The other three agents had developed what might be called social strategies. It is not obvious from the above, but agent-2 has developed its action gene so as to gradually increase the number of ‘NOT’s. By date 100 it had accumulated 9 such ‘NOT’s (so that it actually read NOT [NOT [... NOT [Isaid]...]]). In this way it appears that it has been able to ‘fool’ agent-4 by sometimes lying and sometimes not. Agent-4 has come to rely (at least somewhat) on what agent-2 says and likewise agent-5 uses what agent-4 says (although both mix this with other methods including a degree of randomness).

Thus although all agents were indistinguishable at the start of the run in terms of their resources and computational structure, they evolved not only dif ferent models but also very distinct strategies and roles. They certainly do not all conver ge to the game-theoretic mixed-strategy mentioned above (but a few do). Thus allowing social aspects to emer ge has resulted in a clear dif ference in the behaviour of the model.

Issues and Interpretations

There are some very interesting problems that arise when we try to interpret what occurred. Even given that we can look inside the agents' 'heads' one comes across some of the same problems that philosophers, psychologists and social scientists encounter in trying to account for human communication. The web of cause and effect can be so complex as to impede a straightforward analysis – just as seems to occur in the human case.

One issue in particular is the question of the "meaning" of the agent's utterances to each other. Their utterances do *have* a meaning to each other otherwise they would quickly select out action genes that included "saidBy" clauses. However, these meanings are not obvious. They are not completely determined by their own model structures, but can involve a number of language games whose ultimate grounding is to the practice of such communication in relation to actual decisions. It may be that an approach which uses a Wittgensteinian approach, [21], and describing the state of affairs in terms of 'speech acts' [19] may make for a simpler and more appropriate model of the situation that a traditional AI belief and inference model.

In this particular example it seems that the pragmatics of the situation are the most important for determining meaning, followed by a semantics grounded in the effects of their actions, leaving the syntax to merely distinguish between the two possible messages. This case seems to illustrate Peter Gärdenfors observation about human language:

> *"Action is primary, pragmatics consists of the rules for linguistic actions, semantics is conventionalised pragmatics and syntax adds markers to help disambiguation (when context does not suffice)."*

[9]

CONCLUSION

I have suggested that to model social behaviour one should include the ability to distinguish, identify, model and address other agents individually in a flexible manner, because a social intelligence needs to be grounded in these abilities.

I have exhibited an extension the El Farol model where some of these elements are included. In this model we see the evolution of a society where the intelligence seems to reside, at least somewhat, beyond each agent individually in the web of their interactions - the society of agents. This seems to be due to the co-evolution of the structure together. Such a case illustrates the potential difference between modelling truly social agents where the agents are truly 'bound up' with their society and modelling their cognition individually and then placing them together in a situation where they then interact.

When such social behaviour does occur we may well find ourselves with many of the same difficulties that other social scientists have, namely the tracing of very complex chains of causation if one works in detail and the problem of the meaning and use of our descriptive terms if we attempt a macroscopic approach.

ACKNOWLEDGEMENTS

The original version of this paper was published as: Modelling Socially Intelligent Agents in Organisations, B. Edmonds, 37-42, in Socially Intelligent Agents, AAAI Press Technical Report FS-97-02, ©1997, American Association for Artificial Intelligence.

SDML has been developed in VisualWorks 2.5.1, the Smalltalk-80 environment produced by ParcPlace-Digitalk. Free distribution of SDML for use in academic research is made possible by the sponsorship of ParcPlace-Digitalk (UK) Ltd. The research reported here was funded by the Economic and Social Research Council of the United Kingdom under contract number R000236179 and by the Faculty of Management and Business, Manchester Metropolitan University.

REFERENCES

[1] Akiyama, E. and K. Kaneko, 1996. Evolution of Cooperation, Differentiation, Complexity, and Diversity in an Iterated Three-person Game, *Artificial Life*, 2, 293-304.

[2] Arthur, B. 1994. Inductive Reasoning and Bounded Rationality. *American Economic Association Papers*, 84: 406-411.

[3] Bednarz, J. 1984. Complexity and Intersubjectivity. *Human Studies*, 7: 55-70.

[4] Carneiro, R. L. 1987. The Evolution of Complexity in Human Societies and its Mathematical Expression. *International Journal of Comparative Sociology*, 28: 111-128.

[5] Cooper, R., J. Fox, J. Farringdom and T. Shallice 1997. A Systematic Methodology for Cognitive Modelling. *Artificial Intelligence*, 85: 3-44.

[6] Earley, P. C. and J. Brittain. 1992. Cross-Level Analysis of Organizations - Social Resource-Management Model. *Research in Organizational Behaviour*, 15: 357-408.

[7] Edmonds, B. and S. J. Moss, 1997. Modelling Bounded Rationality using Evolutionary Techniques. AISB'97 workshop on Evolutionary Computation, Manchester. *Lecture Notes in Computer Science*, 1305: 31-42.

[8] Edmonds, B., S. J. Moss, and S. Wallis, 1996. Logic, Reasoning and a Programming Language for Simulating Economic and Business Processes with Artificial Intelligent Agents. In Ein-Dor, P. (ed.): *Artificial Intelligence in Economics and Management*. Boston: Kluwer Academic, 221-230.

[9] Gärdenfors, P. 1997. The pragmatic role of modality in natural language. 20th Wittgenstien Symposium, Kirchberg am Weshel, Lower Austria, 1997.

[10] Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press.

[11] Latane, B., A. Nowak and J. H. Liu. 1994. Measuring Emergent Social Phenomena - Dynamism, Polarization and Clustering as Order Parameters of Social Systems. *Behavioural Science*, 39:1-24.

[12] Luhmann, N. 1996. Soziale Systeme: Grundriß einer allgemeinen Theorie; 6. Auflage; Franfurt am Main, S.45ff.

[13] Montana, D. J. 1995. Strongly Typed Genetic Programming, *Evolutionary Computation*, 3: 199-230.

[14] Moss, S. J. forthcoming. Boundedly versus Procedurally Rational Expectations. In

[15] Moss, S. J. and B. Edmonds, forthcoming. Modelling Economic Learning as Modelling. *Systems and Cybernetics*.

[16] Moss, S. J., B. Edmonds, and S. Wallis, 1997. Validation and Verification of Computational Models with Multiple Cognitive Agents. Centre for Policy Modelling Report 97-25, MMU, Manchester, UK. (http://www.cpm.mmu.ac.uk/cpmrep25.html)

[17] Moss, S. J., H. Gaylard, S. Wallis, and B. Edmonds, forthcoming. SDML: A Multi-Agent Language for Organizational Modelling. *Computational and Mathematical Organization Theory*.

[18] Moss, S. J. and O. Kuznetsova, 1996. Modelling the Process of Market Emergence. In J. W. Owsinski, and Z. Nahorski, (eds.). *Modelling and Analysing Economies in Transition*, MODEST, Warsaw, 125-138.

[19] Searle, J. R. 1969. *Speech acts : an essay in the philosophy of language*. London: Cambridge University Press.

[20] Simon, H. 1957. *Administrative behaviour: Astudy of decision-making processes in administrative organization*. New York: Macmillan.

[21] Wittgenstein, L. 1953. *Philosophical Investigations*. Oxford: Blackwell.

[22] Zambrano, E. 1997. The Revelation Principle of Bounded Rationality. Sante Fe working paper 97-06-060. (http://www.santafe.edu/sfi/publications/Abstracts/97-06-060abs.html)