On Modelling in Memetics

Bruce Edmonds
Centre for Policy Modelling
Manchester Metropolitan University
http://www.cpm.mmu.ac.uk/~bruce

Abstract

The field of memetics is characterised and two types of memetic model analysed: the *a priori* model and the 'black-box' model. These are used to motivate a picture of scientific modelling, where by chains of models are built from abstract models at the 'top' down to data models derived from measurement. These models are linked to each other by *validation* (where the structure of a model comes from another) and *verification* (where the results of a model are checked against another). The role of formal models in clearly establishing these links are discussed.This picture is applied to the process of modelling in memetics. The paper argues that stronger selection are required to enable the field to evolve. In particular it argues that validation and verification constitute a minimum such set for memetic models.

Keywords: meme, modelling, philosophy, science, selection, explanation, constructivism, validation, verification

1 Introduction

Memetics as a serious science has just begun. As in similarly young fields there is an initial proliferation of possible models and lots of informal 'chat' about the proper designation of its subject matter. In this article I take a step back and apply the evolutionary metaphor to the field of memetics itself. I conclude that in order for the field to evolve, better and more stringent selection criteria should be applied. In particular I advocate the validation and verification of models with respect to other models, so that they form a complete explanatory chain from an abstraction of evolutionary and genetic processes to data from the domains of actual cultural communication.

After briefly discussing the content of the field of memetics and its development in section 2, I consider two types of memetic models that have arisen (section 3 and section 4), with examples examined in each. Then in section 5 I argue that the most basic selection criteria should include verification and validation, as a move towards a minimum standard for explanatory models. In section 6 I consider further selection criteria that might aid the judgement of models in complex situations, using a hard-nosed constructivist stance.

2 The Subject Matter of Memetics and its Development

The subject matter of memetics has developed at least partly as a matter of historical accident. Although the etymology of the term 'meme' goes back only to Dawkins in [6], the ideas can be traced back further [11].

Instead of working historically I wish to characterise the field in a functional way; namely:

> the application of models with an evolutionary or genetic *structure* to the *domain* of (cultural) information transmission.

Similar to other developing fields there is a natural progression of questions. An initial question might be *whether such an application is at all useful*. Following on from that is the question of *in what ways is such application useful*. Finally a mature field should be concerned with the question of determining *under exactly what circumstances is such an application useful*. At the moment, although a positive answer to the first question is frequently assumed by academics in the field, the battle for its *general* acceptance is far from won. Work has only just began to touch on the second question.

The field of memetics is a young one. At the moment there are many different memetics models of different kinds. A glance through this journal alone shows many different models of what a meme is or could be. This is healthy, it means that the mechanisms of variation are truly operative, but for the field to *evolve* what we also need good selection criteria for such models. While older fields have had time for appropriate selection mechanisms to be themselves evolved, a new field such as memetics is in the position that there is a great deal of uncertainty about what is a good memetic model.

In fields without good selection criteria that are grounded outside the field (e.g. evidence, the theories of other sciences etc.), there is a strong tendency to adopt a sort of default set of selection criteria, namely: the extent to which the model uses established language and techniques from the field and the extent to which the model is embedded in the field itself, i.e. it builds on other models, criticises other models, examines the thought of other models etc.[1].

Now, while such criteria are fine for maintaining the coherence of a field and for entrenching the position of those already *in* the field, it is not a good set of criteria for ensuring the long term growth and success *of* the field (which explains its widespread use in established but unsuccessful fields like economics). Such long-term success will only come about when people outside the field see the memetic models as useful to them and this will only happen quickly if they feel they can rely on them. The reliability of the models will depend on the extent that the field is grounded in something established (real-world studies and/or accepted models from outside the field). In other words the success of the field (as opposed to the success of academics within the field), will depend largely on the *embedding* of the field within the wider academic landscape. It is notable that there appears to be a strong correlation between the success of a field and its openness to ideas from outside. A susceptibility to Kuhnian paradigm shifts [14] seems to allow fields to adapt better to possible niches in the wider society.

Of course any model selection criteria must take into account the type of model. There are as many types of model as there are reasons for modelling. Two dimensions along which they differ are in their abstractness and their temporal direction. One model may be for the prediction of events (forwards), and another for explanation of events that have already occurred (backwards), some will be metaphors for use in thinking about situations (abstract)

[1] For a more formal definition of such embedding see [9].

and others will be bald data models of some aspect of the real world (concrete). In figure 1, I illustrate this categorization of some models on these two axes.

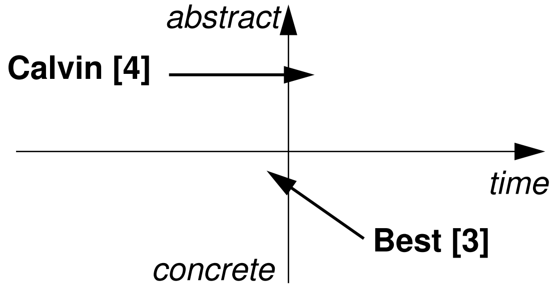

Figure 1. Some different types of memetic models, with their direction of explanation

3 *A Priori* Models

The first type of model I will discuss takes a foundationalist approach. That is it examines the conditions under which an cultural evolutionary process may occur. This is the approach taken by Dawkins in [7], where he states the thesis that in the presence of a new replicator an evolutionary process is “inevitable”. A good example of this type of model is Calvin [4]. Such works attempt to establish that if certain facts are true then a memetic process will (generally) occur. They thus attempt to build forward explanations in a manner that in terms of theory is top-down (i.e. it starting with abstract considerations and concluding with the phenomena of evolutionary process, see figure 1).

One of the weaknesses of such types of models, *when considered in isolation from others*, is that the link to real-world (concrete) models is weak. Typically in such works the examples to support such models are of a very high level involving such complex phenomena as *history* or *religion*. The complete demonstration that the abstractly motivated, *a priori* model would actually show that such complex phenomena are explained by memetic process can, at best, only be sketched and, at worst, is left solely to the imagination due to the fact that it is merely assumed.

This sort of assumption (that the complete explanatory chain *can* be built from such abstract considerations for a wide range of cultural phenomena) can be taken to ridiculous extremes, so that it may be taken *as read* that almost all cultural phenomena *are* due to memetic processes (I am *not* saying that Calvin is necessarily doing so). But the building of such a credible explanatory chain from theory to phenomena is not so easy – this (if and where it is possible) will be long and arduous. It is extremely tempting to attempt to short-cut the process.

One such temptation is to weaken the meaning of key terms in evolutionary models to *ensure* that the phenomena is a ‘memetic’ one. One such weakening is to call any process which includes *some* variation and selection an evolutionary one. This has the effect of losing some of the meaning of memetic models – after all, if one *did* succeed in choosing ones terms such that one could *prove* that many processes were memetic then one would have also proved that such a theory had no empirical foundation; it would be a merely formal proposition. Memetics could become like economics, where certain assumptions, processes and goals are taken as *a priori*[1] so that the gap between the theory and real-life becomes increasingly wide.

[1] Equivalently the field of economics can be defined by the study of such situations where they are true so that the field is no longer designated by the relevant subject matter itself.

If memetics is to be a real theory of cultural processes (as opposed to just one of many equivalent descriptive *frameworks*), there has to be some cultural process that is not designated *by the theory* to be of an evolutionary nature. Otherwise we would be left with non-falisfiable constructs that would not have any true grounding [19].

Let me propose one such process: *inference*. Here, one or more pieces of information act up on each other to form a new piece of information such that the *status* of this new information is such that it derives from the status of the *old* information (and the process by which it is combined). This is an example of variation that is *not blind*; the variant is produced so that we have some assurance that it will be selected *before* that selection process is actually applied to it. Such blindness is Heylighen's criteria for an evolutionary process [12]. The difference between an evolutionary process and, what I call an 'intentional' process is illustrated below in figure 2 and figure 3.

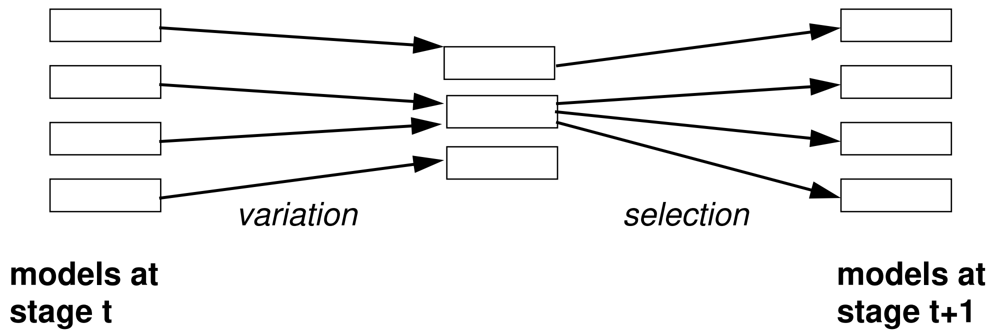

Figure 2. An evolutionary learning mechanism

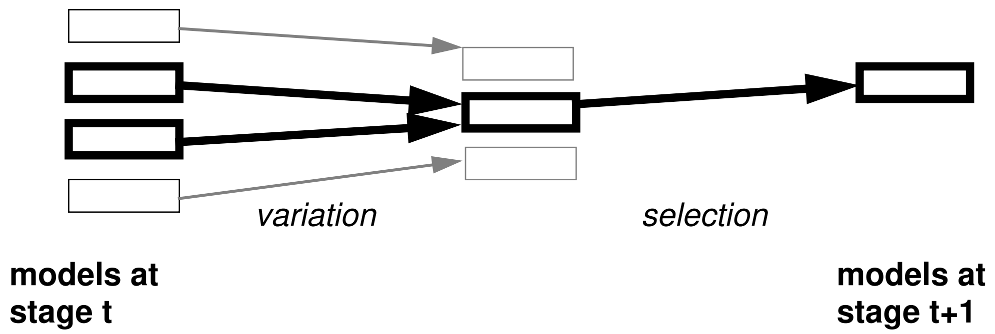

Figure 3. An 'intentional' learning mechanism

Thus such inference (if it occurs) is a good candidate for a process that may act upon information but is *not* a memetic one. Of course, many examples that are sensibly categorised as inference may be as a result of an evolutionary process. However, this does not change the fact that the presence of such a process *as a whol*e would weaken the claim that a *wider* process, of which it this was but a constituent, is essentially *not* an evolutionary one.

Unless we want to follow the sterile route of *a priori* economics then we have to ensure that the explanatory chain from abstract model to real-world relevance is really built rather than merely assumed. The a priori models should be allowed to survive only if such explanatory chains establishing their relevance are established.

4 'Black-Box' Models

The other end of the spectrum of is the 'black-box' model. Here some data from a phenomenon is analysed and a memetic model is fitted to it. So one can say that it is *as if* there were a memetic process at work. A good example of this approach is Best's [3]. Here the phenomena of the repetition of clusters of word in news groups is analyses and evidence is presented to show that some of these clusters are acting as if they are competing with each other as quasi-species. It is unfortunate (for us readers) that Best did not then go into a detailed case study to see the actual mechanisms that caused this apparent anti-correlation, as a biologist might when faced with an apparent predator-prey cycle in some population statistics. Instead he relied on reference to middle-range concepts such as the scarcity of resources. Thus his paper can be summarised as: an analysis of data *indicates* that an evolutionary process is taking place.

Of course the weakness of this type of model, *if it is taken in isolation*, is that one might be mistaken. A process may *appear* to be evolutionary but later turn out to be otherwise. If one took the step of defining evolutionary process by its (high-level) phenomena, then one would no longer a theory but just a description.

To illustrate this case I present a computational model, where a memetic process appears to be occurring, but isn't. It is analysed elsewhere in detail for some of its other properties in [9, 10]. It requires some detail of the model set-up to be understood for the relevance of the example to be become clear.

4.1 The Extended 'El Farol Bar' Model

This model is based upon Brian Arthur's 'El Farol Bar' model [2], but extended in several respects, principally by introducing learning and communication. There is a fixed population of agents (in this case 10). Each week each agent has to decide whether or not to go to El Farol's Bar on thursday night. Generally, it is advantageous for an agent to go unless it is too crowded, which it is if 67% or more of all the agents go (in this case 7 or more). This advantage is expressed as a utility, which each agent uses in the evaluations of their models of their world. Before making their decision agents have a chance to communicate with each other. This model can be seen as an extension of the work in [1], which investigates a three player game.

4.1.1 The environment

Each agent gets the most utility for going to the bar when less than 7 of the other agents go (0.7), they get a fixed utility (0.5) if they do not go and the lowest utility for going when it is crowded (0.4). In this way there is no fixed reward for any particular action because the utility gained from going depends on whether too many other agents also go. In this way there is no fixed goal for the agent's learning, but it is relative to the other agent's behaviour. Under this scheme it is in each agent's interest to discoordinate their action with a majority of the others.

4.1.2 The agents

Each agent has a population of (pairs of) expressions that represent possible behaviours in terms of what to say and what to do). This population is fixed in size but not in content. These expressions are taken from a strongly typed formal language which is specified by the programmer, but the expression can be of any structure and depth. Each agent does not 'know' the meaning or utility of any expression, it can only evaluate each expression as to the utility that would have resulted in if it had used it in the past to determine its action. Each week each agent takes the best such pair of expressions (in terms of its present evaluation against the recent past history) and uses them to determine its communication and action.

Each agent has a fairly small population of such models (in this case 30). This population of expressions is generated initially according to the specified language at random. In subsequent generations the population of expressions is developed by a genetic programming [12] algorithm with a lot of propagation and only a little cross-over.

The formal language that these expressions are examples of is quite expressive. The primitive nodes and terminals allowed are shown in figure 4. It includes: logical operators, arithmetic, stochastic elements, self-referential operations, listening operations, elements to copy the action of others, statistical summaries of past numbers attending, operations for looking back in time, comparisons and the quote operator.

possible nodes for talk gene:
greaterThan lessThan previous times plus minus divide averageOverLast quote saidByLast boundedByPopulation wentLag trendOverLast randomIntegerUpTo
possible terminals for talk gene: wentLastTime maxPopulation IPredictedLastWeek 'barGoer-1' 'barGoer-2' 'barGoer-3' 'barGoer-4' 'barGoer-5' 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

An example expression is shown in figure 5. In this example the agent would say "[IPredictedLastWeek]" if the attendance predicted by the trend in attendances over the last 2 weeks was greater than 5/3 and false otherwise, but it would only actually go if either it said it would or if barGoer-3 said it would.

```
talk:   [OR
            [greaterThan
                [trendOverLast [2]]
                [divide [5] [3]]
            [quote [IPredictedLastWeek]]]
action: [OR
            [saidBy ['barGoer-3']]
            [ISaidYesterday]]
```

Figure 5. An example model

A second example (figure 6) shows the operation of a memetic process. In this example the agent says whatever the expression last uttered by barGoer-5 evaluates to, but goes if it went last time and a random coin-flip turns up heads.

```
talk:     [saidByLast ['barGoer-5']]]
action: [AND
              [randomDecision]
              [IWentLastWeek]
```

Figure 6. A second example model

The primitives are typed (boolean, name or number) so that the algorithm is strictly a strongly-typed genetic program following [15].

4.1.3 Communication

Each agent can communicate with any of the others once a week, immediately before they all decide whether to go to the bar or not. The communication is determined by the evaluation of the talk expression and is usually either 'true' or 'false'. The presence of a quoting operator (quote) in the formal language of the talk expression allows subtrees of the talk expression to be the content of the message. If a quote node is reached in the evaluation of the talk expression then the contents of the subtree are passed down verbatim rather than evaluated. If a quoted tree is returned as the result of an evaluation of the talk expression then this is the message that is communicated. A section of the transcript from the model is shown in figure 7, which shows some of the messages sent.

```
barGoer-2 says true
barGoer-3 says true
barGoer-6 says [AND [AND [AND [IPredictedLastWeek] [IPredictedLastWeek]] [greaterThan
[wentLastTime] [4]]] [quote [quote [lessThan [wentLastTime] [5]]]]]
barGoer-7 says false
barGoer-4 says false
barGoer-9 says true
barGoer-1 says false
barGoer-10 says true
barGoer-8 says [NOT [OR [IPredictedLastWeek] [IPredictedLastWeek]]]
barGoer-5 says true
```

Figure 7. Part of the model transcript showing some messages (week 14 of the run)

The content of the messages can be used by agents via saidBy and saidByLast nodes in the action and talk expressions. If the message is just composed of a boolean value then the saidBy node is just evaluated as this value, but if it is a more complex expression (as a result of a quote node in the sending agents talk expression) then the whole expression will be substituted instead of the saidBy (or saidByLast) node and evaluated as such. The agent can use the output of its own messages by use of other nodes (IPredictedLastWeek and ISaidYesterDay).

The agents are limited in those from whom they can 'listen to', in the sense that at the beginning of the run an arbitrary 'friendship structure' (as in) is imposed and agents can only use saidBy (or saidByLast) nodes referring to agents who are their friends in their models.

In this model 'social imitation' is also enabled. this means that other agents can introduce the content of any message (which is not a mere boolean value) into their own (action) model pool, this would correspond to agents taking the message as a suggestion for an expression to determine their own action. In subsequent generation this expression can be crossed with other expressions in its population of constructs so that its content is mixed in with other *action* models of that agent. Note that social imitation does not incorporate the content of such messages into the models that determine what is said by each agent.

4.1.4 Implementation

The model was implemented in a language called SDML (strictly declarative modelling language), which has been developed at the Centre for Policy Modelling specifically for social modelling [18].

4.1.5 Results, analysis and discussion

If one looks at the occurrence of different types of utterance by the agent and graphs the frequency of different message sub-trees it looks distinctly as if a memetic process is taking place. Consider, for example the occurrence of "IPredictedLastWeek" over the course of the first 40 week of the simulation (figure 8).

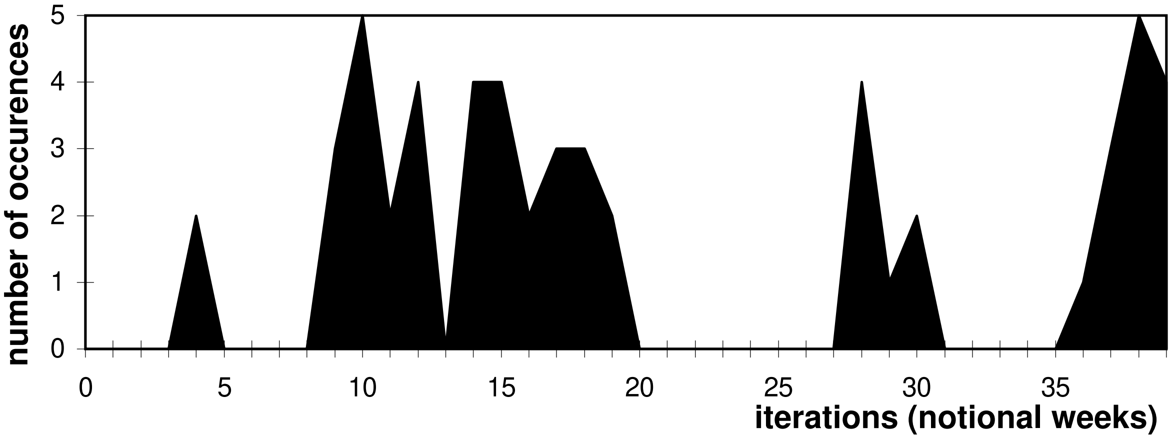

Figure 8. Number of occurrences of "IPredictedLastWeek" in agent utterances

The occurrences occur in distinct ‘bursts’. Futhermore it not restricted to only a few agents. the distribution of who says “IPredictedLastWeek”, when is shown below in figure 9 (Agent-3 has not been included since it never utters this node).

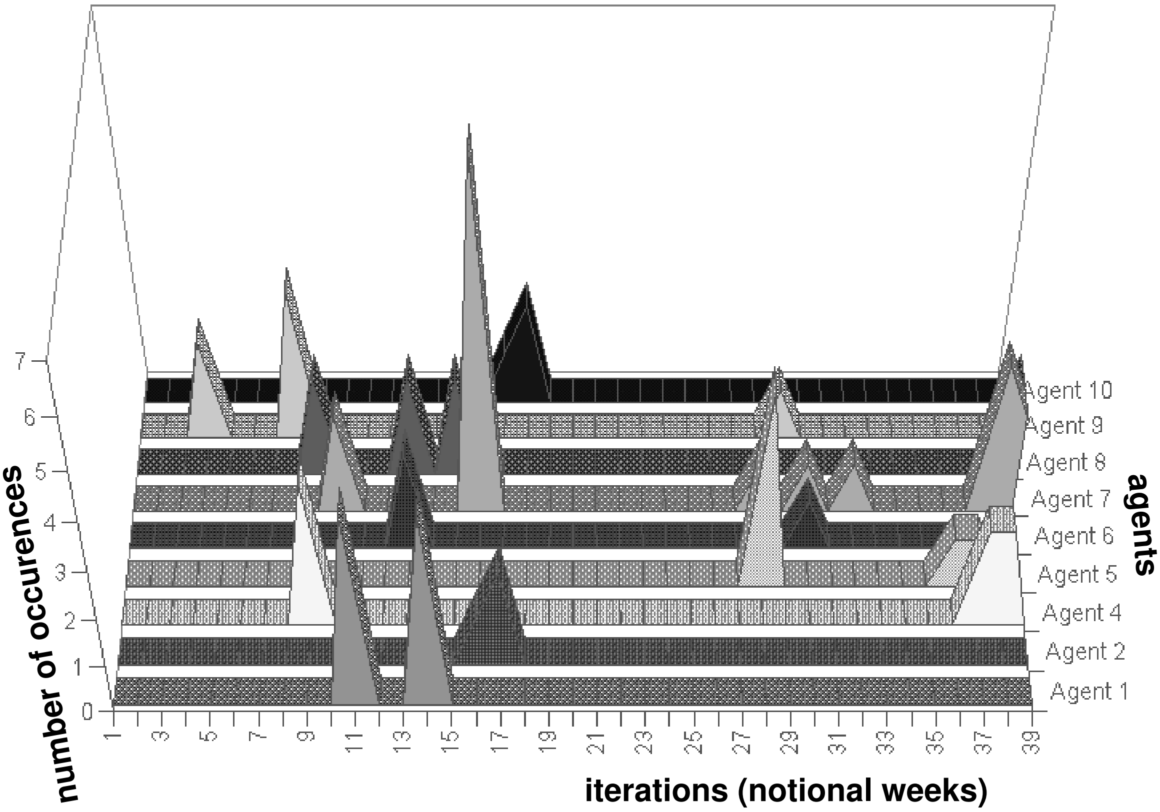

Figure 9. The distribution of utterances of “IPredictedLastWeek” by agent

However, if one analyses the transcripts of the simulation in detail, then one gets a different picture. In figure 10 I exhibit a similar graph to that of figure 8, but separated into those utterances of “IPredictedLastWeek” that were as a result of a saidByLast node in the chosen talk expression by an agent (i.e. those messages that were memetically reproduced) and the rest. These other utterances arose because they were *simultaneously* determined as relevant by the learning of each agent. What memetic imitation there was is attributable to chance copying due to the presence of saidByLast nodes in the talk expressions. These utterances with IPredictedLastWeek were not *selected for* by these agents. In particular one does not find branching chains supporting the presence of memetic copying of messages. In fact there is only *one instance* of such a message being passed down the generations for more than one notional week (this occurs in the middle of the simulation, barGoer-3 at week 16 is copied by barGoer-2 at week 17 which is copied by barGoer-7 at week 18 and then by barGoer-8 at week 19), and even then there is no *replication* of the meme into more than one strand.

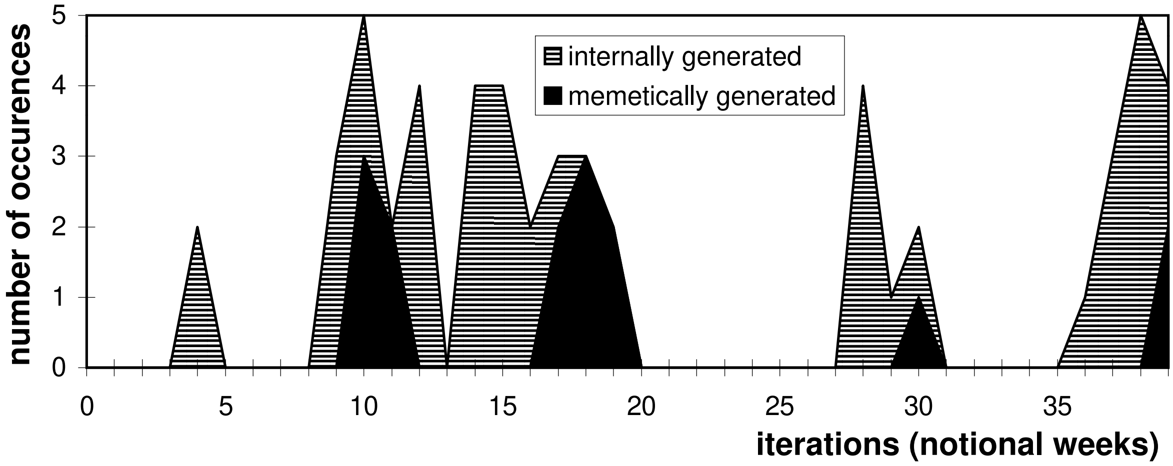

Figure 10. Number of occurrences of "IPredictedLastWeek" in agent utterances analysed by generating process

The fact that the learning processes in this simulation are achieved via an evolutionary mechanisms is beside the point. The fact that inside each agent there is an evolutionary process (by explicit design) does not make the process of messages between agents into a memetic one. A memetic process could have co-evolved if the agents evolved models to do so but this did not occur in this run of the simulation. The agents did copy from each other but not so as these copying processes where chained into a memetic process so that we could meaningfully say the messages themselves underwent an evolutionary process.

If this sort of pseudo-memetic process can occur in an artificial simulation where mechanism are explicitly built in a mechanism for the expression, propagation and selection of memes by agents, then it is possible that similar processes could be happening elsewhere. Thus it is wrong to imply the existence of memetic process, purely as a result of a post-hoc analysis of data and the mere presence of possible mechanisms. Such 'black-box' models are insufficient on their own.

5 Basic Selection Criteria for Memetic Models

It should be clear that what is needed is for such models to be chained into a complete explanation so that the presence of evolutionary mechanisms can be show to *cause* the evolutionary phenomena that results. This would then justify the models that composed that chain. Note that this chain can (and probably will) be composed of more than one model. Even if one did have one model which seems to bridge the gap there would always be, at the top, the metaphor of evolution (that I talked about in section 1) that it must relate to in order to justify its categorisation as a memetic model and the data model at the bottom.

It would be silly to insist that every paper must, by itself, bridge this gap, but what each model should (at a minimum) do is say how it explicitly relates to the models above and below itself in the chain. In other words, each model should be well *verified* and *validated*, as Scott Moss and I suggest for multi-agent models [17]. The validation of a model is where the structure of the model is based on the *results* of another model (even if it is fairly abstract like the idea of evolution). The verification of a model is where the results of the model are checked against something more definite (either the structure of other models that are themselves verified or a data model).

To make this clear let me take an example from a more established science: chemistry. Physics has models of the processes of elementary particle interaction[1]; chemistry builds on this by positing mechanisms of how chemicals interact that use these physical processes (i.e validated) and verify these models against data models of real processes. Sometimes they then use these basic models of how chemicals interact in simulations which involve many chemicals when analytically solving to predict their interaction is too difficult. They may then make approximate numerical models of the results of these simulation and verify these against data models of real processes. The whole chain is illustrated in figure 11.

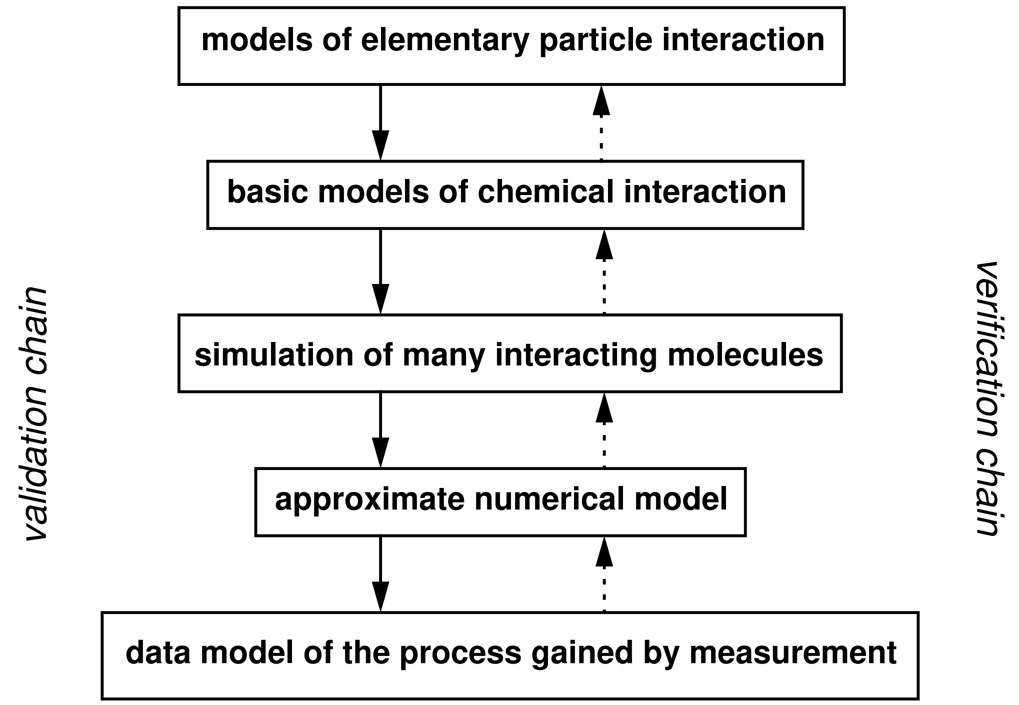

Figure 11. Validation and verification between layers of models in chemistry

Of course, beyond the presence of the validation and verification links, the *strength* of those links is also important. One method for helping to establish strong links is the use of formal models. Even if these formal models are not used to infer anything new they can be immensely useful in pinning down the meaning of terms and so reduce the potential for ambiguity and confusion. It is the bane of the social sciences, that everyone has slightly different meanings for key terms and concepts. A formal model allows the clarification of these meanings, so that the exact relationship between different models can be established. Another way of looking at this is that the links are themselves a sort of model, that are verified against the structure of the model they link: formal models thus allow the validation of these links against formal processes.

My characterisation of memetics in section 2, allows the categorisation a model as *memetic*, by the ultimate validation of its structure of against the evolutionary/genetic metaphor and the

[1] Some of the models in physics are of a slightly different character, because they form the ultimate top-level model for many such chains. Thus despite the fact that these models are often taken as the archetypal model for science they are, in fact quite special cases.

its ultimate verification against data from a process of cultural communication. These processes can be indirect using other intermediate models, as discussed above.

I wish to make clear that I am *not* criticising Calvin, Best or their work. A field will and must have specialists in different sorts of model. However the utility of their work *will ultimately depend* upon whether their models are included in a complete explanatory chain from theory to phenomena. Since memetics has only just began as a serious enterprise it is unsurprising that there are not other convenient intermediate models with which to build these links, but this does not relieve us of the obligation to do so – a situation like economics where the excuse "its still a young science" is used 300 years after its inception should not be allowed. All models that are not part of a full explanatory chain should be mentally tagged as *provisional* and (at least temporarily) discarded if it is not justified in the *near* future.

This is, of course, a sort of bootstrapping process. In order to progress into a "normal science" (in Kuhn's terminology [14]) then we need to establish some complete explanatory chains as quickly as possible. To this end it would help if people in the field concentrated on domains where both the ultimate validation against the evolutionary metaphor and the ultimate verification against evidence of memes was as straight-forward as possible. Thus (at least until the shape and limitations of memetics becomes clearer) it would be far more helpful for studies in domains such as: bird song, nursery rhymes, tunes, and legal phrases in contracts where the mapping between the form and meaning of potential memes is clearly evident. Domains which involve very high level social constructs such as *history* or *religion* should be generally avoided since the construction of a convincing explanatory chain is the most onerous, and hence most suspect if hinted at as support for a model.

6 Further Methods to Ensure Rigorous Selection

Validation and verification are the basic requirements for a model. In themselves they are only sufficient if one is modelling relatively simple processes, where it is possible to effectively and clearly test models by carefully designed experiment. In modelling cultural transmission, we are not (on the whole) so lucky. In approaching our subject matter we are forced to make many choices regarding the selection of data, the modelling framework, modelling goals, etc. that there is a danger of these two selection criteria being insufficient to ensure the continued evolution of our models. It is well known that in evolutionary processes if the mutation rate is too high compared to the rest of the process then this can swamp the process so that it does not learn to successfully adapt; this may occur to the field of memetics – if we let it.

Hence for memetics (and other similar fields), the difficulties of verification mean that we need more stringent selection criteria. We may have to accept that we are *constructing* our memetic reality rather than merely passively *reflecting* it by our modelling choices, since the complexity of our chosen phenomena forces this on us. But just because we take such a constructivist stance[1] we do not have to be vague – we can respond to our chosen challenge by seeking to retain rigour.

These other criteria can be such as: its similarity to other accepted models (a lateral coherence criteria); the complexity of the model; the meaningfulness of the model in human terms; the ease with which the model can be used to calculate its predictions; and how general (or generalisable) it is.

[1] For a good introduction to constructivism and the importance of selection criteria see the pages at the Principia Cybernetica Project at URL: http://pespmc1.vub.ac.be/construc.html

Let me stress again that these are not optional extras – by applying these and making them explicit in our work we are only accepting their inevitability. When dealing with the complexity of phenomena such as social interaction we have no choice but to accept trade-offs in terms of the complexity of our models and (for example) their error-rate when verified against data [8].

7 Conclusion

The complexity of social interaction and the difficulty in directly falsifying memetic models means that unless we apply strong selection criteria the field of memetics will be swamped by the forces of variation and hence not successfully evolve. As a result of this memetics as a *field* may well lose the race to other fields in academia and society. For this reason the application of stringent selection processes (that are relevant to the survival of the field as opposed to relevant to only the survival of academics and their thought) is vital.

This can be done in a number of ways:

1. (Use of formal models) insisting on formal or computational models to pin down the meaning of our theories. This makes the relationships between our theories easier to assess since the formal models will help specify the reference of terms in an unambiguous way;
2. (Validation) make clear the assumptions behind our models in terms of the modelling language, validating theories, modelling goals etc.;
3. (Verification) explicitly state how we are verifying our model, including: how we have selected the model or information to verify against and the precise method of this verification;
4. (Other criteria) in addition to the above, argue for our models in terms of additional criteria such as: simplicity; ease of computation of predictions; etc.
5. (Domains) initially choosing domains where the meaning of our models are clear – i.e. a concentration of phenomena where the identity of memes is transparent due to the clear mapping between their forms and their content (e.g. birdsong, tunes, legal phrases in contracts, nursery rhymes) and treating examples concerning high-level constructs (e.g. religion, history, organisational behaviour) with great care.

Acknowledgements

My thanks to Mark Mills, for his kind comments on a previous paper [9], which provided the impetus for this paper, to Scott Moss for many discussions about the process of modelling, and Steve Wallis for writing SDML.

SDML has been developed in VisualWorks 2.5.1, the Smalltalk-80 environment produced by ObjectShare (formerly ParcPlace-Digitalk). Free distribution of SDML for use in academic research is made possible by the sponsorship of ObjectShare (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] Best, M., L. (1997); Models for Interacting Populations of Memes: Competition and Niche Behavior. *Journal of Memetics - Evolutionary Models of Information Transmission*, **1**. http://www.cpm.mmu.ac.uk/jom-emit/1997/vol1/best_ml.html

[4] Calvin, W.H., 1997; The Six Essentials? Minimal Requirements for the Darwinian Bootstrapping of Quality. *Journal of Memetics - Evolutionary Models of Information Transmission*, **1**. http://www.cpm.mmu.ac.uk/jom-emit/1997/vol1/calvin_wh.html

[5] Campbell, D. T. (1960). Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes. *The Psychological Review*, **67**: 380--300.

[6] Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.

[7] Dawkins, R. (1982). *The extended phenotype*. Oxford University Press.

[8] Edmonds, B. (1997). Complexity and Scientific Modelling. 20th International Wittgenstein Symposium, Kirchberg am Wechsel, Austria, August 1997. Also available as a CPM Report 97-23, MMU, Manchester, UK.

[9] Edmonds. B. (1998). Capturing Social Embeddedness: a constructivist approach. CPM Report 98-34, MMU, Manchester, UK.

[10] Edmonds, B. (forthcoming). Modelling Socially Intelligent Agents. *Applied Artificial Intelligence*.

[11] Gatherer, D. (1997). Macromemetics: Towards a Framework for the Re-unification of Philosophy. *Journal of Memetics - Evolutionary Models of Information Transmission*, **1**. http://www.cpm.mmu.ac.uk/jom-emit/1997/vol1/gatherer_dg.html

[12] Heylighen F. (1992): Principles of Systems and Cybernetics: an evolutionary perspective, in *Cybernetics and Systems '92*, R. Trappl (ed.), Singapore: World Science, 3-10.

[13] Koza, J.R. (1992). *Genetic Programming: On the programming of computers by means of natural selection*. Cambridge, MA: MIT Press.

[14] Kuhn, T.S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

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

[16] Moss, S.J. and Edmonds, B. (1998). Modelling Economic Learning as Modelling. *Systems and Cybernetics*, **29**:5-37.

[17] Moss, S. J., Edmonds, B. and Wallis, S. (1997). Validation and Verification of Computational Models with Multiple Cognitive Agents. CPM Report 97-25, MMU, 1997.

[18] Moss, S. J., H. Gaylard, S. Wallis, and B. Edmonds, (1998). SDML: A Multi-Agent Language for Organizational Modelling. *Computational and Mathematical Organization Theory*, **4**:43-69

[19] Popper, K. R. (1968). *The Logic of Scientific Discover*y. London: Hutchinson.