The Meme as a Design Pattern for Social Learning in Agent-Based Models

Andrew Rixon[1], Peter Robinson[2] and Fabio Boschetti[1]

1-CSIRO
2- Information Technology and Electrical Engineering, The University of Queensland, Australia

Abstract:
Agent-based models have become popular within complex systems science for their ability to model a system of many components with simple interactions. In particular, their ability to demonstrate surprise and emergent properties have lead to their use in many varied situations involving social simulation. Little work so far has been focussed on leveraging the collective wisdom for approaches to building better agent-based models at the architectural level. This paper proposes a useful design pattern for capturing social learning within agent-based models. Moreover, this paper indirectly asserts that this provides some value, functionality and usefulness to the concept of 'meme' which has been lacking since the conception of the field of memetics.

Introduction

Complex systems is a relatively new area of research that is primarily focused on systems made up of many, but typically simple components. Of particular interest are the *emergent* properties of such systems. Emergent properties are global properties or behaviours exhibited by the system that are often difficult, or impossible, to predict from the properties or behaviour of the individual components. Complex systems research covers a wide variety of areas, from biology to economics and air traffic control.

When we study a complex system via Agent Based Modelling (ABM), we need to provide agents with assumptions and behaviours, a social network to interact with and some rules for changing, adapting and evolving their behaviour as a result of the interactions. Similarly, when modellers discuss standards and frameworks for developing, exchanging and sharing ABM related codes, they also need to interact within the research network, negotiate their requirements and adapt their coding practises, possibly by adopting other's people standards. How these processes develop is no different from the dynamics of other items of human culture. It is complex, very non-linear, hardly predictable, and subject to evolutionary forces, which in general guarantee improvements, but not optimality[1].

Thus, assigning and adapting behaviours can be seen, as if via a magnifying glass, at the different levels of a computer agent, a programmer, a negotiating researcher and a social being. This nesting of similar dynamics at multiple levels leads inevitably to complexity and the emergence of unexpected outcomes. It also suggests that each individual level could be better understood by analysing the common features of the larger system and maybe by importing tools and ideas from one level to the other.

[1] Whoever uses a computer today knows the standard he/she is subjected to, though improving over the years, did not develop with optimality in sight, but rather from a complex mixture of serendipity, market forces, historical curiosity and human vagaries

The relative immaturity of complex systems research has tended to mean that researchers have produced ad-hoc simulations and visualizations of systems of interest. For agent-based modelling approaches, drawing insights and common features from the larger system is in its infancy and as the area matures, it is expected that common approaches (patterns) will be identified that will help simplify the design, development and implementation of agent-based model simulations and visualizations.

With our attention directed to social learning for agent-based modelling we propose that the concept of meme, that is "a unit of cultural transmission, or that which is imitated" coupled with Alexander's notion of 'patterns', can be the unifying concept, providing a useful framework for addressing social learning in agent-based modelling and simulation.

Patterns

Patterns were first identified in Christopher Alexander's seminal text "A pattern language" (Alexander *et al.* 1977). In this text Alexander describes a pattern as follows:

"Each pattern describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice".

Since then, the concept of patterns has become extremely influential in software development and design with the Gang of Four's seminal text called "Design Patterns" (Gamma *et al.* 1995).

Using design patterns in software development can speed up development processes by allowing practical lessons learnt to be leveraged out of related or similar software projects. Effective software design requires considering issues that may not become visible until later in the implementation, thus emphasising the importance of formulated design pattern approaches. Reusing design patterns helps to avoid such issues that can cause major problems during the software development process.

There has been a good literature developing on how to build agent based models (Wooldridge *et al.* 1998; Gilbert *et al.* 1999; Gilbert *et al.* 2002; Rixon *et al.* 2005) with (Hare *et al.* 2004) providing a useful taxonomy to guide the use of agent-based models. In particular (Hare *et al.* 2004) suggest that Agent-Based Models are most appropriate when there is:

- Coupling of environmental and social processes
- Individual decision making
- Social interactions and
- Adaptive behaviour.

There is, however, little in the way of specific instruction on how to overcome certain problems encountered when building agent-based models. In particular, with social learning processes becoming fundamental to many agent-based models, there is a real need for robust methods of incorporating social learning elements and processes into agent-based models.

In software development, design patterns allow developers to communicate using well-known, well understood names for software interactions. Common design patterns can be improved over time, making them more robust than ad-hoc designs. It is the aim of this paper to leverage the design pattern approach to inform and create discussion on how to build better agent-based models.

In particular, this paper addresses social learning within agent-based models and proposes that the meme be considered as a design pattern for modelling social learning within such models. The aim is to initiate discussion within the research community so that this pattern can be refined into a useful pattern for application in agent-based modelling.

In the next section we present a brief review that highlights the use of memes for social learning. Following that, we discuss the meme as a design pattern. We conclude with some remarks about memes as a potentially useful design pattern along with the limitations.

Social Learning and Memes

Social learning is described at en.wikipedia.org/wiki/Social_learning as:

"Observational learning or social learning refers to learning that occurs as a function of observing, retaining and replicating behaviour observed in others. It is most associated with the work of psychologist Albert Bandura, who implemented some of the seminal studies in the area and initiated social learning theory."

The principles underlying Bandura's social learning principles (Bandura 1961) are:

> 1. The highest level of observational learning is achieved by first organizing and rehearsing the modeled behavior symbolically and then enacting it overtly. Coding modeled behavior into words, labels or images results in better retention than simply observing.
>
> 2. Individuals are more likely to adopt a modeled behavior if it results in outcomes they value.
>
> 3. Individuals are more likely to adopt a modeled behavior if the model is similar to the observer and has admired status and the behavior has functional value.

It has been suggested that social learning may affect behaviour in the follow ways (Bandura et al 1977):

- Teaching new behaviours
- Increasing or decreasing the frequency of which previously learnt behaviours are carried out
- encouraging previously forbidden behaviours

- increasing or decreasing similar behaviours. For example, observing a model excelling in piano playing may encourage an observer to excel in playing the saxophone.

Bandura's principles of social learning and the associated adoption by individuals leads naturally to the consideration of memes.

Dawkins defined a meme as "a unit of cultural transmission, or a unit of imitation" (Dawkins, 1976). While several other definitions have been proposed, it is generally agreed that a meme is a self-propagating "mental construct" of some kind (e.g. an idea, a recipe, an instruction or a theory), whose dynamics resembles, to some extent, that of genes, since a meme (like a gene) contains 'instructions' for people actions/behaviours, can be copied (possibly with errors) and can undergo selective pressure. Memes have been suggested as drivers of cultural evolution, and their intellectual appeal appears to center on two key concepts. The first considers memes adhering to a strong form of Darwinian evolution (Universal Darwinism); in the second, meme evolution is de-coupled from the dynamics of the humans/communities who share them (Blackmore 1999). Adopting this view, humans are simply hosts or carriers of memes, whose propagation depends on the memes' own selective potential, not on the benefit they bring to the hosts.

Mechanisms of meme transfer

Critics of memetic theory have focused on two core points: first, the existence and dynamics of real-world memes is hard to demonstrate and, so far, has not offered better explanatory power or predictability than other theories (i.e. it has not replaced alternative scientific models); second, it is not clear what mechanism controls a meme's selective potential, without relying on the tautology 'a meme which gets replicated is fit – a meme is fit if it gets replicated'.

We propose a slightly different framework which attempts to replace the abstract notion of a meme's 'selfish interest' detached from the dynamics of the host. We suggest a way in which the selective potential of the memes and of the host may align, thereby conserving the intuitive notions of memes/patterns we discuss in this work.

This can be summarised in the following conjecture: *memes replicate because they are perceived (whether correctly or not) as a means of efficient prediction*. Since efficient prediction is crucial to survival, memes take advantage of agents' inherent propensity towards economic, efficient predictions in order to replicate.

There are three components to this conjecture:

1) 'efficient prediction', which refers to a prediction obtained at low computational cost (not to whether the prediction is accurate or not);
2) 'perceived' refers to the fact that the prediction does not necessarily have to be correct in order to allow a meme to replicated. It means that an agent may be attracted to a meme because of the *perception* that it offers efficient predictability;
3) *'whether the prediction is correct or not*' determines whether a meme is useful or not to the carrier. Cases in which the prediction turns out to be wrong can be seen as an example of 'selfish' meme which is detrimental to the agent.

Crucially, this conjecture allows for the arising of emergent behaviour in agents communities because of the self referentiality which arises when sharing a prediction between a number of agents can affect the outcome of the prediction itself. Because human behaviour depends on its context, which in turn depends on the actions of other humans, the prediction a meme induces can be reinforced, *materialised*, or defeated by its very own replication. Thus the *perception* of the predictability offered by a meme is important, rather than its mere accuracy. Whether good or bad, a meme which propagates because of a *perceived* predictability power, may determine a specific outcome, thereby accounting for the proliferation of so-called self-fulfilling or self-defeating prophesies.

As an example pertinent to this work, imagine a community of ABM modellers who has to agree between 2 coding standards (A and B) with apparently comparable benefits and drawbacks. At the beginning of the decision process, both standards (memes/patterns) have the same probability of spreading and the same potential 'worth' for all agents. Also they offer the same potential effectiveness for the users to predict upon which standard they should develop their codes. During the process, however, the initial (possibly random) adoption of a specific standard (meme) by the majority of agents will improve its prediction capability (since it will suggest it may succeed), its efficiency and, in turn, further enhance its spread. This is a clear example of how a meme can turn itself into a kind of reality by enhancing its own predictive efficiency. Importantly, the succeeding standard (meme/pattern) may not be the most profitable[2], rather the one which gets 'locked' (and it in turn locks in the users) by a first, possibly random, adoption.

Apart for random fluctuations, another reason for the spreading of non optimal memes may lie in the 'cost of adoption'. The cost of prediction is important, because it can favour the spreading of a meme even if its predictions are inaccurate. Given that it is difficult to define a good strategy for an agent, it may be much easier (and more efficient) to stick to a plausible meme (whether effective or not) than to waste a considerably amount of resources in trying to predicting the behaviour of the complex environment. Perception of prediction then becomes more important that accuracy.

There may be other mechanisms that allow an inaccurate (but efficient) meme to propagate and in the following we draw some simple vignettes to describe them.

Vignette 1; Adopt a locally beneficial meme. Meme A may not be the most 'beneficial' for a community, however, if meme A is shared by the majority of people, a 'better' meme B can display its more beneficial power only provided the majority of the community switches to it at the same time. In other words, a meme may not be the most efficient predictor, but *locally* it may offer an efficient way of predicting the environment. In this situation, a non beneficial meme may survive and replicate for a long time, possibly indefinitely, until a major upheaval occurs.

Vignette 2; Copy, rather than experiment. In the most naïve view, an agent imitating a neighbour's behaviour (one of the proposed meme transmission

[2] A real world example of this process has been the adoption of VHS versus Betamax standard in the first days of video-recording.

mechanisms) not only can provide the agent with novel ideas/procedures (which it may have not been able to develop itself) but also can be seen as a very cost effective way to predict an outcome. Say I want to cook a soup. I can experiment with various ingredients, amounts and cooking time, and improve my ability to predict the taste of the soup as my experience increases, or I can just use a recipe from a friend whom I know is a good cook or from a book or which I have tasted already. Not only shall I save the effort of experimenting, but also I can reasonably predict that the soup will be tasty. In terms of the efficiency of prediction, following a recipe improves the predictability *and* decreases the cost of the prediction.

Vignette 3; Adopt non-verifiable beliefs. A religion provides a set of beliefs. Strictly speaking, such beliefs cannot be verified and thus crude prediction is not possible. However, the efficiency of prediction concept can be applied in two ways. First, and more trivially, what is important is not the accuracy of the predication in itself, but rather the *perception* that prediction is possible. Given metaphysical, existential or challenging questions, an agent can either consider becoming a scientist, thus spending an entire life researching the origins of life (with a strong probability of discovering little), or it can adopt the Biblical account at a much lower cost. Since the prediction of each approach is unverifiable anyway (as well as practically irrelevant), the Bible appears to be a more accessible (i.e. efficient) way out of this enigma. On a more subtle level, religion carries a number of ethical values, rules and policies that are shared by a community. Accepting those values means sharing them with your neighbours, whose behaviour then becomes easier to predict. Living in a society based on common assumptions and expectations makes the agents' environment more predictable and consequently easier to live in.

Social Learning within Agent-Based Models

Learning within agent based models can be thought of as the acquisition of knowledge. Knowledge can be broadly divided into two kinds: information (typically called beliefs in the agent community); and behaviour. Acquisition of knowledge can happen in many different ways, for example:

- information from the environment can update the beliefs of the agent
- the agent may deduce or induce new beliefs based on existing beliefs
- agents may learn behaviours using, for example, neural networks or genetic algorithms
- agents may acquire beliefs or behaviours either directly (via communication) or indirectly (via observation) from other agents

For the purposes of this paper, the fourth of the points above can be considered as social learning within the agent. This corresponds closely with Dawkins idea of the meme – a unit of cultural transmission corresponds to a belief that can be acquired, and a unit of imitation corresponds to a behaviour.

From an agent implementation point of view, it is necessary to consider how social learning can be implemented. Acquisition of information such as the water usage of another agent can be implemented simply by updating a belief store. On the other hand, behaviour acquisition requires updating the possible behaviours of the agent. Note that information such as how to compute the water usage for a particular activity or how to play a piece of music is more like a belief that a behaviour in this context.

In either case, the code of the agent needs to be modified in some way when this sort of learning takes place.

One way to implement this in an agent is to provide each agent with all the possible code that might be needed and only "turn on the code" when it is learned. Another option is for the agent to receive code in a message from the agent it is learning from. A third option is to share code, but to modify behaviour by learning new parameter settings for the code. A fourth option is to observe a behaviour of another agent as a series of actions and to imitate that series of actions.

Examples of Agent-Based Models using Memes

Before proposing the meme design pattern it is important to emphasise that the process of proposing patterns is very much more about recognition / elucidation / description of a pattern already in existence, rather than proposal / creation of a new entity or creature of pattern. Indeed, this may be the very nature by which a pattern obtains its power. In this light we seek to demonstrate three examples where the meme as a pattern has emerged within the literature.

Example 1:

A multi-agent system called the water memes model explores the impacts of memetic transfer of water conservation behaviours and social networks on urban water usage conservation (Rixon *et al.* 2004). In the water memes model, the resident agents are considered to have a selection of water memes, some which are water saving memes and some which are not. These memes have a direct mapping between the type of device, the frequency of use and the amount of water used. In this model, ten types of water meme are used (see Table 1).

Table 1. Names and descriptions of the ten water memes used in the water memes model

Name	Description
Garden	Five types of garden watering technology available, each with differing water use. These are bucket, hose, hose sprinkler, fixed sprinkler and drip system
Shower	AAA-rated shower head is 45% more efficient than normal head
Toilet	Dual-flush toilet saves water over the single flush
Brush teeth	Brushing the teeth using a glass saves over having a tap running
Prepare food	Preparing food in the kitchen

	sink with the plug in is more efficient than with a running tap
Wash clothes	Ensuring that the washing machine is full each wash
Leaking tap	The knowledge of how much water is saved by stopping leaking taps
Leaky toilet	The knowledge of how much water is saved by diagnosing a leaky toilet
Dishwasher	Whether dishes are washed in the kitchen sink or in a dishwasher
Rainwater tank	Installing a rainwater tank can reduce external water demands

Water memes are copied based on social interaction. Residents within a household tend to become more like each other, sharing a common set of beliefs. Resident agents who are water worried look to their friendship networks to seek out water savers.

The key algorithms for the propagation of memes (and hence the behaviours for this model) belong to the resident agent and are the SeekSimilarBeliefs and SeekWaterSavers functions.

Example 2:

The next example investigates the appearance of the meme in a model for consumer behaviour that has been developed to address two key aspects of peoples learning behaviours. Those key aspects are:

1. People do not always optimise their outcomes, but often engage in satisfying behaviour.
2. People often learn about new attractive behaviours by using information about other people's behaviours.

The model is one that integrates these assumptions and is called the Consumat approach provided by (Jager *et al.* 2000).

The behaviours available under the consumat model are:

- Deliberation
- Repetition
- Imitation
- Social comparison

These behaviours appear based on the degree of satisfaction and uncertainty which the agent faces in terms of making decisions. The following figure describes the relationships, see (Jager et al 2000) for a full description.

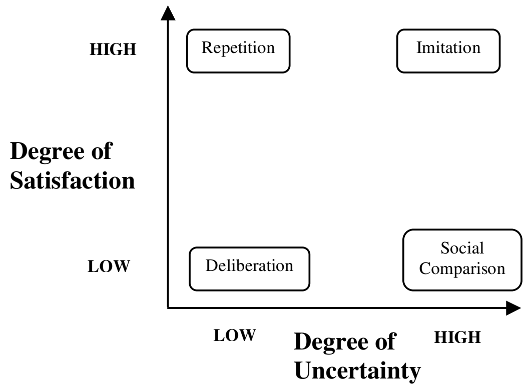

Figure 1: Consumat model for behaviour choice dependent on degrees of satisfaction and uncertainty

The most important aspects of this example from the point of view of this paper are as follows.

Social comparison – stands for reasoned social processing and relates to social comparison theory (Festinger 1953). In particular social comparison theory states that people have a drive to evaluate the behaviour of similar people when they feel uncertain about which behaviour to perform. Moreover, the behaviour of other people is often considered as a norm. The key to this behaviour is the updating of the agents mental map and either the copying of the other agents behaviour or the maintenance of its own behaviour.

Imitation – stands for automatic social processing and relates directly to social learning theory proposed by (Bandura 1977). In this model a cosumat engages in imitation if it is unsatisfied and uncertain. In particular, the agent will do what the imitatory agent did in the previous time-step. It is thus motivated to enact the same behaviour which the comparison / imitatory agent performed to maintain its satisfaction.

Important to our case here is that in both cases, that of imitation and social comparison, another agent's behaviour is copied and becomes part of the behaviour of the current agent.

Example 3:

The final example comes from a discussion from a review article (Hegselmann 1998) for Gaylord's book on Simulating society (Gaylord *et al.* 1998):

"While chapter 1 is basically devoted to modelling and programming individuals which move, chapter 2 equips these individuals with attributes which then change as a result of social interactions. The attributes considered are cultural memes, represented by a meme list, with each component having an integer value taken from a predefined interval. In a first variant of the model, individuals will adapt one of the different memes of a facing other. That will happen with a probability which is higher the more meme values the facing individuals already share. Who adapts to whom is a matter of chance. In further models, exchange of meme values depends on similarity of the values of randomly selected memes in the meme list or on social status which is added as a third feature of an individual in addition to facing direction and the meme list."

Interestingly, the model described above makes use of attributes considered as cultural memes represented by meme lists not dissimilar to those found in Example 1.

The Meme Design Pattern

In the previous section we illustrated three examples where memes or meme equivalents were used within an agent-based model. We now propose the meme design pattern which addresses how to model and implement social learning into agent-based models.

The meme design pattern aims to inform implementations and modeling of social learning, that is, the transfer and distribution of behaviour and beliefs throughout the population / society. In particular this design pattern is proposed for heterogeneous desires of agents and agent systems which require social learning.

This meme design pattern provides a flexible design for the uptake of behaviours and beliefs, addressing social learning within the agents. In particular the meme pattern provides the ability to design agent-based models with:

- Behaviour and belief enhancement and uptake
- Social learning
- Meme hierarchies that provide a scalable approach to implementing decision making behaviour into agents with applicability both at the low local level as well as the global meta level where decisions about decisions can be made.

We suggest that this meme pattern can be used in situations where the meme can be considered as a behavioural or belief unit for an agent, there is a need to manage meme's within the system and a specific need for social learning within agent-based simulations.

Figure 2 illustrates a Unified Modelling Language (UML) description for the meme design pattern. The core elements defined within this pattern are the Meme, Agent and Network. Note: The "behaviour" functionality in the Meme object can also be considered as "belief".

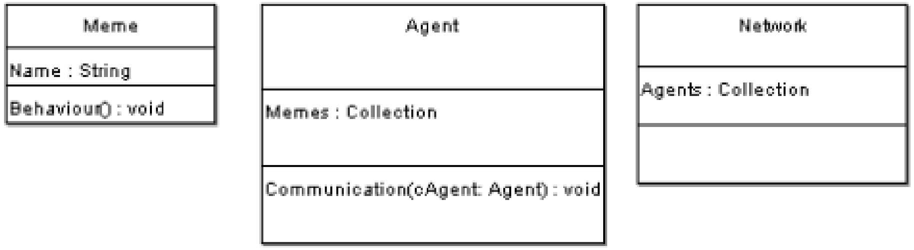

Figure 2: UML Diagram detailing the participants of the Meme proto-pattern

That is, a minimal description is that there needs to exist a network of agents each with communication abilities between each agent. Each agent in particular contains a memes collection and each specific meme within this collection has at least a name and a behaviour. Table 2 below captures the key participants and responsibilities for the objects contained within the meme pattern.

Table 2: Participants and responsibilities for Meme proto-pattern

Class	Responsibilities
Meme	Provide a meme
Agent	Has a collection of memes Communication with other Agents
Network	Has a collection of Agents Provide a network (social) for agents communication

The meme design pattern is considered to provide the following benefits. Firstly, considering the meme as a behaviour or belief provides a useful abstraction for approaching social learning within the design and development of agent-based models. Secondly allowing agents to contain a collection of memes allows for enactment and transfer of memes within the population of agents

An Example of Pseudo-Code for Meme Design Pattern:

As discussed in (Rixon et al. 2005) In the case of the water memes model, the time unit used was a day so the core algorithm used was:

```
For intSteps = 1 To NumberOfTimeSteps.
   For Each oHousehold In oHouseholds
Call oHousehold.DoActivities(oRainfall(intSteps)), oEvaporation(intSteps))
dblMeanTotalWaterUse=dblMeanTotalWaterUse+oHousehold.dblHouseholdWate
rUse
    Next oHousehold
Next intSteps
```

Resident agents are implemented within the model to simply perform their daily activities such as:

- take shower;

- use the toilet;
- brush teeth; and
- wash clothes and dishes.

Each activity in which the agent engages, and the associated use of water, is driven by the agent's beliefs (their memes) for that activity. Thus, for example, the 'take shower' behaviour and associated water use simply becomes:

```
Public Function TakeShower() As Double
        TakeShower = oMemes.GetWaterUse(oMemes.oMyMemes("Shower"))
End Function
```

Using the meme to encapsulate behaviour means that it is simple for agents to swap their memes and have the characterising behaviour follow, since the implementation (here, the actual water consumption) is linked to the object (in this case, the ShowerMeme).

Considerations for the Meme Design Pattern

In deciding to use the meme design pattern it is important for the designer / developer to consider some key questions which the authors consider important to the meme design pattern approach.

Firstly, how large can the network of agents be before there is significant computation impact? It may be the case that the number of agents required within the simulation is simply too large to allow for the meme design pattern to be used.

Secondly, degrees of connectivity of the network of agents need to be considered for impact on agent learning. It may be that the system responds with agents that never learn, or attempt to learn everything!

Thirdly, questions are raised on deciding the 'suitability' of behaviour acquisition and retention in the model. A choice needs to be made on how agents communicate with each other (eg. seeking or broadcasting) and also what kind of filtering occurs at the meme level (acquisition or retention).

Finally, moving the memes in and out of the agents assume that the recipient of a meme can actually use the acquired meme. Thus it is crucial that the behaviours encapsulated by memes remain valid to an agents input/output requirements for 'acting' on a specific behaviour.

Generally, the meme design pattern is not useful for:

- Fixed behaviours or beliefs within agents
- No need for behaviour acquisition
- No network of agents
- No need for learning capabilities within agents

Discussion

Whether ABM modelling will ever be able to model the complexity and emergent properties of real societies is today a subject of heated debate and approaching that aim will require considerable advances in the understanding of social processes as well as theoretical computer science. The contribution from this paper is to suggest that the study of memes can offer a good test case for the convergence of these 2 fields as well as provide some practical avenues to advance applied ABM modelling.

The main benefits involved in this 'meme' approach can be summarised as follows.

1) They provide a framework for modelling the development and spread of ideas which is related to, but not a mere consequence of, the advantages they provide to the agents. This goes past the rigid and inflexible Darwinism as well as the picture of the traditional 'homo economicus' towards a more realistic and unpredictable human dynamics; a meme success is no more a consequence of the agents evolution. The meme is no longer a detached process (as in standard memetics) but instead a dynamical interacting process, affecting and being affected by the agents they inhabit.
2) They provide for a step towards overcoming the dichotomy between belief and behaviour of classic ABM. A meme, being a unit of replication, blurs such distinction. This relates to recent trends in computer science and in particular in the insight that interaction (among agents and between agents and the environment) and the inherent parallelism which comes with it confound the traditional distinction between data and program (Wiedermann J., van Leeuwen, 2002; Milner, 1993).
3) The points above lead naturally to a scenario in which the distinction between agents and memes can also be blurred (memes can become agents with their own dynamics) and the modelling of different combinations of these approach can be modelled and explored.

Our approach is pragmatic. We do not concern ourselves on whether memes 'exists' in Nature, or what they may 'be' in the context of human affairs. Rather, we are interested in using them as 'tools' that may help us elucidate some aspects of human dynamics and provide avenues to streamline the development of both standards and applications in the ABM community.

In this paper we have proposed a meme proto-pattern based on examples of agent based systems that engage in social learning. In particular we have demonstrated that there are already several examples in the literature of how a meme has proven itself useful within the context of providing social learning to agent-based models. Moreover, we believe that treating the meme as a pattern provides a great utility to an otherwise struggling concept.

Our objective with this paper is to stimulate discussion about this proto-pattern with the hope that it can be extended to a useful software pattern. Ultimately, we hope that a meme pattern will be of benefit in applications involving a network of communicating agents that engage in social learning processes.

References

Alexander, C., S. Ishikawa, et al. (1977). A Pattern Language: Towns, Buildings, Construction, Oxford University Press.

Bandura, Albert, Ross, Dorothea, & Ross, Sheila A. (1961). Transmission of aggressions through imitation of aggressive models. Journal of Abnormal and Social Psychology, 63, 575-582

Bandura, A. (1977) Social Learning Theory. New Jersey: Prentice Hall. Hardback: ISBN 0138167443 - Paperback: ISBN 0138167516

Blackmore, S. (1999). The Meme Machine, Oxford University Press.

Dawkins, R. (1976). The Selfish Gene, Oxford University Press.

Edmonds, B. (2002). "Three challenges for the survival of memetics." Journal of Memetics - Evolutionary models of information transmission **6**: 1-6.

Gamma, E., R. Helm, et al. (1995). Design Patterns, Addison-Wesley Professional; 1st edition (January 15, 1995).

Gaylord, R. J. and L. D'Andria (1998). Simulating Society - A Mathematica Toolkit for Modeling Socioeconomic Behavior, New York, NY: TELOS/Springer Verlag.

Gilbert, N. and S. Bankes (2002). "Platforms and methods for agent-based modeling." Proceedings of the National Academy of Sciences **99**(3): 7197-7198.

Gilbert, N. and P. Terna (1999). "How to build and use agent-based models in social science." Discussion paper.

Hare, M. and P. Deadman (2004). "Further towards a taxonomy of agent-based simulation models in environmental management." Mathematics and Computers in Simulation **64**: 25-40.

Hegselmann, R. (1998). "Simulating Society - A Mathematica Toolkit for Modeling Socioeconomic Behavior : Review Article." Journal of Artificial Societies and Social Simulation **1**(3).

Jager, W., M. A. Janssen, et al. (2000). "Behaviour in commons dilemmas: Homo economicus and Homo psychologicus in an ecological-economic model." Ecological Economics **35**: 357-379.

Rixon, A. and S. Burn (2004). "Exploring water conservation behaviour through participatory agent-based modelling." IFAC Wokshop on Modelling and Control for Participatory Planning and Managing Water Systems Venice 29 Set-1stOct 2004, Polipress, Milano, 2004.

Rixon, A., M. Moglia, et al. (2005). "Building an ABM from scratch – Discussing the need for a platform." CABM-HEMA-SMAGET Joint Conference on Multi-Agent Modelling for Environmental Management Bourg St Maurice – Les Arcs - France 21- 25 march / mars 2005.

Wooldridge, M. and N. R. Jennings (1998). Pitfalls of Agent-Oriented Development. 2nd International Conference on Autonomous Agents.

Wiedermann J., van Leeuwen, J: The Emergent Computational Potential of Evolving Artificial Living Systems. Ai Communications, IOS Press, Vol. 15, No. 4, 2002, pp. 205-216

Milner, R., 1993, Communications of the ACM archive, Volume 36 , Issue 1, 78 – 89